



Scientific Domain Knowledge Improves Exoplanet Transit Classification with Deep Learning

Megan Ansdell¹ , Yani Ioannou² , Hugh P. Osborn³ , Michele Sasdelli⁴
(2018 NASA Frontier Development Lab Exoplanet Team),

and

Jeffrey C. Smith^{5,6} , Douglas Caldwell^{5,6} , Jon M. Jenkins⁵ , Chedy Räissi⁷ , and Daniel Angerhausen⁸
(2018 NASA Frontier Development Lab Exoplanet Mentors)

¹ Center for Integrative Planetary Science, University of California at Berkeley, USA; ansdell@berkeley.edu

² Machine Intelligence Lab, University of Cambridge, UK

³ Aix Marseille Univ, CNRS, CNES, Laboratoire d'Astrophysique de Marseille, France

⁴ Australian Institute for Machine Learning, University of Adelaide, Australia

⁵ NASA Ames Research Center, USA

⁶ SETI Institute, USA

⁷ Institut national de recherche en informatique et en automatique, France

⁸ Center for Space and Habitability, University of Bern, Switzerland

Received 2018 October 31; revised 2018 November 18; accepted 2018 November 19; published 2018 December 5

Abstract

Space-based missions such as *Kepler*, and soon the *Transiting Exoplanet Survey Satellite (TESS)*, provide large data sets that must be analyzed efficiently and systematically. Recent work by Shallue & Vanderburg successfully used state-of-the-art deep learning models to automatically classify *Kepler* transit signals as either exoplanets or false positives; our application of their model yielded 95.8% accuracy and 95.5% average precision. Here we expand upon that work by including additional scientific domain knowledge into the network architecture and input representations to significantly increase overall model performance to 97.5% accuracy and 98.0% average precision. Notably, we achieve 15%–20% gains in recall for the lowest signal-to-noise transits that can correspond to rocky planets in the habitable zone. We input into the network centroid time-series information derived from *Kepler* data plus key stellar parameters taken from the *Kepler* data release 25 catalog. We also implement data augmentation techniques to alleviate model over-fitting. These improvements allow us to drastically reduce the size of the model, while still maintaining improved performance; smaller models are better for generalization, for example from *Kepler* to *TESS* data. This Letter illustrates the importance of including expert domain knowledge in even state-of-the-art deep learning models when applying them to scientific research problems that seek to identify weak signals in noisy data. This classification tool will be especially useful for upcoming space-based photometry missions focused on finding small planets, such as *TESS* and *PLATO*.

Key words: methods: data analysis – planets and satellites: detection – techniques: photometric

1. Introduction

The past 25 years have seen the flourishing of two contemporaneous yet disparate fields—that of exoplanets in astronomy, and that of deep learning in computer science; both have rapidly moved from predominantly theoretical to now largely data-driven regimes. For exoplanet science, this has been powered by the launch of wide-field, high-precision space telescopes designed to search for transiting exoplanets. These facilities—in particular NASA’s *Kepler Space Telescope* (Borucki 2016)—have discovered more than 3000 confirmed planets (>70% of the total known; e.g., Borucki et al. 2011; Mayo et al. 2018), enabling exoplanet population statistics that are revolutionizing our understanding of the universe (e.g., Dressing & Charbonneau 2015; Gaidos et al. 2016). For deep learning (see review in LeCun et al. 2015), large labeled data sets, increases in computational power, and modern techniques for training deep neural networks have brought about breakthroughs in computer vision, speech recognition, and natural language processing (e.g., Krizhevsky et al. 2012; Ioffe & Szegedy 2015; He et al. 2016).

These two fields now intersect in the detection and classification of transit-like signals in the large quantities of

data from space-based observatories like *Kepler* and soon the *Transiting Exoplanet Survey Satellite (TESS)*. These data need to be efficiently and reliably vetted for false-positive signals, such as those caused by stellar eclipses and instrumental noise, which largely outnumber the true planet transit signals. In particular, when searching for low signal-to-noise transit signals (e.g., as for small rocky planets in the habitable zone), chance correlations of stochastic instrumental and stellar signals can mimic transiting planet signals, making it extremely difficult to identify real transits just above the noise floor of the data. Deep learning—a machine learning tool named for its use of computational layers—provides a means to tackle these challenges.

For these reasons, exoplanet transit classification was selected as a project for the 2018 NASA Frontier Development Lab⁹ (FDL), an eight-week research incubator aimed at applying cutting-edge machine learning algorithms to challenges in the space sciences. NASA FDL teams consist of two machine learning experts and two space science researchers, with the aim of enabling more effective machine learning models with the help of scientific “domain knowledge”—i.e.,

⁹ <http://frontierdevelopmentlab.org>

the information, insight, or intuition relevant to a specific problem that a domain expert can provide that may not be immediately obvious to others.

The use of deep learning for automatically classifying candidate exoplanet transit signals has been previously explored by Shallue & Vanderburg (2018), who developed a convolutional neural network trained on *Kepler* data (see also Zucker & Giryes 2018 and Pearson et al. 2018 for applications of neural networks to simulated transit data). They clearly demonstrated the successful application of deep learning to transit classification, however improvements could be made with the inclusion of additional scientific domain knowledge. In this Letter, we present results from the 2018 NASA FDL program that investigated these possibilities. All of the code and data used in this Letter are publicly available.¹⁰

2. Data Sets

2.1. Flux Time Series (Light Curves)

In this Letter we use the Q1–Q17 *Kepler* Data Release 24 (DR24) light curves from the Mikulski Archive for Space Telescopes.¹¹ These were produced by the *Kepler* Science Processing Pipeline (Jenkins et al. 2010), which starts by calibrating the time-series target pixel file (TPF) images, then performs fixed-aperture photometry, and removes systematic instrumental errors (e.g., Smith et al. 2012; Stumpe et al. 2012). Each light curve consists of integrated flux measurements spaced at 30 minute intervals spanning up to 4 years ($\approx 70,000$ points) and contains one or more “threshold crossing events” (TCEs) identified by the *Kepler* pipeline. Each TCE is a potential exoplanet transit with a given period, epoch, and duration; however, most TCEs will be false-positive signals, sometimes caused by astrophysical phenomena such as eclipsing binaries (EBs) or background eclipsing binaries (BEBs), but also often by instrumental noise artifacts or other spurious events.

Following Shallue & Vanderburg (2018), we perform additional processing of the light curves for each TCE. First, we flatten the light curve by iteratively fitting a basis spline (with the in-transit points of the TCE excluded to preserve the transit signal), then divide the light curve by the best-fit spline while linearly interpolating over the transit points (see Shallue & Vanderburg 2018 for more details, and also Figure 3 in Vanderburg & Johnson 2014 for an illustration of this process). We implement a different spline-fitting routine than Shallue & Vanderburg (2018), which results in $5\times$ faster data processing times; namely, we use `LSQUnivariateSpline` in `SciPy` rather than `bspline` in `PyDL` (a library of Python replacements for IDL built-in functions). Second, we create “global” and “local” views of each phase-folded TCE following the description in Shallue & Vanderburg (2018). Both views are scaled so that the continuum is at 0 and the maximum transit depth is at -1 . The global view encapsulates the full view of the phase-folded light curve (e.g., including secondary transits of EBs) at the cost of long-period TCEs having poorly sampled transits. The local view, which depends on the transit duration, then provides a more detailed view of the primary transit shape. These two views are illustrated in

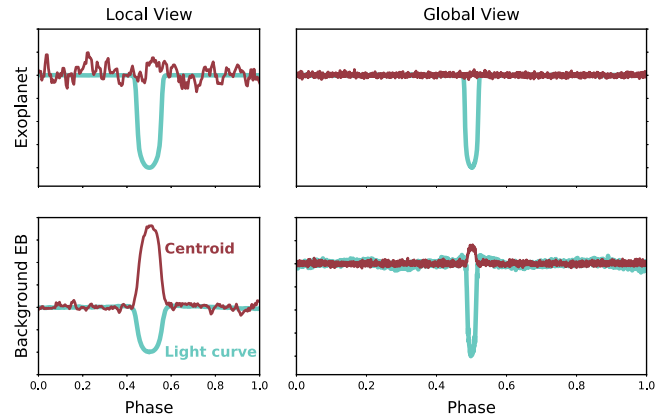


Figure 1. Local (left) and global (right) views of the light curves (cyan) and centroids (maroon) for an example confirmed planet (top) and background eclipsing binary (bottom), illustrating how the centroid curves can be used to identify a common type of astrophysical false positive.

Figure 1 and used as inputs into the deep learning models implemented in this Letter (see Figure 2; Section 3).

2.2. Centroid Time Series (Centroid Curves)

We also use the time series of the pixel position of the center of light (centroid) measured by the *Kepler* pipeline from the same TPF as the flux time series (Section 2.1). Centroids provide information on the position of the source of the transit-like signal, which is particularly useful for identifying BEBs. This is because centroids will shift in the opposite direction of the BEB if both the BEB and target star are contained within the photometric aperture used to measure the flux. We use the flux-weighted (first-moment) centroids rather than the pixel response function (PRF) centroids; although the PRF centroids are more robust against background noise, a significant number of sources do not have PRF centroid information, which complicates implementation of machine learning algorithms.

We use the x and y pixel coordinates of the centroid to compute the absolute magnitude ($r = \sqrt{x^2 + y^2}$) of the centroid displacement from the TPF center. We then follow the same process as the light curves for smoothing, phase-folding, and translating into local and global views. However, rather than normalizing the centroid curve to the maximum transit depth, we subtract the median and divide by the standard deviation, where these values are calculated out-of-transit and across the entire training data set (this standard practice is called “normalization” in machine learning). Moreover, we normalize the standard deviation of the centroid curves by that of the light curves, which ensures that TCEs with no significant centroid shifts show flat lines with noise signal strengths that are similar to that of the light curves (and thus do not dominate the signal strengths). Example phase-folded light curves and associated centroid curves, for both global and local views, are given in Figure 1 for a confirmed exoplanet and BEB.

2.3. Stellar Parameters

We use the updated *Kepler* DR25 catalog of Mathur et al. (2017) to obtain intrinsic stellar parameters. These parameters consist of stellar effective temperature (T_{eff}), surface gravity ($\log g$), metallicity ($[\text{Fe}/\text{H}]$), radius (R_*), mass (M_*), and density (ρ_*). These stellar parameters are normalized (Section 2.2): we subtract the median and divide by the standard deviation, where

¹⁰ <http://gitlab.com/frontierdevelopmentlab/exoplanets>

¹¹ <http://archive.stsci.edu/kepler>

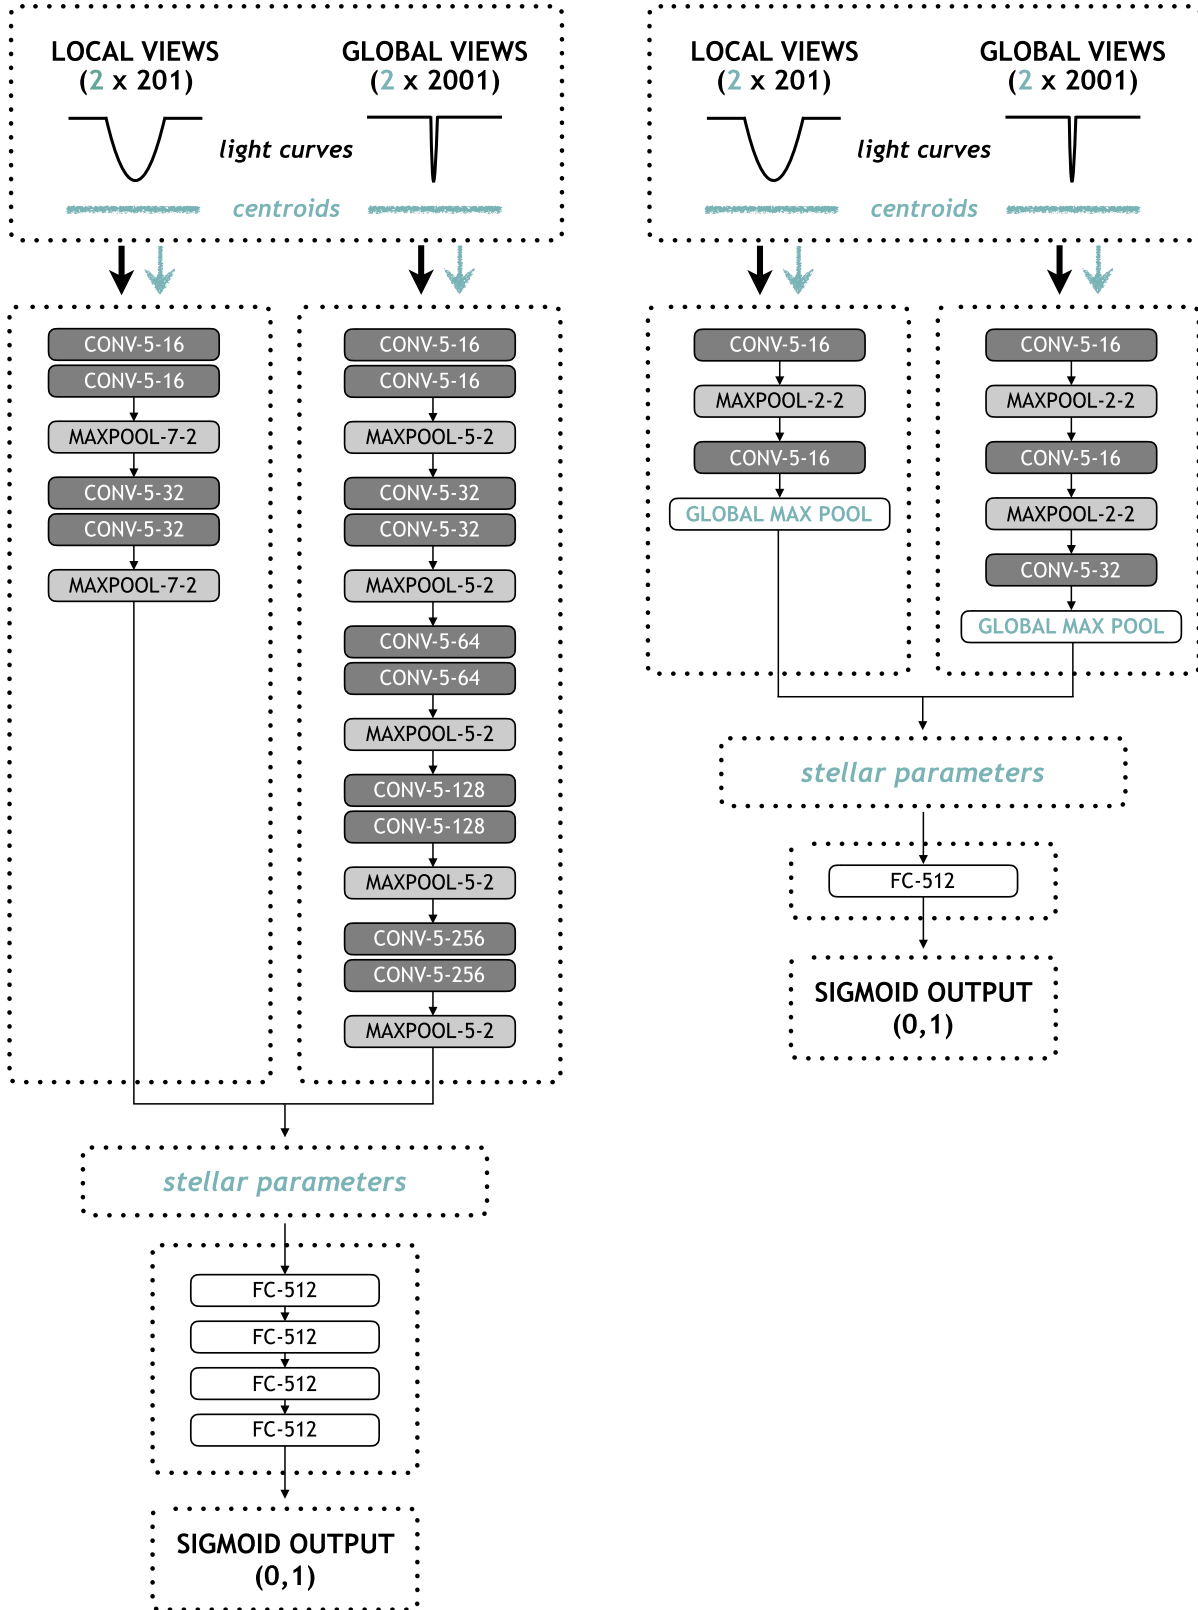


Figure 2. Convolutional neural network architectures used in this Letter. Left: Exonet, where the additions over the baseline Astronet model are shown in blue (Section 3.2). The flattened outputs of the disjoint one-dimensional convolutional columns are concatenated with the stellar parameters, then fed into the fully connected layers ending in a sigmoid function. Following Shallue & Vanderburg (2018), the convolutional layers are denoted as CONV-(kernel size)-(number of feature maps), the max pooling layers are denoted as MAXPOOL-(window length)-(stride length), and the fully connected layers are denoted as FC-(number of units). Right: the significantly reduced Exonet-XS model version described in Section 3.3.

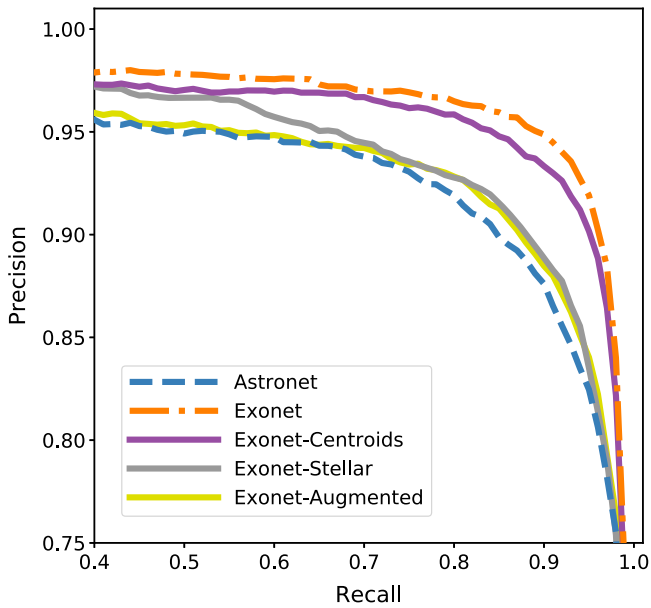


Figure 3. Precision-recall curve of Astronet (Section 3.1) compared to those of Exonet (Section 3.2) with different additions of scientific domain knowledge to show the individual contributions to increases in model performance. Exonet-Centroids is just the addition of the centroid curves, Exonet-Stellar is just the addition of the stellar parameters, and Exonet-Augmented is just the addition of our supplementary data augmentations. Exonet then combines all of these improvements into a single model.

these values are calculated for each parameter across the entire training set, such that the distribution of each parameter has a median of 0.0 and standard deviation of 1.0. We note that we tested the inclusion of proper motions and parallaxes from the *Gaia* DR2 catalog (Gaia Collaboration et al. 2018), but found no improvement in model performance.

2.4. Labels

We use the same labels as Shallue & Vanderburg (2018), which are taken from the *Kepler* DR24 TCE Table available on the NASA Exoplanet Archive.¹² The `av_training_set` column contains the labels used to train the *Autovetter* (McCaulliff et al. 2015) and primarily come from human-vetted KOIs assembled from multiple papers (e.g., Borucki et al. 2011; Batalha et al. 2013; Burke et al. 2014). The `av_training_set` column has four possible values: planet candidate (PC), astrophysical false positive (AFP), non-transiting phenomenon (NTP), and unknown (UNK). Following Shallue & Vanderburg (2018), we ignore the UNK TCEs (4630 entries) and then binarize the remaining labels as “planet” (PC; 3600 entries) or “false positive” (AFP + NTP; 12,137 entries). We then randomly divide the TCEs into training (80%), validation (10%), and test (10%) sets using the same random seed as Shallue & Vanderburg (2018) to preserve comparability. We use the validation set to tune hyperparameters and the test set for our final model performance results (Section 4).

3. Machine Learning Models

3.1. Astronet: Baseline Model

Here we briefly summarize *Astronet*, the deep convolution neural network developed by Shallue & Vanderburg (2018) that

Table 1
Ensembled Results on Test Set

Model	Accuracy	Avg. Precision
Astronet	0.958	0.955
Exonet	0.975	0.980
Astronet-XS	0.953	0.936
Exonet-XS	0.966	0.963

we use as our baseline model. *Astronet* is implemented in TensorFlow (Abadi et al. 2016), an open source software library for machine learning originally developed at Google Brain. As shown in Figure 2, the *Astronet* model architecture has two disjoint one-dimensional convolutional columns (one for the global view and one for the local view) with max pooling, the results of which are concatenated and then fed into a series of fully connected layers ending in a sigmoid function that produces an output in the range (0,1) that loosely represents the likeliness of a given TCE being a true planet transit (1) or false positive (0).

For model training, *Astronet* uses the Adam optimization algorithm (Kingma & Ba 2014) to minimize the cross-entropy error function. During training, data are augmented by applying time inversions to the input light curves with a 50% chance. The Google-Vizier system (Golovin et al. 2017) automatically tunes the hyperparameters of the input representations, model architecture, and training; consequently, the model is trained with a batch size of 64 for 50 epochs, and the Adam optimizer is implemented with a learning rate of $\alpha = 10^{-5}$, exponential decay rates of $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$ (see Kingma & Ba 2014 for details on these parameters). Shallue & Vanderburg (2018) use model “ensembling” to average the results from 10 independently trained versions of the same model with different random parameter initializations. Ensembling makes comparisons between different model architectures more robust because it averages over the stochastic differences in the individual models due to their different random parameter initializations. Moreover, ensembling improves model performance because the individual models can perform slightly better (or worse) in different regions of input space, in particular when the training set is small and thus prone to overfitting.

As part of this Letter, we re-implemented *Astronet* in PyTorch (Paszke et al. 2017) in an effort to expand the user base to those unfamiliar with TensorFlow. Our *Astronet* performance results given in Table 1 are consistent with those reported in Shallue & Vanderburg (2018); for example, we find an accuracy of 0.958 compared to their 0.960 value. We use the values in Table 1 as our baseline for comparison in the rest of this Letter.

3.2. Exonet: Revised Model with Domain Knowledge

Here we use scientific domain knowledge to add several features to our baseline *Astronet* model architecture and input representations in an effort to increase model performance. This modified model, which we call *Exonet*, is illustrated in Figure 2 and compared to the baseline *Astronet* model. The key modifications are described below. For model training, we retain the use of the Adam optimizer, cross-entropy loss function, batch size of 64, and learning rate of $\alpha = 10^{-5}$ used by *Astronet* (Section 3.1).

¹² <http://exoplanetarchive.ipac.caltech.edu/>

Addition of centroid time series: we input our analogous global and local views of the centroid time-series data (Section 2.2) as second channels of the disjoint convolutional columns used for the light curves. The motivation behind this architecture is to help the model learn the connections between the shapes of the light curves and centroid curves, which can be useful for identifying false positives, in particular BEBs (Figure 1). We note that the addition of centroid information was suggested by Shallue & Vanderburg (2018) as a potential avenue for improvement.

Addition of stellar parameters: we concatenate the stellar parameters (Section 2.3) to the flattened outputs of the convolutional layers directly before feeding them into the shared fully connected layers. We add this information because stellar parameters are likely correlated with classification; for example, giant stars with large radii are far more likely to host stellar eclipses than planetary transits (which would be undetectable).

Augmentation of training data: the baseline *Astronet* model augments the data by randomly flipping the time axis of half of the input light curves during training. We adopt this data augmentation technique, also applying the time-axis flip to the associated centroid curves. Because the training set is still quite small compared to typical machine learning problems, and because we found that *Astronet* suffers from model over-fitting, we apply an additional data augmentation technique during training to mimic measurement uncertainties in the flux measurements. Namely, we add random Gaussian noise to each input light curve, where the standard deviation is randomly chosen from a uniform distribution between 0 and 1.

3.3. *Exonet-XS: Reduced Model Size*

With the additions described in Section 3.2, we find that we can drastically reduce the size of the model architecture while still maintaining improved performance. This reduced model architecture, which we call *Exonet-XS*, is illustrated in Figure 2. We reduce the number of convolutional layers from 4 to 2 for the local column and from 10 to 3 for the global column. We also introduce global max pooling at the output of each convolutional column, as global max/average pooling has been shown to reduce the number of model parameters while increasing generalization (Lin et al. 2013; Ioannou et al. 2016), and is used in most state-of-the-art models for ImageNet (e.g., He et al. 2016). *Exonet-XS* has a model size that is $\approx 0.07\%$ that of the full *Astronet* model with $\approx 5 \times 10^{-4}$ fewer trainable parameters. Smaller models are often preferred, as they generalize better (or over-fit less; Hastie et al. 2001). Thus *Exonet-XS* should perform better when re-trained on other data sets, for example those expected from *TESS*.

4. Results

To assess model performance, we use three key metrics: accuracy, average precision, and precision-recall curves. Accuracy is the fraction of correct classifications by the model, for both planets and false positives, at a given threshold for deciding when the model output in the range (0, 1) becomes a positive class; we use a threshold of 0.5 for our accuracy calculations. Precision is the fraction of transits classified as planets that are true planets, while recall is the fraction of true planets recovered by the model; these can then be plotted on a precision-recall curve, which shows the trade-off between

precision and recall for different thresholds. The average precision summarizes the precision-recall curve as the weighted mean of precisions achieved at each threshold.

Table 1 gives the ensembled results on the test set, showing the 1.7% increase in accuracy and 2.5% increase in average precision of *Exonet* over *Astronet*. For the reduced *Exonet-XS* model, we see a 0.8% increase in accuracy and 0.8% increase in average precision. We also show results for the reduced model without the additions described in Section 3.2 (*Astronet-XS*), to illustrate that the gains in performance are still due to the inclusion of new scientific domain knowledge, rather than the change in model architecture.

Figure 3 shows precision-recall curves with each component of scientific domain knowledge added individually to illustrate their separate contributions to improving model performance. For this, we use k -fold cross-validation on the combined training and validation sets. Cross-validation is a method of evaluating model generalization performance that is more robust than using training and test sets; in k -folds cross-validation, the data is instead split repeatedly into k parts of equal size and multiple models are trained (here we use $k = 5$, which is typical). Figure 3 shows that the addition of the centroid time series provides the biggest gain in model performance, while stellar parameters also make a significant impact. Data augmentation does not greatly increase model performance by itself, rather the main benefit is to alleviate model over-fitting.

Figure 4 then shows the precision and recall as a function of a measure of the signal-to-noise of the candidate transits—the so-called “multiple event statistic” (MES; Jenkins et al. 2002) that the *Kepler* pipeline reports with each TCE. Notably, *Exonet* shows 15%–20% increases in recall for low-MES transits that often correspond to Earth-sized planets, some of which are in the habitable zone. The scatter in model performance at a given MES value is also noticeably smaller for *Exonet*, reflecting the robustness of the model. Note we do not perform hyperparameter optimization for *Exonet* and *Exonet-XS*, thus the model results could still be improved.

5. Conclusions

We expanded upon the method of Shallue & Vanderburg (2018) for applying deep learning to automatically classify *Kepler* candidate transit events with the addition of scientific domain knowledge. We used as our baseline model their convolutional neural network architecture, *Astronet*, which inputs “global” and “local” views of the phase-folded light curves through disjoint one-dimensional convolutional columns followed by shared fully connected layers that output a number in the range (0,1) that approximates the likeliness of a transit being a planet (1) or a false positive (0).

For our modified model, which we call *Exonet*, we created analogous global and local views of the centroid time-series data and input them as second channels of the disjoint convolutional columns, mainly to help identify BEBs. We then also concatenated key stellar parameters to the flattened outputs of the convolutional layers before feeding them into the shared fully connected layers; this helped to identify other types of false positives, such as giant star eclipsing binaries. Because we found that *Astronet* was prone to over-fitting, we also implemented data augmentation during training. *Astronet* already performed random time-axis reflections

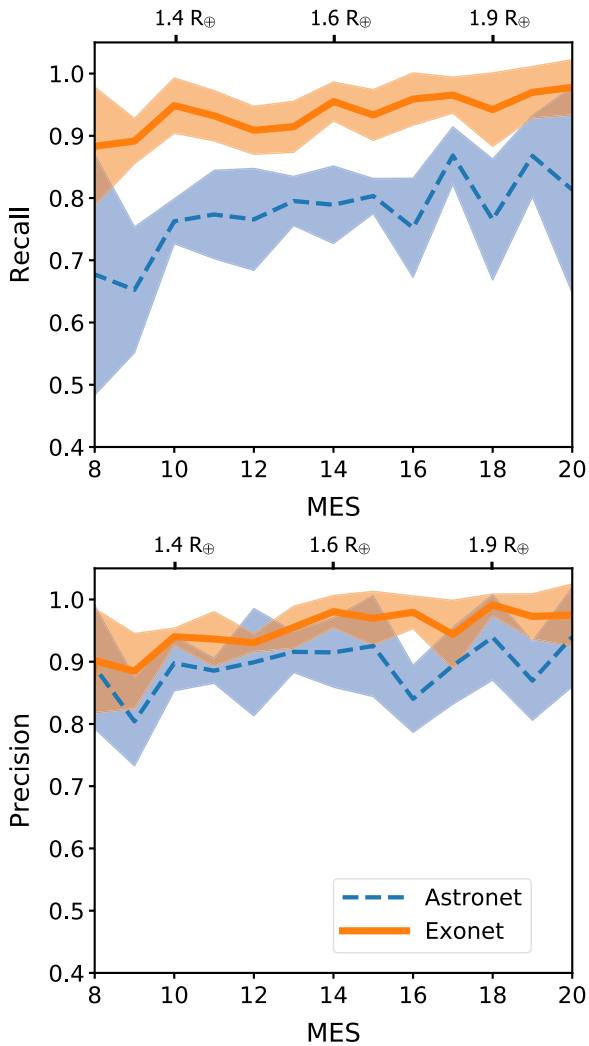


Figure 4. Recall (top) and precision (bottom) as a function of MES, which is a measure of the signal-to-noise of candidate transits. The *Kepler* pipeline only reports $MES > 7.1$ and the cross-validation k -folds sometimes did not contain any planets with $MES < 8$, thus here we only plot $MES \geq 8$. The solid lines are the averages of the cross-validation k -fold results, while the shaded regions show their standard deviations. We use a threshold of 0.7 to calculate the individual recall and precision values. The top axes show the median planet radius for confirmed/candidate KOIs in three MES bins, illustrating that the gains in performance by *Exonet* can be most significant for Earth-sized planets.

on the light curves, which we adopted and also applied to the corresponding centroid time series. Furthermore, we added random Gaussian noise to the input light curves to simulate uncertainties on the flux measurements.

These additions of scientific domain knowledge to the model architecture and input representations significantly improved the model accuracy and average precision by $\simeq 2.0\%$ – 2.5% , which is notable given the already impressive performance of *Astronet* prior to this Letter. Moreover, we showed that these gains in model performance are disproportionately high for low signal-to-noise transits that can represent the most interesting cases of rocky planets in the habitable zone. This demonstrates the importance of including domain knowledge in even state-of-the-art machine learning models when applying them to scientific research problems that seek to identify weak signals in noisy data. This classification tool will be especially useful for upcoming space-based photometry missions focused

on finding small planets, such as *TESS* (Ricker et al. 2014) and *PLATO* (Rauer et al. 2014). A forthcoming paper will document the application of our deep learning model to simulated *TESS* data (H. P. Osborn et al. 2018, in preparation).

We gratefully thank Google Cloud for providing the compute and storage resources critical to completing this work. M.A. also gratefully acknowledges the NVIDIA Corporation for donating the Quadro P6000 GPU used for this research. We thank Adam Lesnikowski, Noa Kel, Hamed Valizadegan, Yarin Gal, and Richard Galvez for their helpful discussions. This Letter includes data from the *Kepler* mission; funding for *Kepler* is provided by the NASA Science Mission directorate. The data used in this Letter were obtained from the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. M.A. also acknowledges support from NSF grant AST-1518332 and NASA grants NNX15AC89G and NNX15AD95G/NEXSS. H.P.O. acknowledges support from Centre National d’Etudes Spatiales (CNES) grant 131425-PLATO.

Software: Astropy (Astropy Collaboration et al. 2013, 2018), PyTorch (Paszke et al. 2017), Astronet (Shallue & Vanderburg 2018), Jupyter (Kluyver et al. 2016), SciPy (Jones et al. 2001), TensorFlow (Abadi et al. 2016), Matplotlib (Hunter 2007).

ORCID iDs

Megan Ansdell <https://orcid.org/0000-0003-4142-9842>
 Yani Ioannou <https://orcid.org/0000-0002-9797-5888>
 Hugh P. Osborn <https://orcid.org/0000-0002-4047-4724>
 Jeffrey C. Smith <https://orcid.org/0000-0002-6148-7903>
 Douglas Caldwell <https://orcid.org/0000-0003-1963-9616>
 Jon M. Jenkins <https://orcid.org/0000-0002-4715-9460>
 Chedy Räissi <https://orcid.org/0000-0002-6100-7519>
 Daniel Angerhausen <https://orcid.org/0000-0001-6138-8633>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, arXiv:1603.04467
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Batalha, N. M., Rowe, J. F., Bryson, S. T., et al. 2013, *ApJS*, **204**, 24
- Borucki, W. J. 2016, *RPPH*, **79**, 036901
- Borucki, W. J., Koch, D. G., Basri, G., et al. 2011, *ApJ*, **736**, 19
- Burke, C. J., Bryson, S. T., Mullally, F., et al. 2014, *ApJS*, **210**, 19
- Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, **807**, 45
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
- Gaidos, E., Mann, A. W., Kraus, A. L., & Ireland, M. 2016, *MNRAS*, **457**, 2877
- Golovin, D., Solnik, B., Moitra, S. et al. (ed.) 2017, Google Vizier: A Service for Black-Box Optimization, <http://www.kdd.org/kdd2017/papers/view/google-vizier-a-service-for-black-box-optimization>
- Hastie, T., Tibshirani, R., & Friedman, J. 2001, *The Elements of Statistical Learning*, Springer Series in Statistics (New York: Springer)
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in XIV European Conf. on Computer Vision (ECCV) 5, ed. B. Leibe et al. (Berlin: Springer), 630
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ioannou, Y., Robertson, D., Shotton, J., Cipolla, R., & Criminisi, A. 2016, International Conference on Learning Representations (ICLR), arXiv:1511.06744
- Ioffe, S., & Szegedy, C. 2015, in Proc. XXXII Int. Conf. on Machine Learning (ICML) 37, ed. F. R. Bach & D. M. Blei (Brookline, MA: Microtome), 448
- Jenkins, J. M., Caldwell, D. A., & Borucki, W. J. 2002, *ApJ*, **564**, 495
- Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010, *ApJL*, **713**, L87

- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, Proc. XXV Int. Conf. on Neural Information Processing Systems (NIPS), ed. F. Pereira et al. (Red Hook, NY: Curran Associates, Inc.), 1097
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, 521, 436
- Lin, M., Chen, Q., & Yan, S. 2013, CoRR, arXiv:1312.4400
- Mathur, S., Huber, D., Batalha, N. M., et al. 2017, *ApJS*, 229, 30
- Mayo, A. W., Vanderburg, A., Latham, D. W., et al. 2018, *AJ*, 155, 136
- McCauliff, S. D., Jenkins, J. M., Catanzarite, J., et al. 2015, *ApJ*, 806, 6
- Paszke, A., Gross, S., Chintala, S., et al. 2017, in Advances in Neural Information Processing Systems 31, ed. S. Bengio et al. (Red Hook, NY: Curran Associates, Inc.)
- Pearson, K. A., Palafox, L., & Griffith, C. A. 2018, *MNRAS*, 474, 478
- Rauer, H., Catala, C., Aerts, C., et al. 2014, *ExA*, 38, 249
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, *Proc. SPIE*, 9143, 914320
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, 155, 94
- Smith, J. C., Stumpe, M. C., Van Cleve, J. E., et al. 2012, *PASP*, 124, 1000
- Stumpe, M. C., Smith, J. C., Van Cleve, J. E., et al. 2012, *PASP*, 124, 985
- Vanderburg, A., & Johnson, J. A. 2014, *PASP*, 126, 948
- Zucker, S., & Giryes, R. 2018, *AJ*, 155, 147