



NotPlaNET: Removing False Positives from Planet Hunters TESS with Machine Learning

Valentina Tardugno Poleo¹ , Nora Eisner² , and David W. Hogg^{1,2,3} ¹ Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA; vt2189@nyu.edu² Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA³ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

Received 2024 March 27; revised 2024 May 24; accepted 2024 July 2; published 2024 August 1

Abstract

Differentiating between real transit events and false-positive signals in photometric time-series data is a bottleneck in the identification of transiting exoplanets, particularly long-period planets. This differentiation typically requires visual inspection of a large number of transit-like signals to rule out instrumental and astrophysical false positives that mimic planetary transit signals. We build a one-dimensional convolutional neural network (CNN) to separate eclipsing binaries and other false positives from potential planet candidates, reducing the number of light curves that require human vetting. Our CNN is trained using the TESS light curves that were identified by Planet Hunters citizen scientists as likely containing a transit. We also include the background flux and centroid information. The light curves are visually inspected and labeled by project scientists and are minimally preprocessed, with only normalization and data augmentation taking place before training. The median percentage of contaminants flagged across the test sectors is 18% with a maximum of 37% and a minimum of 10%. Our model keeps 100% of the planets for 16 of the 18 test sectors, while incorrectly flagging one planet candidate (0.3%) for one sector and two (0.6%) for the remaining sector. Our method shows potential to reduce the number of light curves requiring manual vetting by up to a third with minimal misclassification of planet candidates.

Unified Astronomy Thesaurus concepts: [Exoplanets \(498\)](#); [Exoplanet detection methods \(489\)](#); [Convolutional neural networks \(1938\)](#); [Light curves \(918\)](#); [Light curve classification \(1954\)](#)

1. Introduction

There have been significant advancements in space-based, precise time-series measurements, particularly with missions such as Kepler (Borucki et al. 2010) and the Transiting Exoplanet Survey Satellite (TESS; Ricker et al. 2014). These advances have led to large numbers of exoplanet discoveries. To effectively parse through the growing volume of data, the field requires innovative analysis techniques.

Despite the growing selection of transit-finding and classifier models, the process of identifying transiting exoplanets still relies heavily on human vetting. One of the most successful methods used to identify periodic, transit-like signals in photometric data is the box least squares (Kovács et al. 2002) algorithm. However, this method typically produces more candidates than true planets, with a large fraction of the identified signals being false positives that mimic true transits. These false-positive scenarios include eclipsing binaries and systematic variations caused by the instrumentation, cosmic rays, or temperature fluctuations. Typically, teams of experts examine possible transit signals and assess their disposition (e.g., Crossfield et al. 2018; Eisner et al. 2020; Guerrero et al. 2021). This process is very time consuming, particularly if each candidate is inspected by multiple experts. Other vetting methods consist of fitting physical and probabilistic models to light curves. Foreman-Mackey et al. (2016) present a fully automated model that operates without human intervention. Yet the described method struggles to identify long-period planets orbiting variable stars.

Convolutional neural networks (CNNs; O’Shea & Nash 2015), widely used machine learning tools for image pattern recognition, have also been applied to help the classification process. Progress has been made in employing CNNs and other machine learning techniques to identify light curves that contain transit-like signals as well as to classify these signals as planet candidates or false positives (Shallue & Vanderburg 2018; Dattilo et al. 2019; Yu et al. 2019; Malik et al. 2022; Tey et al. 2023). Nevertheless, human vetting still significantly outperforms current models when identifying long-period, single-transit, and aperiodic events (e.g., Fischer et al. 2012; Schwamb et al. 2013; Wang et al. 2013).

Typically, machine learning classification methods rely on the use of phase-folded time-series data and thus require prior knowledge of the periodicity of the signal (e.g., Salinas et al. 2023). The requirement of two transit events for phase folding also restricts their use to multitransit planets and thus shorter-period planets. Furthermore, most of these methods do not consider background or centroid information, which may help identify background events and eclipsing binaries. Other CNNs make use of heavily preprocessed data. A more flexible model that does not require phase folding or extensive preprocessing, and that considers background and centroid information could thus complement existing models.

We built and trained a one-dimensional CNN to remove as many false-positive signals as possible without discarding planet candidates. In contrast to previous work, (e.g., Zucker & Gyires 2018; Olmschenk et al. 2021; Salinas et al. 2023), our model focuses on contaminant detection and does not rely upon the use of known orbital periods. The open-source CNN, NotPlaNET,⁴ was trained using labels provided by the Planet



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.

⁴ <https://github.com/vtardugno/TESS-CNN>

Hunters TESS citizen science project. The labels consist of times within light curves where a transit-like signal has been flagged, and the associated label of “likely planet candidate” (PC), “eclipsing binary” (EB), or “other signal” (OS). Since the CNN was purposely designed to work with data that has not been phase folded, NotPlanET allows us to vet single-event signals and help identify single-transit, longer-period planets. The vetting and identification of single-transit signals is particularly important in the era of TESS, where the majority of planets with periods longer than 27 days will not exhibit consecutive transits in the TESS data.

The paper is structured as follows: Section 2 describes the data and preprocessing that was carried out. Section 3 provides an overview of the neural network’s architecture. The contaminant selection procedure and the results are outlined in Section 4. Section 5 provides a discussion of the results, and, lastly, Section 6 concludes with a summary of our results and directions for future research.

2. Data, Labels, and Preprocessing

The data were composed of the 2 minute cadence TESS Simple Aperture Photometry (SAP) and Pre-search Data Conditioning SAP light curves, which are produced by the TESS pipeline at the Science Processing Operations Center (Jenkins et al. 2016) and made publicly available by the Mikulski Archive for Space Telescopes (MAST). Each sector observed by TESS measures $24^\circ \times 96^\circ$ oriented along the line of ecliptic longitude and consists of ~ 27 days of photometric observations. The data also included eclipsing binaries from the Prša et al. (2022) TESS Eclipsing Binary Stars catalog and TESS objects of interest (TOI; Guerrero et al. 2021). We used data from sectors 2–6, 10–11, 17–19, and 32–60. Sectors ranging from 2 to 35 were used to train the network. The training sectors were randomly selected from the earliest 35 sectors to assess whether a model trained on early sectors could have predictive power for unseen newer sectors. We left gaps in the selected sectors to ensure our model could generalize to unseen sectors. Sectors 36–42 were used to optimize our model’s hyperparameters and fine-tune our final classification scheme. Sectors 43–60 were left unused until the very end of the evaluating phase so that we could assess how well our model performed on data it had not been trained or validated on.

2.1. Labels

We make use of the light curve labels from the Planet Hunters TESS citizen science project (for details see Eisner et al. 2021). In brief, the online project presents volunteers with individual TESS light curves and asks them to identify the times of possible transit-like signals using their mouse button. Each light curve is seen by 15 volunteers, who independently identify transit signals. For each light curve, the markings from the individual volunteers are combined using a density-based clustering algorithm, which allows us to identify light curves where multiple volunteers identified the same signal. Taking into consideration the number of volunteers who identified a given signal and their user weights (see Eisner et al. 2021) we calculate a transit score for each TESS light curve for a given sector (a higher score indicates a higher confidence in the signal being real).

For each TESS sector, the light curves with the 500 highest ranked scores are visually inspected by three members of the Planet Hunters TESS science team to differentiate between possible PCs, EBs, and OSs. This manual vetting of the light curves has been completed for sectors 1–65.

It is important to note that while we relied on the scientists’ labels as ground truth during the training and validation of our model, there remains a degree of uncertainty in these labels. In particular, the scientists are encouraged to vet light curves “optimistically,” meaning that if the nature of the signal is ambiguous, it should be classified as a PC. Identifying ambiguous signals as potential planets allows for these light curves to be inspected in the next round of vetting, minimizing the loss of potential candidates, but this can lead to the mislabeling of several targets.

Notably, while some of the light curves and their labels were taken from the TESS EB catalog and TOIs, the exact locations of the transits were derived from the volunteer markings in Planet Hunters TESS. This ensured that the methodology for determining signal locations was consistent across planets, EBs, and other signals, with all signal locations being determined via Planet Hunters TESS.

2.2. Preprocessing

All of the TESS light curves with a PC, EB, or OS label were first normalized and then split into smaller chunks for data augmentation purposes as well as to ensure translational invariance. The light curves and accompanying data, except for the background flux, were normalized by dividing by the median value of each array. Conversely, the background flux was normalized by applying a median filter to correct for the time-variable scattered light from the Earth and the Moon. Then, all the data were parsed using a stride of 800 data points and chunked into 2 days fragments. The choice of a 2 days fragment size stems from our main objective: to avoid discarding planets, especially single-transit and long-period ones. Given that such planets may have orbital periods spanning several days, we opted for a fragment size capable of capturing these extended transit events.

Gaps in the light curves were handled by appending the data as if the gaps did not exist. This approach was chosen because our focus is on single events within narrow windows, rather than periodic signals, making the influence of these gaps minimal on our training process. Interpolating over large data gaps was avoided to prevent introducing bias and distorting the signal characteristics. It remains an open question how to best handle data gaps, and further research is required to explore techniques to best address the problem of missing data.

To assign the labels to each chunk, we examined whether each fragment contained a citizen scientist marking. If so, the fragment was assigned the label given during manual vetting. If not, the fragment was assigned an OS label. To create training and validation sets, the data were shuffled and split using sklearn’s (Pedregosa et al. 2011) `train_test_split` function with its default settings.

3. Neural Network Architecture

The CNN was built with PyTorch (Paszke et al. 2019) and consists of six 1D convolutional blocks. Each convolutional block has two 1D convolutions with a kernel size of 5, followed by the ReLU activation function, layer normalization,

Table 1
NotPlaNET Convolutional Neural Network Parameters

Layer	Output Size (LxC)	Activation	Normalization	Pooling	Dropout
Conv1_1/Conv1_2	1440×32	ReLU	Layer Norm	Max Pooling	10%
Conv2_1/Conv2_2	720×64	ReLU	Layer Norm	Max Pooling	10%
Conv3_1/Conv3_2	360×128	ReLU	Layer Norm	Max Pooling	10%
Conv4_1C/conv4_2	180×256	ReLU	Layer Norm	Max Pooling	10%
Conv5_1/Conv5_2	90×64	ReLU	Layer Norm	Max Pooling	10%
Conv6_1	45×32	ReLU
Conv6_2	45×16	ReLU	Layer Norm	Max Pooling	...
Flatten	16×1
Linear	16×3	Softmax

Note. All values were selected through hyperparameter optimization by measuring which configurations yielded the lowest validation loss.




Figure 1. Convolutional block diagram. The block consists of two 1D convolutional layers with the ReLU activation function, layer normalization, and max pooling.

and max pooling. Lastly, we added one linear layer with the Softmax activation function. Figure 1 illustrates a single convolutional block, and Table 1 describes the entire architecture. The architecture was chosen using hyperparameter optimization with the Weights and Biases development platform Biewald (2020). The network’s inputs were the light curve chunks along with their background fluxes and centroids, and the outputs consisted of three numbers which were scores assigned to each category.

The Adam optimizer (Kingma & Ba 2014) with a learning rate of 0.0005 was used to minimize the cross-entropy loss. Due to the imbalanced nature of the data, with EBs being the most represented, we assigned loss weights to each category. A learning rate scheduler was set to decrease the learning rate by a factor of 0.8 every 50 epochs and the network was trained for 200 epochs. During every training loop, 10% of the nodes at the end of each convolutional block were randomly zeroed, leading to a different configuration after every epoch. Dropout regularization helps the network learn patterns in the data rather than memorizing the training set, avoiding overfitting. In every epoch the training and validation loss were calculated and the weights associated with the lowest validation loss were stored as our final model parameters.

4. Classification Methods and Results

We devised a selection process that used the CNN scores to determine whether a given light curve should be flagged as a




Figure 2. Fraction of planet candidates (top) and true contaminants (bottom) flagged as contaminants as a function of the PC score threshold for the seven validation sectors. For each validation sector, we chose a threshold that maximized the fraction of discarded contaminants while discarding zero planets. We then chose our final threshold (dashed line) by performing a weighted average of each validation sector’s threshold. The weights were assigned as the inverse square of the number of contaminants found in a given sector, allowing for a more conservative cutoff.

contaminant. We aimed to find a threshold that found as many contaminants as possible while removing no planet candidates in the validation sectors. For each validation sector, we probed CNN prediction thresholds in all three categories and found that the best results were achieved when placing a threshold on the PC score. Figure 2 illustrates the effect of different PC thresholds on the fraction of planet candidates and true contaminants discarded.

The threshold for each validation sector was chosen so that no planet candidates were discarded, and the number of true contaminants found was maximized. We then found the average of the sectors’ thresholds, weighted by the inverse square of the number of contaminants found per sector. The weight equation was chosen empirically, using the validation sectors to identify what final threshold resulted in the highest combined fraction of true contaminants found while imposing the condition that no planet candidate is incorrectly classified. We suspect the reason behind the success of this weighting scheme was that it resulted in lower weights for thresholds that

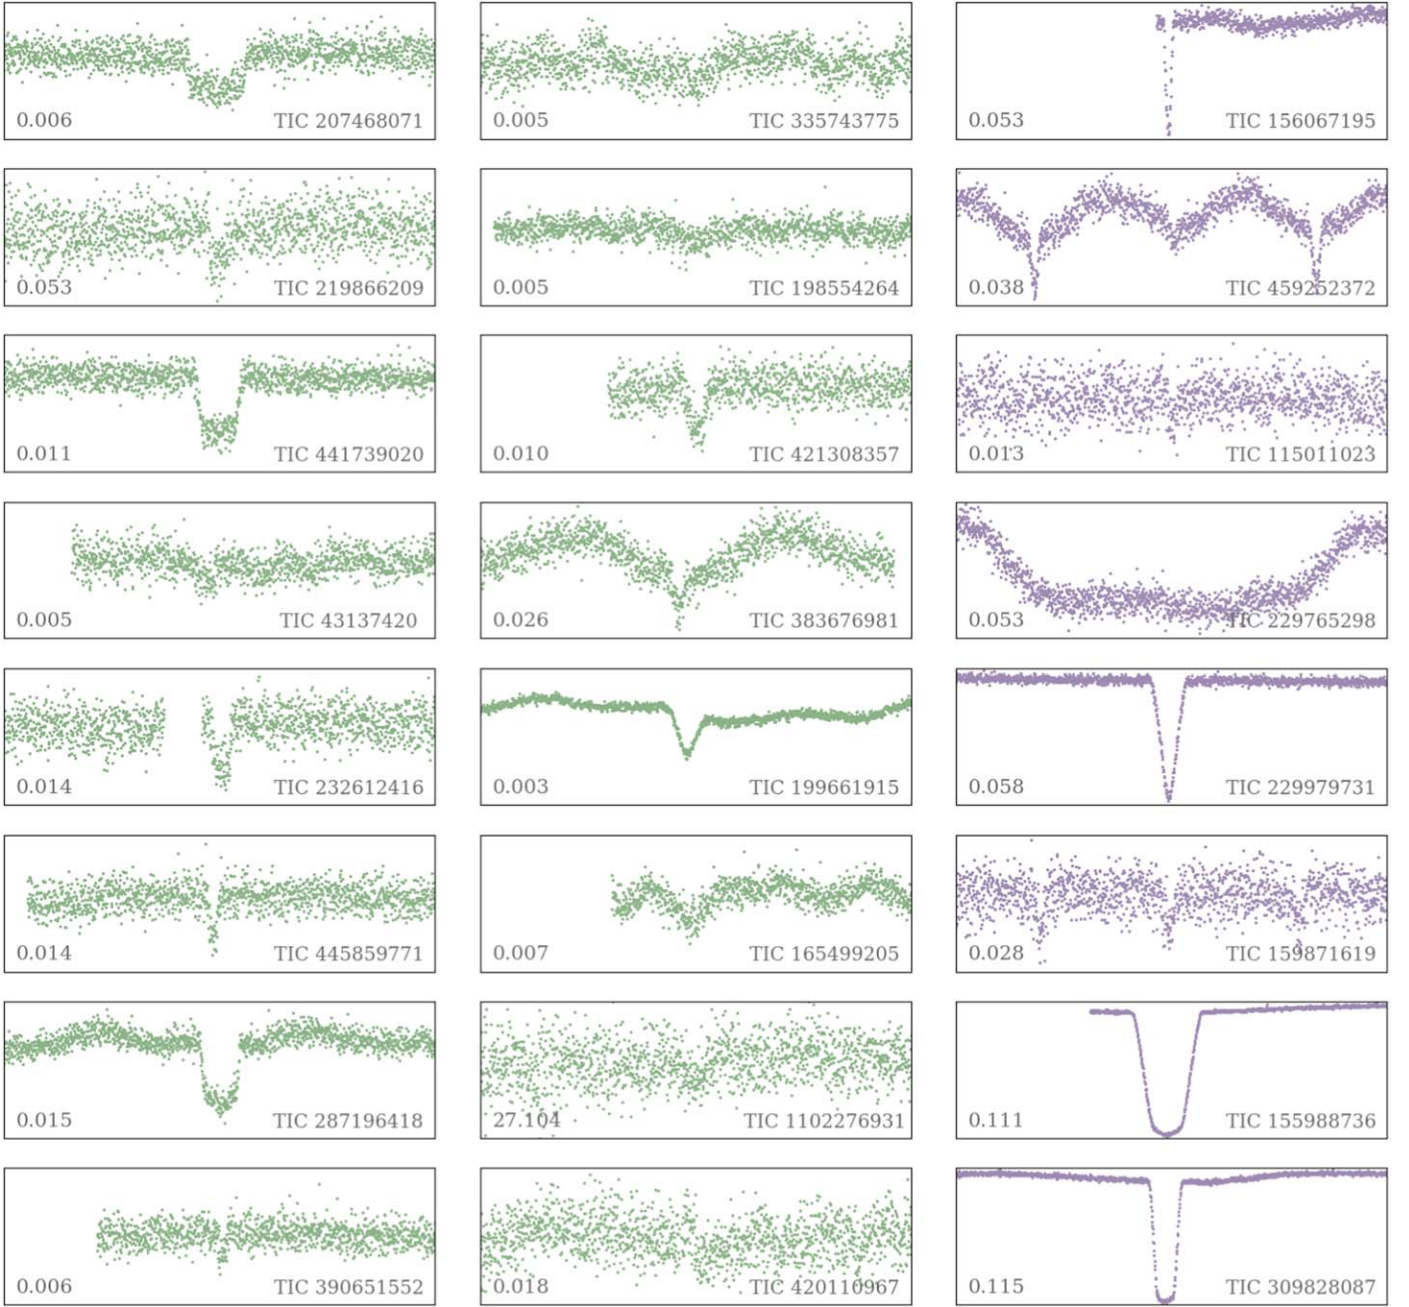


Figure 3. Selected light curve chunks from sector 50. The rightmost column (purple) contains the light curves flagged as contaminants by our model. The leftmost and middle columns (green) contain light curves classified as “keep for further vetting.” The left column shows true planet candidates, while the middle one shows true contaminants that were incorrectly classified as “keep for further vetting.” The y-axis range is displayed in the lower-left corner of each panel, while the lower-right corner displays the TIC ID of the light curve.

discarded more contaminants, yielding a more conservative cutoff.

Before making a classification, we looked at all the scores assigned to each chunk of a given TIC ID and sector. Light curves with a PC score lower than the determined threshold in all chunks were labeled as contaminants. As long as at least one chunk in the light curve had a PC score greater than the threshold, the light curve was assigned to the “keep for further inspection” category. Figure 3 displays selected light curve chunks for the “contaminants” and “keep for further vetting” categories. Although we selected a very conservative threshold, the model can easily be modified to more aggressively reduce

contamination at the expense of a few missed planet candidates.

We applied NotPlaNET along with our chosen classification method and threshold on the test set, containing data from sectors that had not been previously used for training, threshold selection, or any other decision-making steps. The classification results are summarized in Table 2.

For our test set, the percentage of EBs in the found contaminants had a median value of 82% across all test sectors with a maximum of 100% and a minimum of 68%. However, the median, maximum, and minimum percentages of EBs discarded out of the total number of EBs per sector were 23%,

Table 2
Light Curves Flagged as Contaminants by Our Model

Sec	PC	EB+OS (%)	Sec	PC	EB+OS (%)
43	0	68 (11%)	52	0	128 (18%)
44	0	115 (17%)	53	2	162 (20%)
45	0	73 (17%)	54	0	10 (20%)
46	0	15 (20%)	55	0	46 (14%)
47	0	125 (16%)	56	0	175 (22%)
48	0	124 (16%)	57	0	113 (22%)
49	0	92 (12%)	58	1	288 (36%)
50	0	82 (10%)	59	0	89 (19%)
51	0	43 (17%)	60	0	55 (15%)

Note. The planets wrongly classified as contaminants are listed on the PC column. The EB+OS column lists all the true contaminants discarded in each sector along with the percentage of total contaminants discarded.

44%, and 16%, respectively. While our model cannot yet find a high fraction of the EBs in the data, a high fraction of the flagged contaminants we found were EBs. In other words, the contaminant category has high precision and low recall rates when looking for eclipsing binaries. We suspect our model could be modified and applied to projects focusing on eclipsing binary stars. Moreover, we found no significant difference in the morphology distributions of eclipsing binaries that were labeled as contaminants and the ones that were kept for further vetting.

5. Discussion

NotPlaNET demonstrated the potential to identify contaminants while retaining nearly all potential planet candidates. Our model correctly identified 10% to 37% of the true contaminants, with a median identification rate of 18% across all 18 test sectors. Notably, in 16 of the 18 sectors, no planet candidates were classified as contamination. In sector 53, two planet candidates were erroneously flagged as contaminants, while in sector 58, one planet candidate was misclassified. Our contaminant sample predominantly comprised eclipsing binaries, with a median percentage of 82%.

5.1. Challenges and Limitations

Despite our efforts to avoid incorrectly flagging planet candidates, three were misclassified by our model. All three are TOIs, two of them (TIC 243335710, TIC 372207084) have very short periods of 14.7 and 1.2 days respectively, and the other (TIC 219462190) has a low signal-to-noise ratio (S/N) of 7. This suggests that short-period and low-S/N planet candidates may be at higher risk of being incorrectly labeled by our model. However, the periods and S/Ns are only available for candidates designated as TOIs. Thus, we were unable to perform a meaningful analysis across the entire data set to formally test whether the planet’s S/N or period relates to its likelihood of being incorrectly flagged by our model.

Moreover, there is some uncertainty in our labels, which may be introducing some error to our model. The subjective nature of light curve labels and the occasional incomplete marking of transits by citizen scientists contribute to potential gaps in the training/validation data, potentially leading to biased or incomplete models. To mitigate this, project scientists could manually add transit marks to light curves before training and validating the model. This approach could improve model

performance and its ability to generalize to unseen data. However, it is essential to consider the resource-intensive nature of manual vetting, balancing the benefits of improved model performance against the costs and practical constraints associated with the process. Despite the ambiguity in what we interpret as “ground truth” in our training/validation, NotPlaNET can be a helpful tool to reduce the number of light curves requiring manual classification.

5.2. Comparative Analysis

The primary goal of NotPlaNET is to effectively differentiate between real transit events and false positives in photometric time-series data from the TESS mission. While injection and recovery tests using synthetic data are common practice in testing models to identify planetary signals, the complexity and variety of false-positive scenarios encountered in real TESS light curves is challenging to capture using simulations. These contaminants can include artifacts, asteroid crossings, background events, and satellite systematics, which are difficult to simulate accurately.

Given the limitations of simulated data in replicating these diverse scenarios, training or testing the model using a potentially limited set of simulations runs the risk of introducing biases. Instead, we have focused on training and validating the model with real TESS light curves identified by Planet Hunters TESS citizen scientists and confirmed by project scientists.

As shown in Figure 2, the fraction of recovered contaminants is notably lower for Sector 42 compared to the other test sectors. This may be due to Sector 42 containing the first TESS observation consisting primarily of the ecliptic plane, and as such contains an increased number of contaminants and a higher level of scattered light.

5.3. Future Directions

A follow-up study could assess whether model performance improves if employed earlier in the vetting process. NotPlaNET was trained on and applied to light curves at the end of the Planet Hunters pipeline, where differentiating between a planet candidate and a false positive becomes increasingly difficult. Thus, we suspect our model could also be applied to light curves earlier in the classification process with potentially better results. Employing NotPlaNET in different stages of the vetting process while probing a variety of classification thresholds is a possible direction for future research.

Additionally, our model could be adapted for other science goals by revisiting the scoring method and threshold selection. An alternative method could be to create a scoring system that rewards or penalizes finding contaminants and removing planet candidates differently, depending on the science goals for which the model is being applied. The score S can be written in the form:

$$S = aN_c - bN_p, \quad (1)$$

where N_c is the number of removed contaminants, N_p is the number of removed planets, and a and b are weighing constants.

NotPlaNET will be incorporated into the late stages of the Planet Hunters TESS pipeline. After the citizen scientist vetting process described in Section 2, a neural network designed by Malik et al. (2022; PlaNET) will be employed to calculate a

score for the vetted light curves and rank them accordingly. Our network will be applied after `PlaNET` to remove light curves that are confidently contaminants before the project scientist vetting stage.

6. Conclusion

We designed `NotPlaNET`, a CNN trained to identify contaminants from TESS light curves to reduce the number of light curves requiring human vetting. Our model was able to identify a maximum of 37% and a minimum of 10% of the contaminants across the 18 test sectors, with a median percentage of 18%. None of the planet candidates for 16 of the 18 test sectors were incorrectly flagged. For sector 53, two planets (0.6%) were flagged as contaminants. In sector 58, one planet candidate (0.3%) was incorrectly classified as a contaminant. Moreover, we found our contaminant sample to have a median eclipsing binary percentage of 82%.

The described methodology has shown potential to help in the identification of exoplanet candidates in large surveys. Our model can be applied to almost any light curve data set with a variety of science goals. Modifying `NotPlaNET` and the classification thresholds to accommodate light curves of other dispositions or to find other kinds of astronomical objects like eclipsing binary stars is a potentially interesting follow-up project. Nevertheless, when employed to remove contaminants from TESS data as described in this paper, our model can significantly help reduce the number of light curves that require manual vetting by project scientists and potentially increase our chances of finding and studying interesting long-period planets.

Although human vetting is, as of now, indispensable in the classification and identification of transiting exoplanets, machine learning continues to show great potential to help speed up the process. Integrating machine learning methods in the vetting pipeline can allow project and citizen scientists to more efficiently parse the data from today's large-scale astronomical surveys.

Acknowledgments

It is a pleasure to thank Shirley Ho (Flatiron), Artemis (cat), and the members of the Blanton–Hogg group meeting at NYU for valuable advice and input. Special thanks to Cole for his healthy skepticism, which motivated V.T.P. to prove him wrong and finish this paper.

This paper includes data collected by the TESS mission. Funding for the TESS mission is provided by the NASA Science Mission Directorate. This publication also uses data

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

The Flatiron Institute is a division of the Simons Foundation.
Facilities: TESS, MAST.

Software: NumPy (Harris et al. 2020), Matplotlib (Hunter 2007), PyTorch (Paszke et al. 2019), scikit-learn (Pedregosa et al. 2011).

ORCID iDs

Valentina Tardugno Poleo  <https://orcid.org/0009-0002-3883-7490>

Nora Eisner  <https://orcid.org/0000-0002-9138-9028>

David W. Hogg  <https://orcid.org/0000-0003-2866-9403>

References

- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
- Biewald, L. 2020, Experiment Tracking with Weights and Biases, <https://www.wandb.com/>
- Crossfield, I. J. M., Guerrero, N., David, T., et al. 2018, *ApJS*, **239**, 5
- Dattilo, A., Vanderburg, A., Shallue, C. J., et al. 2019, *AJ*, **157**, 169
- Eisner, N. L., Barragán, O., Aigrain, S., et al. 2020, *MNRAS*, **494**, 750
- Eisner, N. L., Barragán, O., Lintott, C., et al. 2021, *MNRAS*, **501**, 4669
- Fischer, D. A., Schwamb, M. E., Schawinski, K., et al. 2012, *MNRAS*, **419**, 2900
- Foreman-Mackey, D., Morton, T. D., Hogg, D. W., Agol, E., & Schölkopf, B. 2016, *AJ*, **152**, 206
- Guerrero, N. M., Seager, S., Huang, C. X., et al. 2021, *ApJS*, **254**, 39
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jenkins, J. M., Twicken, J. D., McCauliff, S., et al. 2016, *Proc. SPIE*, **9913**, 99133E
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, **391**, 369
- Malik, S. A., Eisner, N. L., Lintott, C. J., & Gal, Y. 2022, arXiv:2211.06903
- Olmschenk, G., Ishitani Silva, S., Rau, G., et al. 2021, *AJ*, **161**, 273
- O'Shea, K., & Nash, R. 2015, arXiv:1511.08458
- Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv:1912.01703
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825, <http://jmlr.org/papers/v12/pedregosa11a.html>
- Prša, A., Kochoska, A., Conroy, K. E., et al. 2022, *ApJS*, **258**, 16
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, *JATIS*, **1**, 014003
- Salinas, H., Pichara, K., Brahm, R., Pérez-Galarce, F., & Mery, D. 2023, *MNRAS*, **522**, 3201
- Schwamb, M. E., Orosz, J. A., Carter, J. A., et al. 2013, *ApJ*, **768**, 127
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, **155**, 94
- Tey, E., Moldovan, D., Kunimoto, M., et al. 2023, *AJ*, **165**, 95
- Wang, J., Fischer, D. A., Barclay, T., et al. 2013, *ApJ*, **776**, 10
- Yu, L., Vanderburg, A., Huang, C., et al. 2019, *AJ*, **158**, 25
- Zucker, S., & Gyires, R. 2018, *AJ*, **155**, 147