The Astrophysical Journal, 927:219 (16pp), 2022 March 10 https://doi.org/10.3847/1538-4357/ac498f

Recovery of TESS Stellar Rotation Periods Using Deep Learning

Zachary R. Claytor[1], Jennifer L. van Saders[1], Joe Llama[2], Peter Sadowski[3], Brandon Quach[4], and Ellis A. Avallone[1]

[1] Institute for Astronomy, University of Hawai'i at Mānoa, 2680 Woodlawn Drive, Honolulu, HI 96822, USA; zclaytor@hawaii.edu

[2] Lowell Observatory, 1400 West Mars Hill Road, Flagstaff, AZ 86001, USA

[3] Department of Information and Computer Sciences, University of Hawai'i at Mānoa, 1680 East-West Road, Honolulu, HI 96822, USA

[4] Department of Computing and Mathematical Sciences, California Institute of Technology, 1200 East California Boulevard, MC 305-16 Pasadena, CA 91125, USA

Received 2021 April 28; revised 2022 January 7; accepted 2022 January 7; published 2022 March 17

Abstract

We used a convolutional neural network to infer stellar rotation periods from a set of synthetic light curves simulated with realistic spot-evolution patterns. We convolved these simulated light curves with real TESS light curves containing minimal intrinsic astrophysical variability to allow the network to learn TESS systematics and estimate rotation periods despite them. In addition to periods, we predict uncertainties via heteroskedastic regression to estimate the credibility of the period predictions. In the most credible half of the test data, we recover 10% accurate periods for 46% of the targets, and 20% accurate periods for 69% of the targets. Using our trained network, we successfully recover periods of real stars with literature rotation measurements, even past the 13.7 day limit generally encountered by TESS rotation searches using conventional period-finding techniques. Our method also demonstrates resistance to half-period aliases. We present the neural network and simulated training data, and introduce the software `butterpy` used to synthesize the light curves using realistic starspot evolution.

Unified Astronomy Thesaurus concepts: Starspots (1572); Stellar physics (1621); Stellar rotation (1629); Stellar properties (1624); Convolutional neural networks (1938); Computational astronomy (293); Astronomical simulations (1857); Light curves (918)

Supporting material: interactive figure, tar.gz file

1. Introduction

Stellar rotation is fundamentally linked to the structure and evolution of stars. In the decade since the Kepler mission, much has been learned about rotation, feeding into asteroseismology, empowering gyrochronology, and changing the way we think about stellar-evolution codes. Rotation period estimates are made possible through a variety of methods. Historically, spectroscopy enabled estimates of rotation velocity due to Doppler red/blueshift from the receding/approaching halves of a stellar disk. The projected rotation velocity could then be used to compute an upper limit on the period if the stellar radius was known. Missions like Convection Rotation and planetary Transits (CoRoT; Baglin et al. 2006) and Kepler (Borucki et al. 2010) have shifted the paradigm: the majority of period estimates now employ photometry instead of spectroscopy. This works particularly for stars which, like the Sun, exhibit magnetic dark and bright spots that induce periodic variations to the light curves as the stars rotate. Several techniques have been developed in recent years to extract rotation information from spot-modulated stellar light curves. Namely, Lomb–Scargle periodograms (LSPs; Marilli et al. 2007; Feiden et al. 2011), autocorrelation analysis (McQuillan et al. 2013, 2014), wavelet transforms (Mathur et al. 2010; García et al. 2014), Gaussian processes (Angus et al. 2018), and combinations of these (Ceillier et al. 2017; Santos et al. 2019; Reinhold & Hekker 2020) have all been used to infer rotation periods from light curves.

Period-finding methods have paved the way for large studies of stellar rotation. Applied to CoRoT and Kepler, these techniques have delivered tens of thousands of rotation period estimates, which in turn have been used to advance our understanding of stellar and Galactic evolution (e.g., McQuillan et al. 2014; van Saders et al. 2016; Davenport 2017; van Saders et al. 2019; Amard et al. 2020; Claytor et al. 2020). The Transiting Exoplanet Survey Satellite (TESS; Ricker et al. 2015) stands to increase the number of inferred periods by an order of magnitude in its ongoing all-sky survey.

Rotation studies have also brought to light the limitations of period-detection methods. For example, conventional methods are subject to aliases from multiple groups of starspots (e.g.,"double-dipper" stars; Basri & Nguyen 2018), and they still struggle to detect rotation in quiet, Sun-like stars (McQuillan et al. 2014; van Saders et al. 2019; Reinhold & Hekker 2020). Furthermore, the traditional methods do not necessarily reveal a star's true period. Rather, they reveal the period(s) of latitudes at which starspots form, which may rotate faster or slower than the star's equator due to surface differential rotation.

Finally, the systematics of TESS have made traditional period searches difficult (Oelkers & Stassun 2018; Canto Martins et al. 2020; Holcomb 2020; Avallone et al. 2021). The lunar-synchronous orbit of TESS has a 13.7 day period, and the telescope is subject to background variations from reflected sunlight causing periodic contamination that is difficult to remove. As a result, dedicated rotation studies struggle to obtain reliable periods longer than about 13 days (e.g., Canto Martins et al. 2020; Holcomb 2020; Avallone et al. 2021). New, data-driven methods are needed to overcome these systematics and recover periods.

Deep learning is relatively new to astronomy, but in a short time deep-learning methods have proven to be valuable at mining information from large data sets. Perhaps the most well-known use of deep learning is image recognition (e.g., He et al.

2015). Deep neural networks are extremely effective and efficient at processing and classifying two-dimensional (2D) images, making them ideal tools to use on astronomical data. Neural networks, and in particular convolutional neural networks (CNNs), efficiently extract information from time series, spectra, and image data. The strength of CNNs comes from their local connectivity, which incorporates the knowledge that neighboring input points are highly correlated. Examples of success using CNNs in astronomy include Hezaveh et al. (2017), who used CNNs to characterize gravitational lenses from image data. Within the realm of stellar astrophysics, Guiglion et al. (2020) used the same techniques to obtain stellar parameters from spectra. Moreover, Feinstein et al. (2020) and Blancato et al. (2020) used CNNs to infer stellar parameters and flare statistics from light curves.

Using CNNs, we predict stellar rotation periods from wavelet transforms of light curves. The use of supervised machine learning requires the existence of a training data set for which the target, in this case the rotation period, is known. This is not yet possible with TESS due to the difficulties of obtaining reliable periods using traditional techniques. Furthermore, while there is some overlap between TESS and Kepler, the TESS observations of the Kepler field are short, spanning only 27 days for most targets. This is enough to recover and validate short rotation periods (Blancato et al. 2020), and possibly some subset of longer periods (Lu et al. 2020), but not enough to be useful for the broader population of stars in our Galaxy. Moreover, even large rotation samples from Kepler are likely contaminated with mismeasured periods (Aigrain et al. 2015). Using a training set of periods obtained with conventional techniques risks imprinting this contamination onto the neural network. To avoid this, we followed the approach of Aigrain et al. (2015) and used a set of synthetic light curves generated from physically motivated starspot-emergence models. This is an example of simulation-based inference (e.g., Cranmer et al. 2020), wherein we simulate a physical process and use machine learning to address the inverse problem of inferring the rotation.

We introduce `butterpy`,[5] an open-source Python package designed to simulate realistic starspot emergence and synthesize light curves, followed by a description of the input physics of the simulations. After describing our training set, we outline our CNN and the methods we use to train, validate, and test the network. We evaluate our trained neural network on synthesized data sets spanning different period ranges to identify for what periods the network is most predictive. Next, we discuss the network's performance on a small set of real light curves for which rotation periods are known. We also compare our network predictions to periods recovered using conventional methods before finally concluding with thoughts on the feasibility of our methods to recover stellar rotation periods from real TESS light curves.

2. Synthetic Light Curves: `butterpy`

Synthesized light curves have several advantages over observed light curves: (1) the true, equatorial period of the simulated star is known, rather than an estimate of the period (which may be wrong), (2) data of any length and cadence can be synthesized, and (3) other physical properties like spot characteristics, differential rotation, and surface activity are known and can be independently probed.

To simulate light curves, we developed `butterpy`, a Python package designed to generate realistic, physically motivated spot-emergence patterns faster than conventional surface-flux-transport codes. We invite the community to extend and modify the spot-emergence model in the open-source GitHub repository. The name `butterpy` comes from the butterfly-shaped pattern of spot emergence with time exhibited by the Sun (e.g., Hathaway 2015). We built upon the software developed by Aigrain et al. (2015), which relied on the flux-transport models of Mackay et al. (2004) and Llama et al. (2012) to generate dark-spot-emergence distributions which were in turn used to compute light curves. Aigrain et al. (2015) discuss in detail how well these model light curves match those seen in nature; but, broadly, they reproduce the solar irradiance and resemble light-curve ensembles observed by CoRoT and Kepler. The original model of Mackay et al. (2004) was designed to reproduce spot-emergence patterns of the Sun as well as Zeeman Doppler images of the pre-main-sequence star AB Doradus. Later, Llama et al. (2012) used this model in tandem with exoplanet transit observations to trace the migration of active latitude bands across the surfaces of stars. Aigrain et al. (2015) used the light curves generated from these spot distributions to test the recovery rates of various period-detection techniques, which we seek to emulate. We discuss the method and assumptions here for clarity.

2.1. Spot Emergence and Light-curve Computation

Like Aigrain et al. (2015), we model starspots as uniformly dark circles on the surface of the star. For this exercise, we do not include bright plage or faculae, nor do we allow for nonuniformity of spot intensity or shape. However, we expect our model to capture most of the variance contained in real light curves. Beyond these basic assumptions, there are several variables to consider regarding the emergence of starspots and their effect on the star's light curve, including the latitudes and rates of emergence, the spot lifetimes, and the rotation speed at the latitude of emergence if the star rotates differentially. While observations of these characteristics on stars other than the Sun are limited, they are very well characterized for the Sun (Hathaway 2015, and references therein). For our model, we therefore start with the properties that are known for the Sun and allow them to vary.

2.1.1. Location and Rate of Spot Emergence

The latitudes of spot emergence on the Sun vary with the Sun's 11 year activity cycle. At the beginning of a cycle, spots emerge within active regions at high latitudes ($\lambda \approx \pm 30^\circ$; Hathaway 2015), and the latitude of emergence migrates toward the equator throughout the rest of the cycle. Before the cycle ends, new spot groups begin forming again at high latitudes, indicating some amount of overlap between consecutive cycles. The repeating decay in spot latitude with time gives rise to a butterfly-like pattern known as a "butterfly diagram". Butterfly patterns have been observed in other stars, as well (Bazot et al. 2018; Nielsen et al. 2019; Netto & Valio 2020), but some stars show a random distribution of spot-emergence latitudes with time (e.g., Mackay et al. 2004). The width of active latitude bands has also been shown to differ even for Sun-like stars (Thomas et al. 2019). In our model, we

---

[5] https://github.com/zclaytor/butterpy (Claytor et al. 2021).

allow for either a random or butterfly-like spot emergence between a minimum and maximum latitude that are unique for each star. For the butterfly pattern, spots begin the cycle emerging at a latitude $\lambda_{\max}$, decaying exponentially with time to latitude $\lambda_{\min}$ at the end of the cycle (Hathaway 2011).

As for longitude, spots on the Sun tend to emerge in groups, either next to existing spots, or in some cases antipolar to existing spots (Hathaway 2015). Less often, spots will emerge at random longitudes, not necessarily associated with any existing spot groups. We respectively refer to these two cases as correlated and uncorrelated active regions. In our simulations, we follow the approach of Aigrain et al. (2015), dividing the stellar surface into 16 latitude and 36 longitude bins; the probability of spot emergence is distributed across these bins. To account for the relative likelihood of correlated and uncorrelated emergence, bins already containing active regions are assigned a higher probability of emergence.

The rates of sunspot emergence change with spot area and with time throughout the activity cycle. Schrijver & Harvey (1994) expressed the number of spots emerging in area interval $(a, a + da)$ and time interval $(t, t + dt)$ as $r(t)a^{-2}dadt$, where $r(t)$ represents the time-varying emergence rate amplitude, the active-region area a is in square degrees, and t is the time elapsed in the activity cycle, ranging from 0 to 1. For the time dependence, spots emerge very slowly at the beginning, more rapidly in the middle, and slowly again at the end (Hathaway et al. 1994). Mackay et al. (2004) modeled this using a squared sine function: $r(t) = A\sin^2(\pi t)$. The activity level A is an adjustable scale factor controlling both the average rate of spot emergence and the amplitude of light-curve modulation of a single spot. It is defined such that $A = 1$ for the Sun.

2.1.2. Latitudinal Differential Rotation

We define asterographic longitude such that $\phi = 0$ always faces the Earth. As a consequence, spots move in longitude as the star rotates. The Sun rotates more rapidly near the equator than at the poles, a phenomenon known as "latitudinal differential rotation" (henceforth just "differential rotation"). While differential rotation is more difficult to observe on other stars, some stars may exhibit "antisolar" differential rotation, wherein the equator rotates more slowly than the poles. This has been observed particularly in slowly rotating stars (e.g., Rüdiger et al. 2019). In our model, we allow for solar-like, antisolar, and solid-body rotation. Following Aigrain et al. (2015), we model the differential rotation profile as

$$\begin{aligned} \phi_k(t) &= \phi_k(t_{\max,k}) + \Omega(\lambda_k)(t - t_{\max,k}), \\ \Omega(\lambda_k) &= \frac{2\pi}{P_{\rm eq}}(1 + \alpha\sin^2\lambda_k). \end{aligned} \tag{1}$$

Here, λ_k and ϕ_k denote the asterographic latitude and longitude of spot k. $t_{\max,k}$ is the time at which the spot achieves maximum flux, Ω is the angular velocity at latitude λ_k, and α is the differential rotation shear parameter. To include antisolar, solid-body, and solar-like profiles, we allow α to range from −1 to 1.

2.1.3. Spot-induced Flux Modulation

Once spot emergence is determined, we simulate spot evolution and flux modulation based on the simplified model of Aigrain et al. (2012, 2015). They take the photometric signature $\delta F_k(t)$ of a single spot k to be

$$\begin{aligned} \delta F_k(t) &= f_k(t)\max\{\cos\beta_k(t), 0\}, \\ \cos\beta_k(t) &= \cos\phi_k(t)\cos\lambda_k\sin i + \sin\lambda_k\cos i, \end{aligned} \tag{2}$$

where $\beta_k(t)$ is the angle between the spot normal and the line of sight, accounting for projection on the stellar surface. The inclination i is the angle between the rotation axis and the line of sight, and λ_k and $\phi_k(t)$ are again the latitude and longitude of the spot. The factor f_k is the amount of luminous flux removed if spot k is observed at the center of the stellar disk. Aigrain et al. (2015) used an exponential rise and decay to model the rapid emergence and slow decay of single spots, but we employ a two-sided Gaussian to avoid cusps in the curve of spot area with time while preserving the same emergence and decay behavior:

$$\begin{aligned} f_k(t) &= f_k^{\max}\exp[-(t - t_{\max,k})^2/\tau^2], \\ \tau &= \begin{cases} \tau_{\rm emerge} = \max\left\{2\ {\rm d}, \dfrac{\tau_{\rm spot}}{5}P_{\rm eq}\right\}, & t \leqslant t_{\max,k} \\ \tau_{\rm decay} = \tau_{\rm spot}P_{\rm eq}, & t > t_{\max,k}, \end{cases} \end{aligned} \tag{3}$$

where $f_k^{\max}$ is the flux removed by spot k at the time of maximum emergence $t_{\max,k}$, τ is the relevant emergence or decay timescale, and $\tau_{\rm spot}$ is a dimensionless parameter used to relate the emergence and decay timescales to the equatorial rotation period $P_{\rm eq}$. We verified that our results were not sensitive to the choice of exponential or Gaussian emergence and decay, so the choice is purely philosophical. Like Aigrain et al. (2015), we parametrized the emergence and decay timescales as multiples of the equatorial rotation period. The form of their emergence timescale was chosen so that, in general, emergence is five times faster than decay, with a minimum possible emergence timescale of 2 days. We adopt the same behavior for our simulations. In the simple model of Aigrain et al. (2012), $f_k^{\max}$ takes into account the spot area and contrast, but the model of Aigrain et al. (2015) relates this factor to the strength of the magnetic field:

$$f_k^{\max} = 3 \times 10^{-4}\, A\, B_{r,k}/\langle B_{r,k}\rangle_k, \tag{4}$$

where A is the activity level, and the constant is chosen such that $A = 1$ reproduces approximately Sun-like behavior.

With this expression, the single-spot luminous flux modulation is proportional to the radial (or vertical) magnetic field component at that spot. The magnetic field strength or magnetic flux is proportional to the area of the active region, which van Ballegooijen et al. (1998) derive using the angular width of magnetic bipoles emerging from the active region:

$$B_r^{\pm}(\theta, \phi) = B_{\max}\left(\frac{\beta_{\rm init}}{\beta_0}\right)^2 \exp\left[-\frac{2(1 - \cos\beta_{\pm}(\theta, \phi))}{\beta_0^2}\right], \tag{5}$$

where $B_r^{\pm}$ is the radial component of the magnetic field near either the positive or negative bipole, $B_{\max}$ is the initial peak magnetic field strength in the active region, $\beta_{\rm init}$ is the angular width of a single bipole, β_0 is the angular width (in degrees) of the bipole at the time the active region is inserted into the model, accounting for diffusion, and $\beta_{\pm}$ is the asterocentric angle between a field point and one of the bipoles.

van Ballegooijen et al. (1998) assumed that the bipole width $\beta_{\rm init}$ is proportional to the angular separation between the positive and negative poles, which they call $\Delta\beta$, with a proportionality factor of 0.4. Assuming spots form within 10 degrees of the active-region bipoles, the value of the exponential factor differs from unity by less than 1%. For this reason, we approximate the exponential factor as unity. Thus, at the location of a starspot, Equation (5) simplifies to

$$B_r \approx B_{\max}\left(\frac{0.4\Delta\beta}{\beta_0}\right)^2. \tag{6}$$

Combining this with Equation (4),

$$f_k^{\max} = 3 \times 10^{-4} A \left(\frac{\Delta\beta_k}{\langle\Delta\beta_k\rangle_k}\right)^2. \tag{7}$$

van Ballegooijen et al. (1998) and Mackay et al. (2004) consider a range of bipole widths from about 3.5° to 10°, which we adopt for our models. The distribution of $\Delta\beta$ is the same for every star, so $\langle\Delta\beta_k\rangle_k$ is effectively constant. While some low-mass stars have been suggested to have much larger spots covering substantial fractions of the surface, these stars will still be represented in our training set by models with many small spots. Putting it all together, we have a final system of equations to describe the change in luminous flux from a single spot:

$$\begin{aligned}
\delta F_k(t) &= f_k(t) \exp\left[-(t - t_{\max,k})^2 / \tau^2\right],\\
f_k(t) &= 3 \times 10^{-4} A \left(\frac{\Delta\beta_k}{\langle\Delta\beta_k\rangle_k}\right)^2 \max\{\cos\beta_k(t), 0\},\\
\tau &= \begin{cases} \max\left\{2\ \mathrm{d}, \dfrac{\tau_{\rm spot}}{5} P_{\rm eq}\right\}, & t \leqslant t_{\max,k}\\ \tau_{\rm spot} P_{\rm eq}, & t > t_{\max,k} \end{cases},\\
\cos\beta_k(t) &= \cos\phi_k(t) \cos\lambda_k \sin i + \sin\lambda_k \cos i.
\end{aligned} \tag{8}$$

This depends on $\Delta\beta_k$, λ_k, $\phi_k(t)$, and $t_{\max,k}$, which are unique to each spot; and A, $\tau_{\rm spot}$, $P_{\rm eq}$, and i, which are unique to each star.

2.2. Training Set

Using the model in `butterpy`, we generated one million light curves at 30 minute cadence and 1 yr duration to match the TESS full-frame images (FFIs) in the continuous viewing zones (CVZs). The simulation input parameters and their chosen distributions are listed in Table 1. While these parameters for real stars tend to correlate with one another (e.g., faster-rotating stars tend to be more active), we intentionally left the parameters mutually uncorrelated so as not to imprint our choices on the neural network predictions. The input distribution acts as a prior on the predictions of a neural network. Uniform distributions serve as uninformative ("flat") priors. We therefore chose all our distributions to be uniform or log-uniform (i.e., uniform under a log-transformation) to avoid applying a strong prior in the neural network predictions. We sampled periods uniformly from the range [0.1, 180] days, assuming no relationship between period and activity. The period range was chosen to be as wide as possible to simulate the fastest-rotating stars ($P_{\rm eq} \approx 0.1$ day) while also capturing anything that would go through at least two rotations under observation in the TESS CVZs (the total baseline is 350 days, so an object with $P_{\rm eq} = 180$ will go through almost two rotations in that time). We chose the remaining distributions and ranges to reflect those of Aigrain et al. (2015), which were chosen to resemble and slightly exceed the distributions seen in Kepler stars. We made minor adjustments to the ranges of activity level and differential rotation shear to search a broader parameter space. Figure 1 illustrates an example simulation, showing the distribution of spots on the surface as well as their impact on the observed light curve.

Table 1
Distribution of Simulation Input Parameters

Parameter	Range	Distribution
Equatorial rotation period, $P_{\rm eq}$	0.1–180 days	Uniform
Activity level, A	0.1–10× solar	Log-uniform
Activity cycle length, $T_{\rm cycle}$	1–40 yr	Log-uniform
Activity cycle overlap, $T_{\rm overlap}$	0.1 yr–$T_{\rm cycle}$	Log-uniform
Minimum spot latitutde, $\lambda_{\min}$	0°–40°	Uniform
Maximum spot latitude, $\lambda_{\max}$	$\lambda_{\min}$ + 5°–80 °	Uniform
Spot lifetime, $\tau_{\rm spot}$	1–10	Log-uniform
Inclination, i	0°–90°	Uniform in $\sin^2 i$
Latitudinal rotation shear, $\Delta\Omega/\Omega_{\rm eq}$	0.1–1 (50%)	Log-uniform
	0 (25%)	
	−1–0.1 (25%)	Log-uniform

Note. We adopted the distributions used by Aigrain et al. (2015) with minor modifications: (1) we sampled a broader range of periods and activity levels, (2) we used a uniform distribution of periods so as not to impart an unwanted prior on the neural network prediction, (3) we include antisolar differential rotation by allowing the shear parameter to be negative.

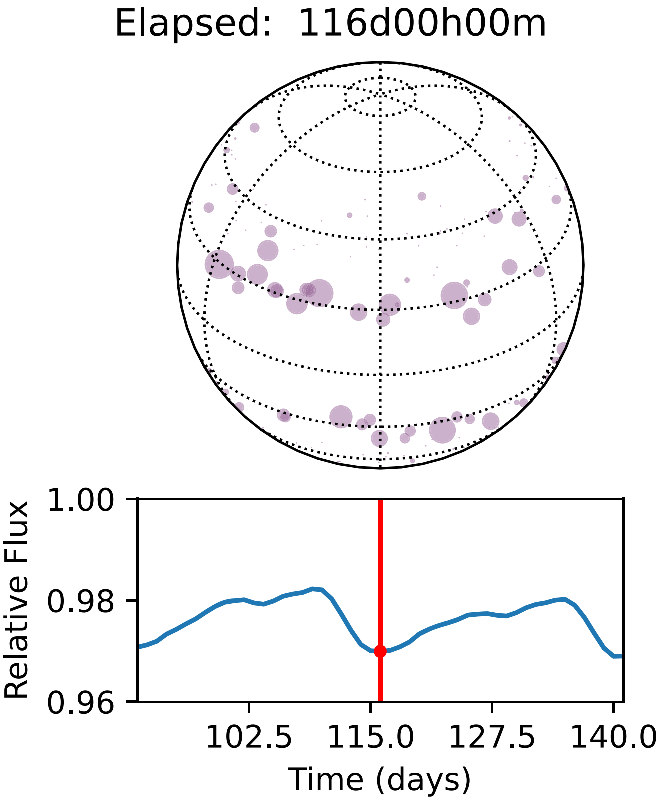

Figure 1. An example of a `butterpy` simulation of spot evolution and light-curve generation. This figure is available online as an interactive figure. The online figure has an interactive slider and play/pause buttons that allow the user to move the figure through time and see changes in the light curve as spots rotate into and out of view.

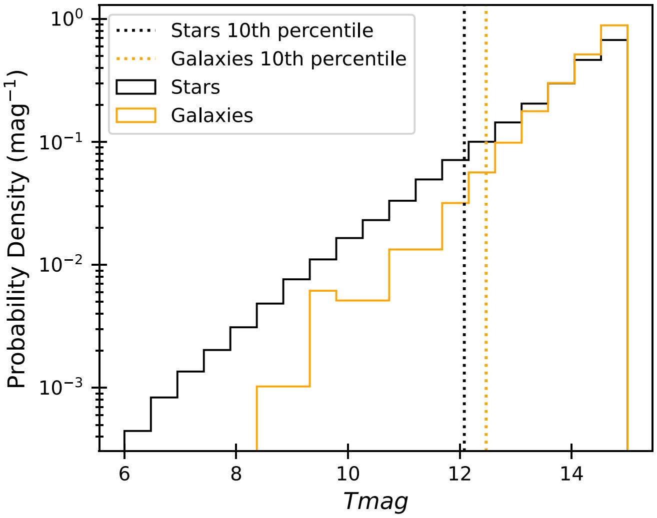

Figure 2. Brightness distributions of stars and galaxies in the TESS Southern Continuous Viewing Zone (SCVZ). The galaxies follow a similar brightness distribution to the stars in the SCVZ, especially toward the faint end where the majority of objects lie. Because the distributions are similar, the training light curves, which use the galaxy light curves to emulate noise, are representative of real distributions of stars' brightness and noise properties.

2.3. TESS Noise Model

To ensure the training light curves properly emulate real TESS light curves, the training set must exhibit TESS-like noise. Aigrain et al. (2015) used light curves from quiescent Kepler stars to achieve this. In their study, a sample of stars from McQuillan et al. (2014) with no significant period detection served as the quiescent data set. Because there are no existing bulk period measurements for stars in the CVZs, we must find another means of simulating TESS noise.

While TESS is a planet-finding mission, the southern CVZ contains thousands of galaxies which should have roughly constant brightness with time. Any changes in the light curves of these galaxies would be due solely to TESS instrument systematics. Furthermore, while galaxies are extended on the sky and not point sources, TESS's large (21″) pixel size means that a typical aperture is larger than most galaxies in the TESS Input Catalog. Indeed, 85% of SCVZ galaxies fit in the width of two pixels, and 93% fit within three pixels. Thus, as long as we choose sufficiently large apertures, the galaxies are effectively point-like, and the galaxy light curves should reasonably resemble light curves of quiescent stars in TESS.

We selected roughly 2000 galaxies in the southern CVZ with $T_{\mathrm{mag}} \leqslant 15$ as our quiescent sample, removing a handful of galaxies known to be active and in the Half-Million Quasars catalog (Flesch 2015). Figure 2 shows the brightness distributions of galaxies and stars in the TESS SCVZ, with vertical lines indicating the 10th percentiles of each distribution. The galaxies are distributed in apparent brightness similarly to the stars in the TESS SCVZ, so the training light curves will have distributions of apparent brightness and photometric precision resembling the distributions of real stars. We queried FFI cutouts from the Mikulski Archive for Space Telescopes (MAST) using `Lightkurve` and `TESScut` (Lightkurve Collaboration et al. 2018; Brasseur et al. 2019). Then, we performed background subtraction and aperture photometry on each source using `Lightkurve` regression correctors, following `Lightkurve` Collaboration (2020). To summarize, aperture masks were chosen using the `create_threshold_mask` function in `Lightkurve`. This method selects pixels with fluxes brighter than a specified threshold number of standard deviations above the image's median flux value. We specified thresholds based on the target's brightness to exclude background pixels from the aperture. Once the raw light curve was computed, the regression correctors fit principal components of the time-series images and subtracted the strongest components from the raw light curve. All sector light curves for a source were then median-normalized and stitched together to form the final "pure noise" light curve.

The galaxy light curves were linearly interpolated to each TESS cadence to fill gaps, whether for missing observations or entire missing sectors. Cadences missing at the beginning or end of the light curve were filled with the light curve's mean flux value. Finally, a galaxy light curve was chosen at random to be convolved with each of the synthetic light curves, yielding our final set of simulated TESS-like light curves. We note that while the light curves were median-normalized, this should not affect the intrinsic brightness or effective signal-to-noise ratio (S/N) of the light curves. The S/N can be thought of as the spot modulation amplitude divided by the photometric precision, both of which are defined relatively and are therefore preserved under median normalization. Thus, the ranges of brightness and photometric precision of our light curves reflect the underlying set of galaxy light curves, which are representative of stars in the TESS SCVZ. We divided the set of 2000 galaxies into two sets of 1000: one set to be convolved with light curves from the training partition, and one for the validation and test partitions (see Section 4 for more about data partitioning).

3. Data Processing/Wavelet Transform

There are several options for input to a neural network to predict rotation periods. One could use the light curve directly; Blancato et al. (2020) suggest this as the best way to obtain periods using neural networks without loss of information. However, using the light curve as input means that the information conveying periodicity is temporally spread out. While neural networks can certainly learn to predict periods this way, a frequency representation concentrates the period information to one location in input space. LSPs (Lomb 1976; Scargle 1982; Feiden et al. 2011) and autocorrelation functions (ACFs; McQuillan et al. 2013, 2014) are two tried-and-true methods of period estimation that have some promise as input to neural networks. While these methods are effective at concentrating periodicity information to one location, real stars' observed periodicities can change with time due to differential rotation. Lomb–Scargle and autocorrelation methods average over these changes, potentially blurring out interesting evolution.

The continuous wavelet transform (Torrence & Compo 1998) has also been used to identify rotation periods from stellar light curves (Mathur et al. 2010; García et al. 2014; Santos et al. 2019). When applied to time series, the wavelet transform creates a 2D image of frequency versus time; thus, it has the bonus of elucidating changes in periodicity with time. Additionally, the wavelet transform's 2D nature makes it an ideal partner to deep-learning approaches. Deep learning has been applied to wavelet transform representations in the fields of physiology (Zhao & Zhang 2005), economics (Bao et al. 2017), and more. We chose the wavelet transform as the best representation to localize periodic information while allowing the tracing of spot evolution, and as the most natural representation to leverage the decades of advances in image classification and analysis with deep learning.

We used the continuous wavelet transform implemented in `SciPy` (Virtanen et al. 2020) with the power spectral density correction of Liu et al. (2007). Using a Morlet wavelet, we computed wavelet power spectra for both the noiseless and noisy light curves in our training set. Examples of both noiseless and noisy power spectra are shown for the same simulated star in Figure 3. We then rebinned the power spectra to 64 × 64 pixels and saved them as arrays for fast access. Higher-resolution arrays were tested (e.g., 128 × 128) and showed no significant change in performance.

We ran several tests with the period axis of the wavelet power spectra, trying maximum periods of 128, 150, and 180 days, before settling on 180 days (i.e., half the observing window) for the final data products. In the tests with 128 and 150 days, the neural network successfully predicted periods longer than the maximum visible period in the periodogram, albeit at lower success rates. This suggests that neural networks can predict periods even when the period at maximum power is beyond the range of the plot, consistent with the results of Lu et al. (2020). This is encouraging for period predictions for stars outside the TESS continuous viewing zones, where observations are substantially less than a year in duration. In the end, we chose 180 days as the maximum value on the period axis to preserve the strongest rotation signals in as many objects as possible.

In addition to `butterpy`, the training light curves and wavelet power spectra are available on MAST as a High Level Science Product via doi:10.17909/davg-m919.[6]

4. Convolutional Neural Network

We used a CNN to predict rotation periods from wavelet transforms. Table 2 outlines the CNN architecture. We used a sequence of 2D convolution layers with rectified linear activation ("rectifier", or "ReLU") followed by 1D max-pooling in the time dimension. The ReLU activation function has the form $f(x) = \max(0, x)$. Its nonlinearity allows the model to represent complex functions, and ReLU learns faster than other nonlinear activation functions. Max-pooling is used to down-sample input and impart a small amount of translational invariance. The shapes of the convolution and pooling kernels were chosen to impart equivariance in the frequency dimension (no pooling, since frequency is what we want to estimate) and translational invariance in the time dimension. This means that the CNN will treat periodic signals the same regardless of when they occur in the wavelet power spectrum (see Ch. 9 of Goodfellow et al. 2016).

The output of the convolution layers is then flattened to one dimension and fed into a series of three fully connected layers, also with ReLU activation. The final layer uses softplus activation, which has the form $f(x) = \ln(1 + e^x)$. A smooth approximation to the rectifier, softplus activation ensures positive output while preserving differentiability. The final layer outputs two numbers, which represent the rotation period and the period uncertainty.

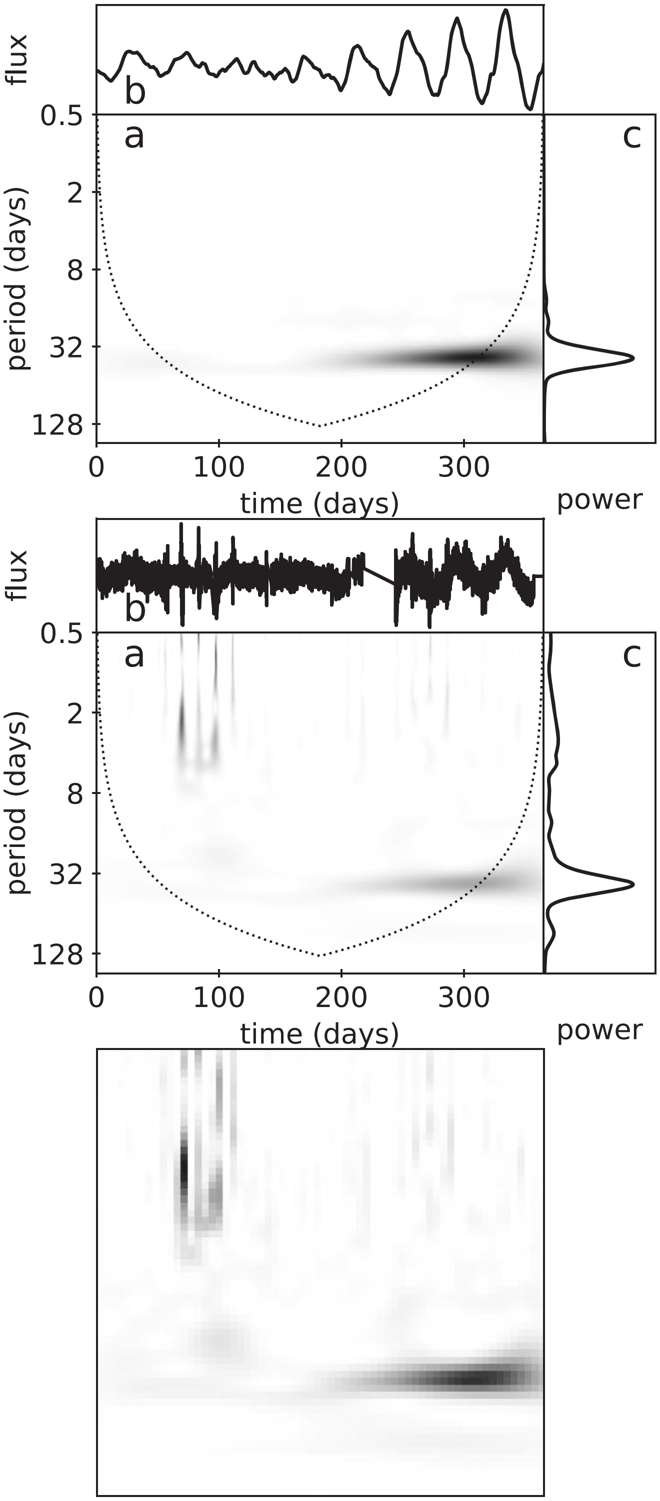

Figure 3. Top: Morlet wavelet transform (a) of a noiseless light curve, shown with the light curve (b) and global wavelet power spectrum (c). Center: plots for the same light curve convolved with TESS noise and restitched. The dotted curve marks the cone of influence, below which the power spectrum is susceptible to edge effects. Bottom: example of a binned wavelet power spectrum we used to train our neural network. The 64 × 64 pixel array, without labels, is what is fed into the network. Neural networks can learn to ignore the noise and pick out stellar signals.

The 2D wavelet power spectra were used as input to our neural network, while the corresponding model rotation periods served as the target output. Target periods were min-max scaled over the entire data set to the range [0, 1]. Each power spectrum array was min-max scaled to the range [0, 1] separately—using

---

[6] https://archive.stsci.edu/hlsp/smarts/

Table 2
Convolutional Neural Network Architecture

Layer Type	Number of Filters	Filter Size	Stride	Activation	Output Size
Input image	⋯	⋯	⋯	⋯	64 × 64
Conv2D	8	3 × 3	1 × 1	ReLU	62 × 62 × 8
MaxPool2D	1	1 × 3	1 × 3	⋯	62 × 20 × 8
Conv2D	16	3 × 3	1 × 1	ReLU	60 × 18 × 16
MaxPool2D	1	1 × 3	1 × 3	⋯	60 × 6 × 16
Conv2D	32	3 × 3	1 × 1	ReLU	58 × 4 × 32
MaxPool2D	1	1 × 4	1 × 4	⋯	58 × 1 × 32
Flatten	⋯	⋯	⋯	⋯	1856
Dense	⋯	⋯	⋯	ReLU	256
Dense	⋯	⋯	⋯	ReLU	64
Dense	⋯	⋯	⋯	Softplus	2

Note. We use three 2D convolution layers, each with ReLU activation and max-pooling. Our implementation uses 2D max-pooling with a one-dimensional (1D) kernel to achieve pooling in the time dimension but not the frequency dimension. This choice preserves frequency resolution but achieves a small amount of translational invariance in the time dimension. The output of the convolution block is flattened to a 1D array and passed through three fully connected (dense) layers, with ReLU and finally softplus output, to yield two numbers: the rotation period and its uncertainty.

the min and max over the entire data set suppressed lower-amplitude signals and substantially impaired performance.

Our full data set of one million examples was partitioned into three sets for model training, validation, and testing. The training set consisted of 80% and was used to fit the model weights. The validation set (10%) was used for early stopping (we stop training when the validation loss does not improve over a window of 10 training epochs to avoid overfitting) and for choosing the optimal hyperparameters. The test set (10%) was used for final model evaluation.

We used the Adam optimizer (Kingma & Ba 2014), which allows for a variable learning rate, with negative log-Laplacian likelihood as the loss function. This loss function allows us to predict both the rotation period and its error (a process known as heteroskedastic regression), indicating which period predictions are more reliable than others. It has the form

$$\mathcal{L} = \ln(2b) + \frac{|P_{\rm true} - P_{\rm pred}|}{b}, \tag{9}$$

where b is taken to represent the predicted uncertainty.

Maximizing the log-likelihood of the Laplace distribution is equivalent to minimizing the mean absolute error instead of the mean-squared error, or predicting the median period instead of the mean. This also means that in cases where the neural network cannot predict with high confidence, predictions will be biased toward the median of the period range. Formally, the Laplace distribution has variance $2b^2$ and standard deviation $\sqrt{2}\,b$. Since we use the predicted uncertainty only to determine the relative credence of predictions, the number should not be considered statistically formal. With this in mind, we leave the notation in terms of b to represent the uncertainty for simplicity and to discourage a statistically formal interpretation of our uncertainty.

With 800,000 input–output pairs in the training set, our model takes roughly 3 hr until fully trained on a single NVIDIA RTX2080. Once trained, evaluation on the test input of 100,000 wavelet power spectrum plots takes less than a minute.

5. Results

We trained and evaluated the neural network on year-long simulations of both noiseless and noisy wavelet transform images. We additionally used LSPs, ACFs, and wavelet transforms to obtain independent period estimates from both the noiseless and noisy data.

Aigrain et al. (2015) performed blind injection-recovery exercises on synthesized Kepler-like light curves to assess the reliability of conventional period-detection methods. On average, the teams recovered periods with 10% accuracy in ∼70% of cases in which periods were obtained. We adopt this 10% accuracy threshold as our success metric, which we designate “acc10”. In addition to acc10, we also quantify results with “acc20”, mean absolute percentage error (MAPE), and median absolute percentage error (MedAPE), defined as follows. If we define the absolute percentage error of example i to be $\epsilon_i = |P_{{\rm pred},i} - P_{{\rm true},i}|/P_{{\rm true},i}$, then our recovery metrics are

$$\begin{aligned} {\rm MAPE} &= \frac{1}{N}\sum_i^N \epsilon_i, \\ {\rm MedAPE} &= {\rm median}\{\epsilon_i\}, \\ {\rm acc10} &= \frac{1}{N}\sum_i^N H(0.1 - \epsilon_i), \\ {\rm acc20} &= \frac{1}{N}\sum_i^N H(0.2 - \epsilon_i), \end{aligned} \tag{10}$$

where $H(x - \epsilon_i)$ is the Heaviside or unit step function.

Before commenting on our period recovery, it is important to note the differences in our light-curve sample from that of Aigrain et al. (2015). The most important differences are in the range of activity level and the light-curve preprocessing. Our sample spans a wider range of activity levels, ranging from 0.1 to 10 times solar, as opposed to 0.3–3 times solar in Aigrain et al. (2015). The logarithmic scale of the distribution ensures that the increase in range evenly adds examples to the high- and low-activity ends. Thus, despite having higher-amplitude examples in our sample, there should be enough lower-amplitude examples to compensate, preserving the comparability of our summary statistics to those of Aigrain et al. (2015). Light-curve preprocessing differs because of the differences in the Kepler pipeline and our custom TESS FFI pipeline. In principle, the Kepler pipeline more aggressively removes systematics, so the Aigrain et al. (2015) simulated light curves are cleaner than ours.

5.1. Period Recovery Using Conventional Methods

We recovered periods from our sample of noisy light curves using LSPs, as implemented in `Lightkurve` (Lightkurve Collaboration et al. 2018), ACFs (McQuillan et al. 2013; and, as implemented in `starspot`, Angus et al. 2018), and global wavelet power spectrum (GWPS, as implemented in `SciPy;` Virtanen et al. 2020). The recovery results are summarized in Figure 4. In each panel, objects falling within 10% of the line $y = x$ are successfully recovered according to our metric. In addition to showing predictions for the entire test set (left column), we also filtered predictions based on various periodogram peak significance metrics to demonstrate how recovery rates improve with confidence in the measurement. We removed the worst half of the data based on peak false-alarm probability for LSP, and peak prominence (as implemented in `SciPy`, Virtanen et al. 2020) for both GWPS and ACF. The remaining half of the test set is shown in the right column of Figure 4.

All three conventional methods struggle to recover periods longer than about 50 days. Longer than this, LSP and GWPS mistakenly recover signals approaching 30 days in period, which we suspect represents the TESS sector length of 27 days. LSP and GWPS are also susceptible to half-period aliases, which fall along the line $y = \frac{1}{2}x$. ACF is less susceptible to half-period aliases and appears to be the most successful method overall. However, the ACF and GWPS often misidentify signals at 5, 13, and 27 days (all well-known frequencies associated with TESS telescope systematics) as the rotation period. Interestingly, the ACF has small pockets of higher success at integer multiples of 27 days, beginning at 54 days. These occur when a star's rotation period coincides with an alias of the TESS sector length. Because the ACF is more likely to be detecting the systematic alias than the period, these coincidences should not be considered actual successes, suggesting that at long periods the ACF's true recovery probability is lower than it appears to be.

In general, recovery was better for targets with higher light-curve amplitude for all three methods, as one would expect. Filtering the sample based on periodogram peak significance marginally improved the recovery rates, but the predictions were still dominated by aliases at true periods longer than 50 days. The recovery rates also improved when limited to shorter periods (see Table 3). We have assumed no rotation–activity relation, so the improved recovery at shorter periods occurs when more rotations are observed in the given baseline, resulting in higher power in the periodograms. Moreover, at shorter periods rotation signals are less easily lost in the telescope systematics. In the unfiltered case, for example, if we limit to periods between 0 and 50 days, acc10 and acc20 improve to 45% and 61% for LSP, 43% and 56% for GWPS, and 42% and 54% for ACF. Filtering on periodogram peak significance and limiting to periods shorter than 50 days, acc10 and acc20 improve to roughly 60% and 80%, respectively, for all three techniques. We note that these short-period recovery rates are on par with those achieved by Aigrain et al. (2015) despite the differences in underlying period and amplitude distributions.

Even with filtering, conventional period-recovery techniques are subject to aliases, leaving much to be desired for long-period recovery in TESS. This is not surprising: LSP false-alarm probability indicates how likely a power spectrum peak is to arise stochastically. It does not account for the systematics and aliases like those seen in TESS. Similarly, the peak prominence used for GWPS and ACF is simply a measurement of the peak height relative to neighboring troughs. It offers little insight into whether a peak arises from astrophysical periodicity, as peaks from strong aliases easily overpower those from astrophysical signals. This motivates the use of a CNN, which should be able to discern between aliases and rotation signals.

5.2. CNN Performance on Noiseless Data

Our neural network's predictions on noiseless test data are shown in the left panel of Figure 5. The predicted periods have a mean absolute percentage error of 14% and a median absolute percentage error of 7%. 61% of periods were successfully recovered to within 10%, setting the bar for comparison to results for the noisy data.

5.3. CNN Performance on Noisy Data

We present the neural network predictions on the noisy test data in the right panel of Figure 5. The predicted periods have a mean absolute percentage error of 246% and a median absolute percentage error of 24%. Only 28% of periods are successfully recovered to within 10%. The horizontal band at the predicted period of 90 days represents simulated stars for which the network could not predict the period at all, instead assigning it the median of the period range. We discuss the filtering and removal of these spurious predictions later in this section.

The addition of TESS-like noise severely inhibits the performance of the neural network. Like the conventional methods, the neural network predictions are more accurate at shorter periods. When limiting to periods of 50 days or less, the median absolute percentage error is 12%, and 44% of targets are recovered to within 10%. The introduction of noise to the light curves also affects the amplitudes at which the network is most reliably predictive. The left panel of Figure 6 shows CNN recovery rate as a function of amplitude R_{per} (as defined by Basri et al. 2011) and equatorial rotation period. Here, "recovered" means the prediction is within 10% of the true period. As expected, the network performs better with higher-amplitude modulations, where the stellar signals are more easily picked out of the noise. For lower-amplitude modulations ($R_{\text{per}} \lesssim 10^4$ ppm), the light-curve noise dominates the stellar signal. At the lowest amplitudes, the recovery fraction approaches zero except for the range near 90 days where uncertain predictions are assigned the median of the period distribution. These low-amplitude light curves are analogous to light curves with no discernible spot modulation and provide an interesting look at how the CNN would handle light curves with no spots present. The predictions for these low-amplitude signals generally have large uncertainty, as we expect.

There are a handful of bins in Figure 6 with seemingly high amplitudes but unsuccessful period predictions and large uncertainties. Most, if not all, of these high-amplitude failed recoveries are due to the noise template used. In these few cases (<1% of galaxy light curves), the light curves all have one extremely noisy sector. This noisy sector dominates the wavelet transform, drowning out any rotation signal; it also dominates the amplitude measurement, masking the actual amplitude of the light curve. These noise templates will be removed from use in future work.

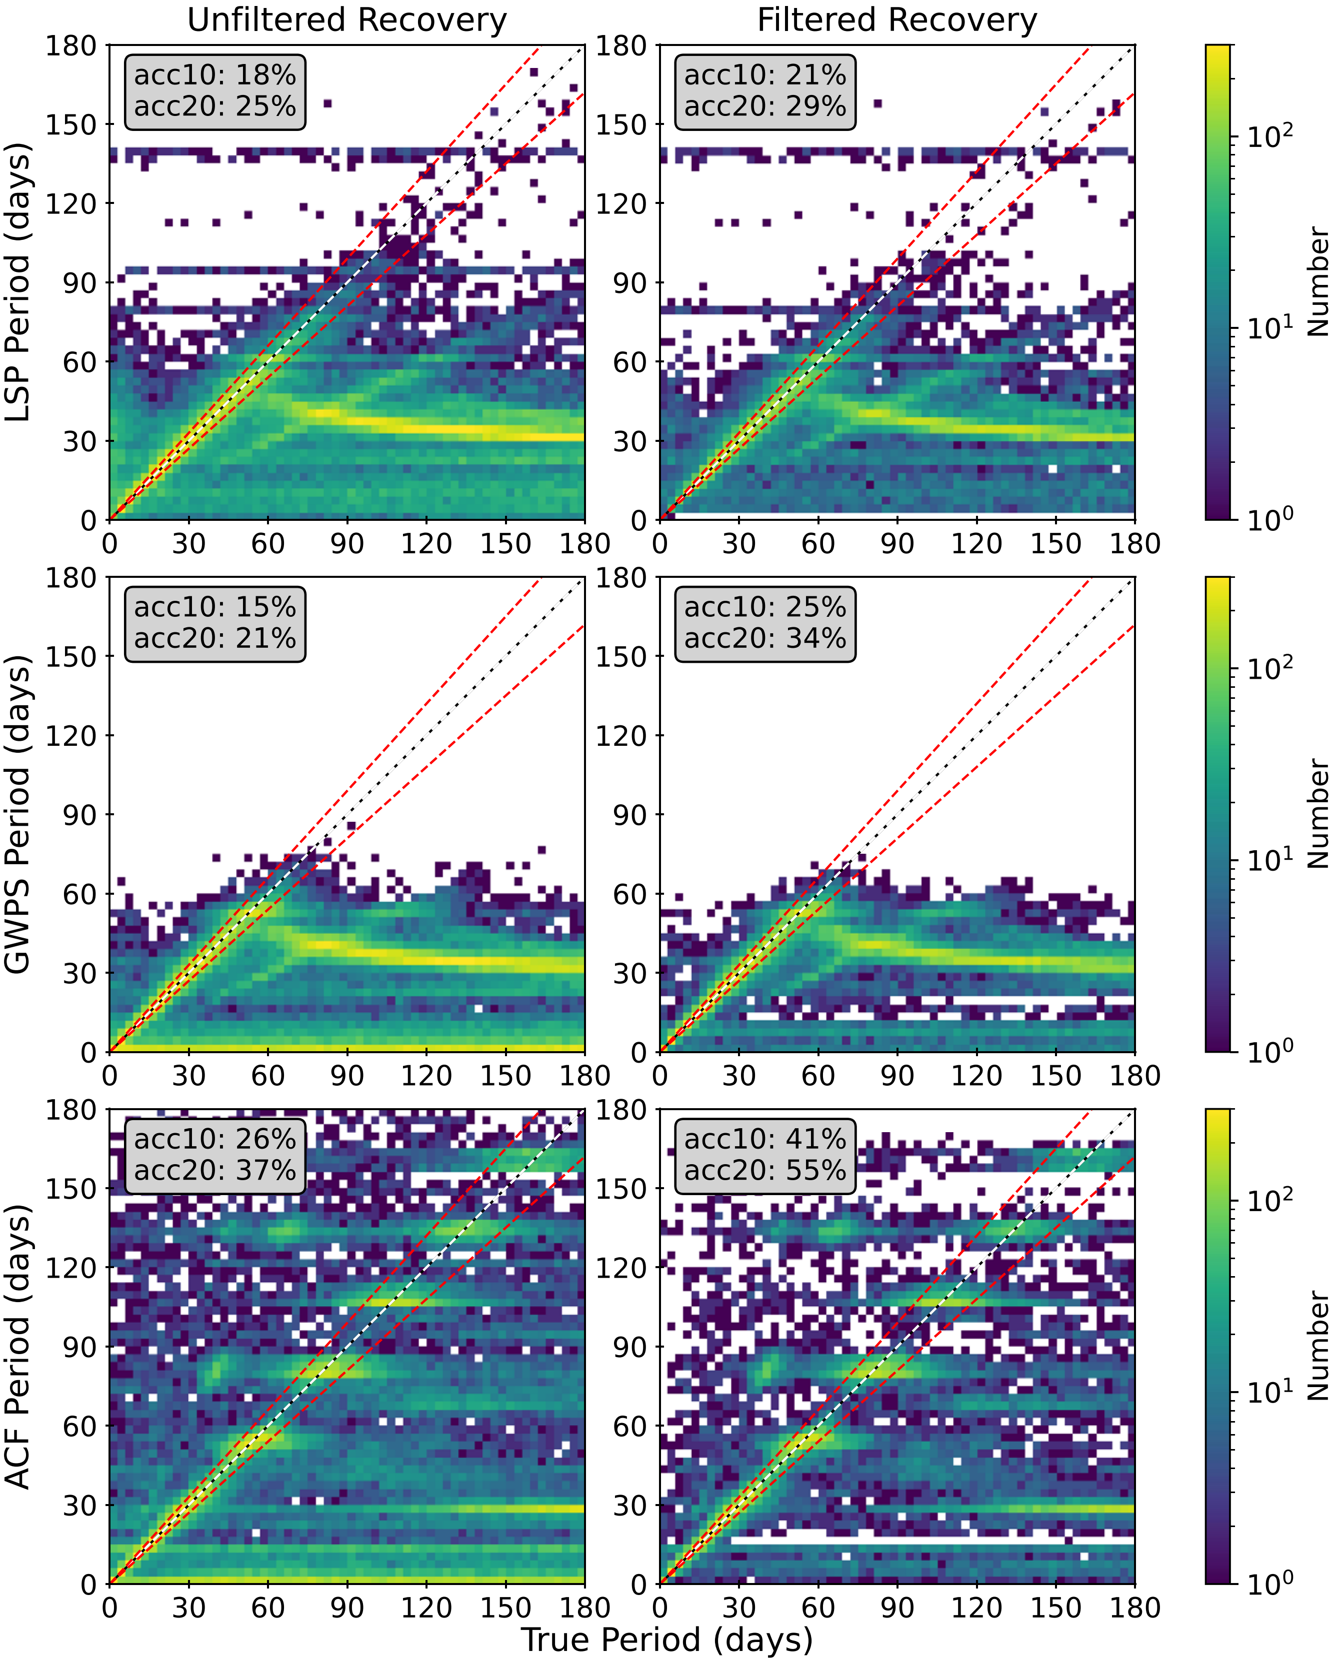

Figure 4. Period recovery using a Lomb–Scargle periodogram (LSP, top), global wavelet power spectrum (GWPS, middle), and autocorrelation function (ACF, bottom). The left column shows the unfiltered period recovery, while the right column shows recovery filtered or "precleaned" using periodogram peak significance metrics: false-alarm probability for LSP and peak prominence for GWPS and ACF. "acc10" represents fraction of periods recovered to within 10% accuracy, while "acc20" is the recovery to within 20%. ACF has the highest overall success, but the recovery worsens significantly at periods longer than about 30 days. Precleaning the sample based on peak significance improves recovery, but only marginally in most cases; aliases still dominate the predictions for true periods longer than about 50 days.

Table 3
Metrics of Period Recovery on Simulated Light Curves

	$P_{\max} = 180$ days				$P_{\max} = 50$ days			
Method	MAPE (%)	MedAPE (%)	acc10 (%)	acc20 (%)	MAPE (%)	MedAPE (%)	acc10 (%)	acc20 (%)
LSP (noiseless)	160	11	48	66	150	7	64	83
GWPS (noiseless)	52	8	56	76	52	6	67	85
ACF (noiseless)	88	6	63	81	32	5	70	88
CNN (noiseless)	14	7	61	81	11	5	69	86
LSP (noisy, unfiltered)	93	64	18	25	146	12	45	61
GWPS (noisy, unfiltered)	69	73	15	21	59	14	43	56
ACF (noisy, unfiltered)	94	53	26	37	96	16	42	54
CNN (noisy, unfiltered)	246	24	28	45	149	12	44	64
LSP (noisy, filtered)	63	54	21	29	54	8	56	73
GWPS (noisy, filtered)	50	51	25	34	27	6	63	80
ACF (noisy, filtered)	47	15	41	55	21	6	66	83
CNN (noisy, filtered)	57	11	46	69	11	7	63	86

Note. Recovery metrics for both the full 0.1–180 day period set, and for the subset with $P_{\rm rot} \leqslant 50$ days. All methods perform better on shorter-period stars, and filtering the period predictions by peak significance (or predicted uncertainty in the CNN case) improves recovery substantially for all methods. The CNN performs just as well as the conventional methods for periods less than 50 days. While for longer periods it appears the ACF outperforms the CNN, the ACF statistics are artificially inflated because nearly all the ACF's long-period "successes" occur where the rotation period coincides with an alias from the TESS sector length (see Figure 4). Thus, the CNN approach consistently outperforms the conventional techniques on simulated light curves with real TESS systematics.

Also noteworthy is that the performance of the neural network worsens at both extremes of the period range. Period recovery worsens at periods longer than about 165 days, but also at periods shorter than 13 days where conventional period-measuring methods succeed. As we discuss further in Section 6.1, this is due to the choice of loss function and its impact on predictions with uncertainty.

In addition to the rotation period, our choice of loss function allows us to predict the period uncertainty, which serves as a metric for how well the network is predicting the period. The right panel of Figure 6 shows the predicted uncertainty versus period and amplitude. The predicted uncertainty, like the recovery rate, is better at higher amplitudes. Since the predicted uncertainty correlates with the recovery rate, the predicted uncertainty is a reliable metric for successful period recovery without already knowing the period. We can then use the predicted uncertainty to select a part of the sample recovered to a desired accuracy.

For our analysis, we selected the half of the test set with the lowest predicted fractional uncertainty. The median predicted uncertainty for the sample before the cut was $\sigma_{\rm pred}/P_{\rm pred} = 0.35$. The period recovery for the best-predicted half of the sample is shown in Figure 7. The cut removed 90% of the spurious 90 day period detections, and all summary statistics were improved. 46% of periods were correctly predicted to within 10%, and 69% were accurate to within 20%. The predicted periods had a mean absolute percentage error of 57% and a median absolute percentage error of 11%. A few targets with incorrectly predicted periods between 100 and 150 days remained after the cut. These had low predicted fractional uncertainty due to their large predicted periods, so they made the cut despite being poorly predicted. They accounted for about 4% of the sample after the cut.

As with all other methods, the CNN performed better on the noisy data when limited to targets with periods less than 50 days. For this subset of the sample, the median predicted fractional uncertainty was $\sigma_{\rm pred}/P_{\rm pred} = 0.2$. Making the same cut as before (using the median fractional uncertainty of 0.2), the recovery of short-period stars improved to acc10 of 58% and median absolute percentage error of 8%. Table 3 shows the complete summary of our recovery results.

6. Discussion

We have demonstrated that CNNs are capable of extracting period information from noisy light curves or, more precisely, transformations of noisy light curves. Our model also predicts the uncertainty in the period estimate, enabling us to see where the network is most successful and determine which period predictions are most reliable. Here we discuss the strengths and weaknesses of our approach and compare them to those of conventional period detection methods. We then comment on the prospects of estimating rotation periods from real TESS light curves.

6.1. Strengths and Weaknesses of the Deep-learning Approach

We note that while `butterpy` includes the evolution of uniformly dark spots, it does not include bright plage or faculae, which are known to contribute significantly to the light curves of some active stars (Gondoin 2008). At this time `butterpy` also does not include nonuniformity of spot intensity, which is observed in sunspots (Hathaway 2015). While these are features we eventually wish to include in the `butterpy` models, the spot models in their current form should capture most of the variance seen in stellar rotational light curves. We do not believe the lack of these additional features significantly impacts our ability to recover rotation periods with deep learning.

Approximating nonuniform spots as uniformly dark circles may make subtle changes to the shape of the light curve, but the period information contained in the wavelet decomposition should be preserved. Similarly, we expect plage and faculae to affect the light-curve shape, but not the overall period information content of the wavelet transform, unless the bright spots perfectly cancel all dark spots on the stellar surface. This

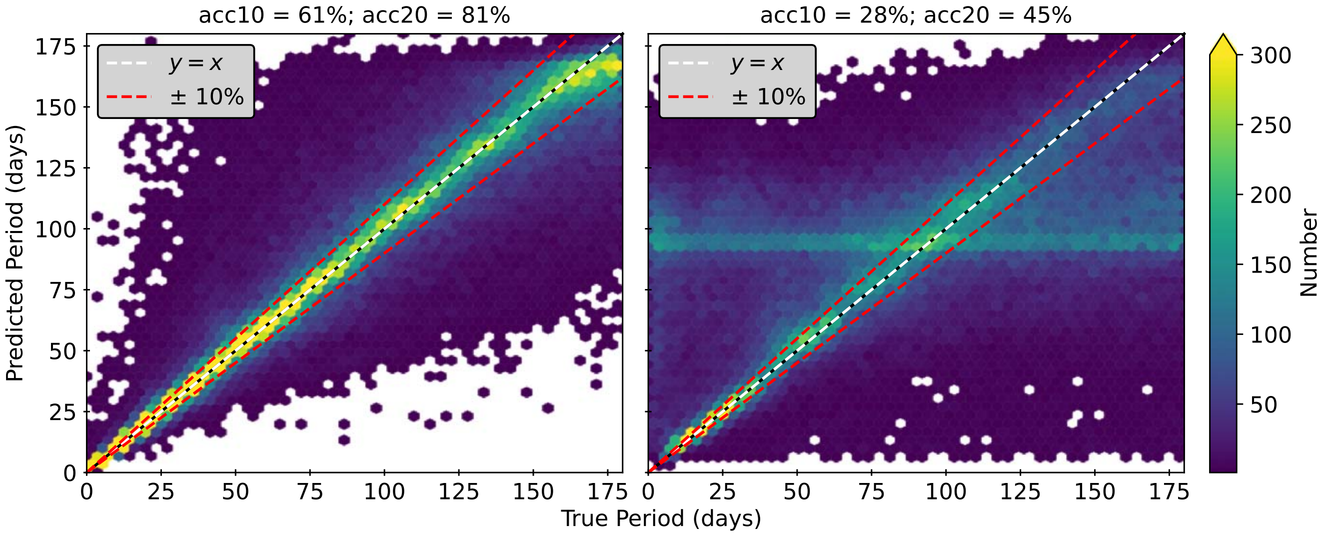

Figure 5. Left: period predictions by our CNN trained on wavelet transforms of noiseless light curves. Predicted periods have a mean absolute percentage error of 14%, median absolute percentage error of 7%, acc10 of 61%, and acc20 of 81%. Right: period predictions from noisy data, where recovery is significantly worse. Predicted periods have mean absolute percentage error of 246%, median absolute percentage error of 24%, acc10 of 28%, and acc20 of 45%. The horizontal band at 90 days represents targets where the model struggled to predict the period. In these cases, the prediction was biased toward the distribution median, or 90 days.

may actually occur in a small subset of stars (Reinhold et al. 2019), but for now we leave that investigation to future work.

Our CNN outperformed conventional techniques in the recovery of rotation periods from synthetic light curves. Whereas the conventional methods failed to recover periods longer than ∼2 TESS sectors, our method successfully recovered simulated star periods across the full simulation range of 0.1–180 days. Simulated stars in the range of periods yet unprobed by TESS—13.7 days up to 90 days and beyond—were recovered with the highest success rate. The recovery rate trails off at each end of the range ($P_{\rm rot} < 10$ days and $P_{\rm rot} > 170$ days) because of the choice of loss function: predicting the period with uncertainty biases predictions toward the median of the ensemble distribution and away from the ends of the range.

The challenge for classic period-recovery methods in TESS light curves is mostly due to sector-to-sector stitching and the presence of scattered Earth/Moon light (repeating every 27 and 13.7 days, respectively). Other effects such as temperature changes and momentum dumps appear at periods of 1.5, 2, 2.5, 3, 5, and 13.7 days (Vanderspek et al. 2018), all of which are clearly visible in Figure 4. All these effects combine to leave periodic imprints in the data that dominate stellar rotation signals and are difficult to remove. All three of our conventional method tests significantly misidentified 27 days as the rotation period. Different methods latch onto different signals, as well. For example, ACF has significant misidentifications at 2.5 and 13.7 days and a weak twice-period alias, while WPS mistakes 1 and 5 day signals as the rotation period. LSP mistakes these high-frequency signals less often, but often falls prey to half-period aliases, as does WPS.

It is noteworthy that, unlike with LSP and WPS, our neural network has no significant misidentification of half-period aliases, double-dipper stars (Basri & Nguyen 2018), or high-frequency systematic aliases. This is especially encouraging since we use WPS as the basis for our training data. The fact that these aliases certainly exist in the training set but are not chosen as the period reveals the biggest strength of the CNN approach: neural networks can learn and bypass systematic and false-period signals to successfully regress rotation periods. Another strength of deep learning is the ability to prefilter results using predicted uncertainties. If the rotation period is ambiguous, the network predicts a large uncertainty, allowing us to disregard the prediction. With the ability to regress long periods from TESS light curves in spite of systematics, we make a significant step toward large-scale stellar rotation studies with TESS.

6.2. Comparisons to Other Period Recovery Attempts

Our results are on par with or better than other recent attempts to estimate >13 day rotation periods from TESS light curves. Canto Martins et al. (2020) used a combination of fast Fourier transform, Lomb–Scargle, and wavelet techniques to estimate periods for 1000 TESS objects of interest. They obtained unambiguous rotation periods for 131 stars, but all were shorter than the 13.7 day TESS orbital period.

Lu et al. (2020) trained a random forest (RF) regressor to predict rotation periods from 27 day sections of Kepler light curves coupled with Gaia stellar parameters. They then evaluated the trained model on single sectors of TESS data for the same stars. Despite the stark differences in light-curve systematics, they were able to recover rotation periods up to ∼50 days with 55% accuracy. Their accuracy is on par with the 57% overall mean error achieved by our deep-learning model. However, when we limit our sample to simulated stars with periods less than 50 days, our mean error is significantly better at 11%. There are caveats to this comparison. First, the RF regressor relied primarily on the light-curve variability amplitude and secondarily on effective temperature; light-curve periodicity was not used for their period regression. Second, Lu et al. (2020) used 2 minute cadence, pre-search data conditioned simple aperture photometry (PDCSAP) TESS light curves, while our light curves were 30 minute cadence, and our processing pipeline was more similar to simple aperture photometry (SAP). PDCSAP light curves are subjected to much heavier detrending than those produced by SAP methods. Finally, Lu et al. (2020) used real TESS data, while we used simulated light curves. Because the two sets comprise different underlying distributions of rotation period and amplitude, they

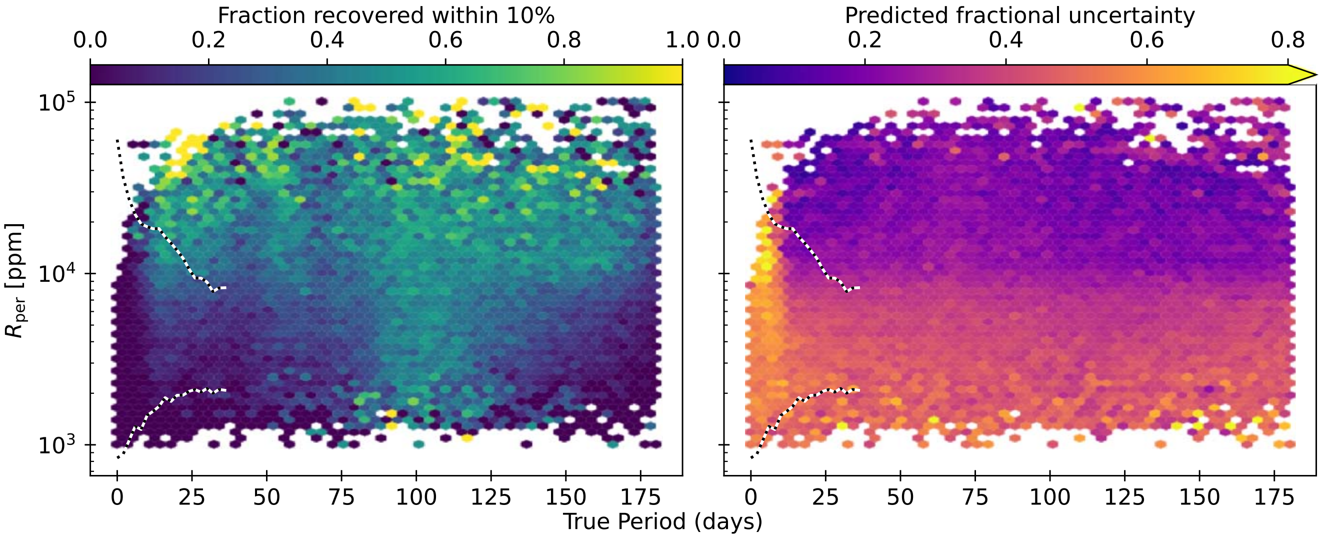

Figure 6. Neural network performance across the full simulation space of periods and amplitudes. In both panels, the dotted lines represent the 10th and 90th percentiles of the distributions from McQuillan et al. (2014), to gauge where stars from Kepler would fall. Left: period recovery rate as a function of period and amplitude for the noisy data. The neural network performs better at higher amplitudes, where the rotation signal overpowers instrumental noise. The high-amplitude failed recoveries (i.e., the darkest blue bins with $R_{\rm per} > 30{,}000$ ppm) are from light curves whose galaxy noise templates had an extremely noisy sector. The noise dominated the light curve and wavelet transform, resulting in mismeasured amplitude and period. These noise templates will be removed in future work. Right: the same data, now colored by the neural-network-predicted fractional uncertainty in rotation period. The prediction is more certain for higher amplitudes. Furthermore, the prediction is most certain in the region with the highest recovery rate, indicating the predicted uncertainty is a reliable metric for period recovery without already knowing the true period.

likely have different recovery probabilities under the same method.

6.3. *Prospects for Measuring Periodicity in TESS*

We have so far demonstrated the ability to recover photometric rotation periods from simulated TESS-like stellar light curves using deep learning. But the biggest question remains: can we reliably measure long periods from real TESS data?

This is a difficult question to answer definitively for several reasons. First, validation of any method requires a set of real stars for which rotation periods are already known. The ideal data set for comparison is Kepler, where tens of thousands of periods have been recorded (McQuillan et al. 2014; Santos et al. 2019). Unfortunately, the overlap between TESS and Kepler is small: most Kepler stars were observed for only a single sector at a time in TESS. With only a 27 day baseline, it is impossible to validate a method of obtaining long periods. Stars in the TESS CVZs were monitored continuously for almost a year, but only a handful of these stars have previously known rotation periods.

Despite the limitations, we attempted to recover rotation periods for a handful of stars observed by Kepler, the Kilodegree Extremely Little Telescope survey (KELT; Pepper et al. 2007), the MEarth Project (Berta et al. 2012), and the All-Sky Automated Survey for Supernovae (ASAS-SN; Shappee et al. 2014; Kochanek et al. 2017).

6.3.1. *Kepler*

We targeted about 1000 stars in the Kepler field that had prior rotation period measurements from McQuillan et al. (2014) as well as two consecutive sectors in TESS, offering a baseline of roughly 50 days. We simulated an entirely new training set with periods spanning 0.1–50 days, using a sample of galaxies in the Kepler field as the noise model. With a 50 day baseline, only periods of up to 25 days might be recoverable, as timescales longer than this may be dominated by edge effects in the wavelet transform. Our CNN was unable to recover Kepler periods reliably. Even when filtering based on predicted uncertainty, the recovery was no better than a random draw from the period distribution. This suggests the CNN/wavelet approach is strongest when multiple (i.e., more than two) contiguous sectors of data are available.

6.3.2. *Kilodegree Extremely Little Telescope Survey*

We similarly targeted 106 stars in the TESS SCVZ also observed by the KELT survey. Oelkers et al. (2018) obtained rotation periods for these stars using Lomb–Scargle periodograms of their KELT light curves. We specifically selected stars with a measured period greater than 13.7 days to test recovery of long periods. To maximize the chances of recovering rotation periods, we used the TESS Science Processing Operations Center (SPOC) FFI SAP light curves (Caldwell et al. 2020). At the time of writing, only sectors 1–6 were available, but we trained our CNN using year-long (13 sector) light curves. However, the construction of our wavelet power spectrum used the same vertical (frequency) axis regardless of light-curve length, so any length of light curve could be used without needing to retrain the neural network.

Upon visual inspection of the TESS-SPOC light curves, we noticed that many did not show obvious rotational modulation. Selecting only those light curves with unambiguous rotational modulation, we were left with 26 light curves with KELT rotation periods spanning 13.7–47 days. While visually inspecting the light curves, we also ensured the apparent rotational modulation was consistent with the periods obtained by Oelkers et al. (2018). We generated wavelet power spectra and evaluated our neural network on the light curves. Figure 8 shows our predictions for these 26 KELT stars.

It is interesting that the rotation periods tend to clump around 17, 27, 37, and 47 days. These periods are not associated with

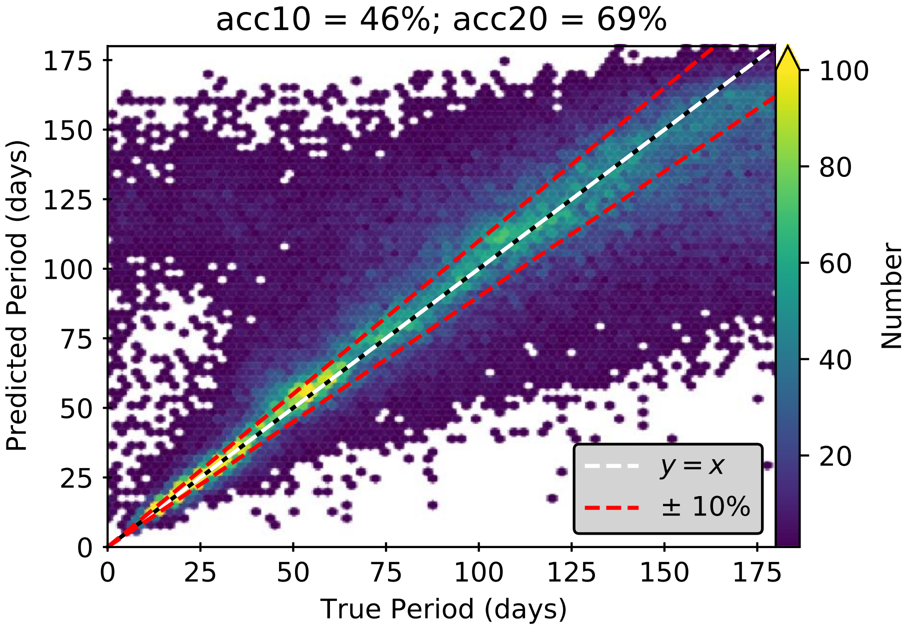

Figure 7. Period recovery for the half of the test set with the lowest predicted fractional uncertainty. Predicted periods have a mean absolute percentage error of 57%, median absolute percentage error of 11%, acc10 of 46%, and acc20 of 69%. The predicted error cut removes the cluster of predicted periods at 90 days, giving credence to our cut to remove spurious period predictions. The cloud of objects with short true periods and predicted periods between 100 and 150 days have low fractional error because their predicted periods are large, but they account for only 4% of the objects remaining after the cut.

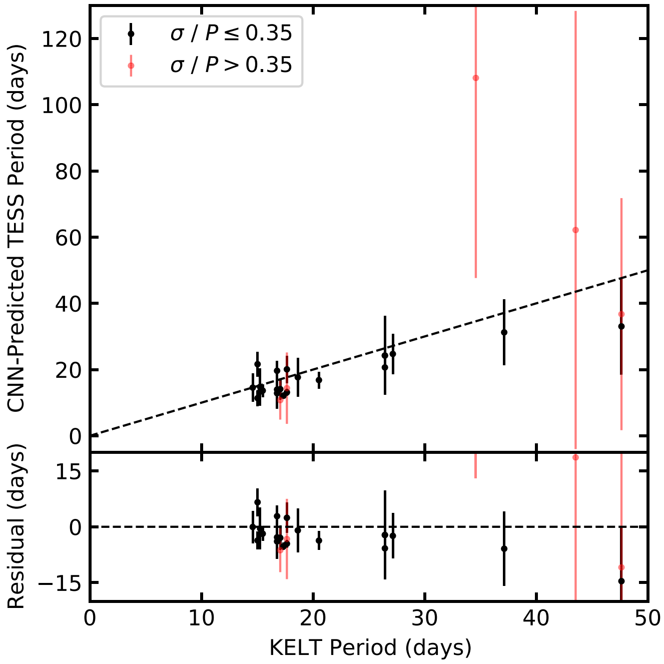

Figure 8. Period recovery of stars in both TESS and KELT for which rotational modulation is apparent in the light curve. We applied the same fractional uncertainty cut applied to the simulation recovery results; the 21 stars that made the cut are in black, while the five stars with unreliable period predictions are in red. We successfully recovered periods longer than 13.7 and even 27 days from real TESS light curves. Even when TESS aliases are the dominant sources of power in the wavelet transform, our neural network was able to recover the correct rotation period.

any of the known KELT aliases discussed by Oelkers et al. (2018), and besides 27 days, they are not associated with any TESS aliases. Furthermore, given that 17, 37, and 47 are prime, they are unlikely to be associated with a single common alias. These are more likely associated with the spacing of periods searched by Oelkers et al. (2018). Their overall period distribution seems to have more stars at these periods than at other values, so the clumping likely arises either from sensitivity of their method to those particular periods, or from the spacing of periods used in their search for rotation.

After running these KELT stars through our CNN, we subjected the predictions to the same cut in predicted fractional uncertainty as in Figure 7. Stars that made the cut are displayed in black, while stars whose predicted uncertainties were too large are shown in red. The KELT stars for which we detected periods range from 0.2% to 2% in variability amplitude. We successfully recovered stars with rotation periods longer than 13.7 and 27 days, even when TESS systematics were the dominant source of power in the wavelet power spectra. Furthermore, using the predicted fractional uncertainty as a quality cut successfully removed stars whose predictions were unreliable or wrong while ensuring the most reliable predictions remained in the sample.

6.3.3. MEarth and TOI-700

Only one long-period target was observed by MEarth in the TESS SCVZ: TIC 149423103. Newton et al. (2018) measured a rotation period of 111 days for this target from MEarth data. Using our neural network on the FFI data from TESS, we obtained $P_{\rm rot} = 116 \pm 48$ days. While the predicted period was within 5% of the “true” period, the relatively large uncertainty (41%) means this target would fail our quality cut, and an ensemble period recovery attempt would miss it.

TOI-700 is another well-characterized star in the SCVZ. Using ASAS-SN data, Gilbert et al. (2020) estimated a precise rotation period of 54.0 ± 0.8 days. Hedges et al. (2020) used a systematics-insensitive periodogram of its TESS light curve to obtain a period of 52.8 days. With our model we predicted a period of 59 ± 53 days. Our period prediction was accurate to within 10%, but the large uncertainty would cause this target to be missed, as well.

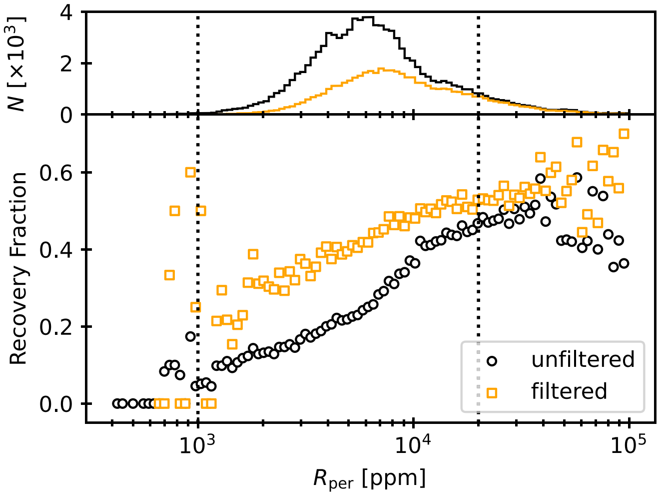

Figure 9. Top: histogram of amplitudes of synthesized light curves in our test set, both before filtering (black) and after (orange). Bottom: recovery rate as a function of amplitude for both the unfiltered (black circles) and filtered (orange squares) period predictions on the synthesized test set. The vertical dotted lines show the approximate range of amplitudes seen by McQuillan et al. (2014), where "real stars" might lie. According to the basic picture of period–amplitude relationships, slowly rotating stars will be toward the left, and fast stars will be toward the right. Applying our neural network to real stars, we expect to detect fewer low-amplitude (slow rotating) stars than high-amplitude (fast rotating) stars, but filtering on predicted uncertainty somewhat mitigates the disparity. The large scatter in recovery fraction of the filtered data at $R_{\rm per} \sim 10^3$ ppm is due to the very small number of objects in these bins. In this range, the unfiltered data bins have an order of magnitude more objects than the filtered data bins (hundreds in the unfiltered versus tens in the filtered).

6.3.4. General Period Recovery and Improvements

While we successfully recovered the rotation periods of these few hand-picked stars, robust recovery on larger, statistical scales will require more work and vetting. Our method allows us to see beyond the 13.7 day barrier, but mostly stars with large amplitudes were reliably recovered: the detection fraction in the synthesized test set picks up around 1% amplitude, and the successfully recovered KELT stars range from 0.2% to 2% in amplitude. We are still limited by the TESS noise floor. TESS is less precise than Kepler at all magnitudes (Vanderspek et al. 2018), so spot modulations require higher amplitudes to rise above the noise. Both panels of Figure 6 show the 10th and 90th percentile variability amplitudes of the rotating stars from McQuillan et al. (2014). The bulk of Kepler's rotating population fall between $R_{\rm per}$ of 1 and 20 parts per thousand and lie in a region where our recovery was less successful. In addition, slowly rotating stars are typically old and therefore less active than their young, rapidly spinning counterparts. The difficulty of recovering periods at smaller amplitudes will adversely affect our ability to detect these slow rotators. These kinds of stars will be more difficult to recover with TESS, whatever the method. However, another view suggests recovery of real stars may not be so farfetched. Figure 9 shows the period recovery rate for the synthesized test set as a function of $R_{\rm per}$, the light-curve modulation amplitude. It is essentially the left panel of Figure 6 collapsed onto the vertical axis. The vertical dotted lines show the range of amplitudes for the Kepler stars of McQuillan et al. (2014). While the recovery is worse at lower amplitudes, filtering by predicted uncertainty improves the recovery rate at all amplitudes and somewhat lessens the disparity between high and low amplitudes. Without improvements to our method, given a log-uniform distribution of amplitudes in the range seen in Kepler, we expect to recover low-amplitude stars about half as frequently as high-amplitude stars.

Still, we believe improvements to our method will maximize what is recoverable from TESS. There are several ways to extend the predictability of our neural network to lower amplitudes and enhance the predictability at high amplitudes. The first and perhaps most useful improvement will come through light-curve processing. Our processing pipeline followed the regression-corrector documentation of the `Lightkurve` Collaboration (2020) using a magnitude-dependent aperture threshold. In practice, a more carefully developed pipeline should be preferred. At the time of writing, the FFIs of sectors 1–6 have been reduced by both the TESS Asteroseismic Science Operations Center pipeline and the TESS-SPOC pipeline. Once sectors 7–13 are processed, the SCVZ will be complete, providing year-long light curves for hundreds of thousands of targets. These light curves will feature more careful systematics removal and should contain cleaner examples to use as "pure noise" light curves in our sample. We leave the use of these light curves to a future paper.

Another improvement may come with the inclusion of observation metadata. In TESS data, certain systematic effects are specific to particular cameras, CCDs, or even locations on a CCD. We could include camera number, CCD number, and x and y pixel coordinates in the training data set, perhaps as extra input nodes to be included alongside the input wavelet transforms. This will allow neural networks to learn where to expect certain features and more easily ignore them in favor of astrophysical signals.

In this work we have assumed no period–amplitude relationship in our simulated training set. This was to avoid imprinting strong priors on the predictions of the neural network. In future efforts, we may see improvements from incorporating period–activity relations in the underlying training set. These might serve to enhance the reliability of predictions, as amplitudes may help inform the network of whether a star is rotating slowly or quickly. We might also see improvements in training efficiency as the network learns more rapidly what short-period and long-period wavelet transforms look like.

Improvements can be made to the neural network, as well. In its current form, our model assumes that all input signals have rotation signatures, but not all real light curves display rotational modulation. In the future we may include a classification step like Lu et al. (2020) to determine which signals contain rotational modulation. Adding this classification step will allow the regressor to focus on signals with recoverable rotation, making for more efficient training.

Based on recovery from the simulated test set, our approach is currently most effective above periods of about 13 days due to the loss function's tendency to bias uncertain predictions toward the ensemble mean. While experimenting with other loss functions may result in improved short-period recovery, preliminary tests using training sets with smaller period ranges have shown promise, as well. One might consider using two neural network regressors—one for short periods, and one for long periods—and a preceding classification step to decide which regressor to use. We leave this investigation to future work.

It is important to note that our implementations of conventional period-recovery techniques perform better on our simulated test set than the best efforts on real light curves (e.g., Canto Martins et al. 2020 and Avallone et al. 2021, who were unsuccessful in recovering anything past 13.7 days). This indicates that, despite all our attempts to create as realistic a training set as possible, our simulations are not perfectly representative of real stars. It could be that our stitching routine fails to suppress long-period signals as the real light curves do. Another possibility is that our spot model, while tuned to the Sun, may not be representative of real spots on other stars. Whatever the reason, we have demonstrated the ability to recover periods even when the systematics that *are* present in our simulations make conventional techniques fail.

Even though our spot-evolution simulations include latitudinal differential rotation, we were unable to recover differential rotation in this study. In some wavelet power spectra of our simulated light curves, the differential rotation is apparent as a slope in the frequency of maximum power versus time. When binning the power spectra to 64 × 64 pixels, the slope is more difficult to resolve. While increasing the resolution of the wavelet power spectrum images should enable the recovery of differential rotation, it will come at the expense of longer training time. We will investigate the recovery of differential rotation, activity levels, and spot properties in future work.

If we can see beyond the complicated systematics, TESS will deliver the largest set of rotation periods yet. McQuillan et al. (2014) obtained rotation periods for 34,000 stars in the Kepler field. The TESS continuous viewing zones combine to cover 900 square degrees around the ecliptic poles, representing about eight times the sky coverage of Kepler during its primary mission. We can therefore expect hundreds of thousands of new stars with rotation period estimates from the TESS CVZs, and perhaps more from lower ecliptic latitudes. Because of TESS's lower precision compared to Kepler, the true number will likely be somewhat smaller, but the prospect of hundreds of thousands of new periods is worth continued refinements of this technique. We leave the application of this tool to the full CVZ samples to a future paper.

7. Summary and Conclusion

We used a CNN to recover rotation periods and uncertainties from simulated light curves with real TESS systematics. Despite the systematics, we successfully recovered periods even for targets whose periods were longer than the 13.7 day barrier encountered by conventional period-recovery methods. In the half of the simulated test data with the smallest predicted fractional uncertainty, we recovered 10% accurate periods for 46% of the sample, and 20% accurate periods for 69% of the sample. We also found no significant misidentification of half-period aliases, unlike the Lomb–Scargle and wavelet methods. While periods were retrieved more successfully from higher-amplitude signals, the ability to predict uncertainties allows us to probe lower-amplitude rotation signals, as well.

We plan to use this method to produce a catalog of rotation periods from TESS FFI light curves. We will also add output options to our neural network to predict latitudinal differential rotation and understand more of the properties used to produce the training set. With deep learning, we hope to maximize the output of TESS in spite of the frustrations that arise from its systematics. The ability to recover rotation periods, especially long periods, from TESS data will finally enable large studies of rotation across diverse populations of stars in the Galaxy if only the systematics can be learned.

We thank the anonymous reviewer for thoughtful and useful comments that improved the quality and results of this work.

We also wish to acknowledge Gagandeep Anand, Connor Auge, Ashley Chontos, Aidan Chun, Curt Dodds, Ryan Dungee, Kyle Hart, Jason Hinkle, Rae Holcomb, Daniel Huber, Miles Lucas, Sushant Mahajan, Anna Payne, Nicholas Saunders, Jessica Schonhut-Stasik, Benjamin Shappee, Xudong Sun, and Jamie Tayar for fruitful conversations that improved the quality of this work.

The technical support and advanced computing resources from the University of Hawaii Information Technology Services Cyberinfrastructure are gratefully acknowledged.

This research was supported in part by the National Science Foundation under grant No. NSF PHY-1748958.

J.v.S. and Z.R.C. acknowledge support from the National Aeronautics and Space Administration (grant Nos. 80NSSC21K0246, 80NSSC18K18584).

J.L. acknowledges support from NASA through an Astrophysics Data Analysis Program grant to Lowell Observatory (grant No. 80NSSC20K1001).

This paper includes data collected by the TESS mission. Funding for the TESS mission is provided by the NASA's Science Mission Directorate.

Software: `NumPy` (Harris et al. 2020), `Pandas` (McKinney 2010; The pandas development team 2020), `Matplotlib` (Hunter 2007), `AstroPy` (Astropy Collaboration et al. 2013, 2018), `SciPy` (Virtanen et al. 2020), `PyTorch` (Paszke et al. 2019), `Lightkurve` (Lightkurve Collaboration et al. 2018), `TESScut` (Brasseur et al. 2019), `iPython` (Perez & Granger 2007), `butterpy` (Claytor et al. 2021), `starspot` (Angus 2021).

ORCID iDs

Zachary R. Claytor https://orcid.org/0000-0002-9879-3904
Jennifer L. van Saders https://orcid.org/0000-0002-4284-8638
Joe Llama https://orcid.org/0000-0003-4450-0368
Peter Sadowski https://orcid.org/0000-0002-7354-5461
Brandon Quach https://orcid.org/0000-0001-9848-7483
Ellis A. Avallone https://orcid.org/0000-0003-1719-5046

References

Aigrain, S., Pont, F., & Zucker, S. 2012, MNRAS, 419, 3147
Aigrain, S., Llama, J., Ceillier, T., et al. 2015, MNRAS, 450, 3211
Amard, L., Roquette, J., & Matt, S. P. 2020, MNRAS, 499, 3481
Angus, R. 2021, starspot: code for measuring stellar rotation periods, v0.2, Zenodo, doi: 10.5281/zenodo.4613887
Angus, R., Morton, T., Aigrain, S., Foreman-Mackey, D., & Rajpaul, V. 2018, MNRAS, 474, 2094
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Avallone, E. A., Tayar, J., Van Saders, J., Berger, T., & Claytor, Z. 2021, AAS Meeting, 238, 314.07
Baglin, A., Auvergne, M., Barge, P., et al. 2006, in ESA Special Publication 1306, The CoRoT Mission Pre-Launch Status—Stellar Seismology and Planet Finding, ed. M. Fridlund et al. (Paris: ESA), 33
Bao, W., Yue, J., & Rao, Y. 2017, PLoSO, 12, e0180944
Basri, G., & Nguyen, H. T. 2018, ApJ, 863, 190
Basri, G., Walkowicz, L. M., Batalha, N., et al. 2011, AJ, 141, 20
Bazot, M., Nielsen, M. B., Mary, D., et al. 2018, A&A, 619, L9
Berta, Z. K., Irwin, J., Charbonneau, D., Burke, C. J., & Falco, E. E. 2012, AJ, 144, 145

Blancato, K., Ness, M., Huber, D., Lu, Y., & Angus, R. 2020, arXiv:2005.09682

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Sci, 327, 977

Brasseur, C. E., Phillip, C., Fleming, S. W., Mullally, S. E., & White, R. L. 2019, Astrocut: Tools for creating cutouts of TESS images, Astrophysics Source Code Library, http://ascl.net/1905.007

Caldwell, D. A., Tenenbaum, P., Twicken, J. D., et al. 2020, RNAAS, 4, 201

Canto Martins, B. L., Gomes, R. L., Messias, Y. S., et al. 2020, ApJS, 250, 20

Ceillier, T., Tayar, J., Mathur, S., et al. 2017, A&A, 605, A111

Claytor, Z. R., Lucas, M., & Llama, J. 2021, Butterpy: realistic star spot evolution and light curves in Python, v0.1.0, Zenodo, doi:10.5281/zenodo.4722052

Claytor, Z. R., van Saders, J. L., Santos, Â. R. G., et al. 2020, ApJ, 888, 43

Cranmer, K., Brehmer, J., & Louppe, G. 2020, PNAS, 117, 30055

Davenport, J. R. A. 2017, ApJ, 835, 16

Feiden, G., Guinan, E., Boyajian, T., et al. 2011, AAS Meeting, 217, 140.18

Feinstein, A. D., Montet, B. T., Ansdell, M., et al. 2020, AJ, 160, 219

Flesch, E. W. 2015, PASA, 32, e010

García, R. A., Ceillier, T., Salabert, D., et al. 2014, A&A, 572, A34

Gilbert, E. A., Barclay, T., Schlieder, J. E., et al. 2020, AJ, 160, 116

Gondoin, P. 2008, A&A, 478, 883

Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep Learning (Cambridge, MA: MIT Press)

Guiglion, G., Matijevic, G., Queiroz, A. B. A., et al. 2020, A&A, 644, A168

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357

Hathaway, D. H. 2011, SoPh, 273, 221

Hathaway, D. H. 2015, LRSP, 12, 4

Hathaway, D. H., Wilson, R. M., & Reichmann, E. J. 1994, SoPh, 151, 177

He, K., Zhang, X., Ren, S., & Sun, J. 2015, arXiv:1512.03385

Hedges, C., Angus, R., Barentsen, G., et al. 2020, RNAAS, 4, 220

Hezaveh, Y. D., Perreault Levasseur, L., & Marshall, P. J. 2017, Natur, 548, 555

Holcomb, R. 2020, AAS Meeting, 235, 274.04

Hunter, J. D. 2007, CSE, 9, 90

Kingma, D. P., & Ba, J. 2014, Proc. of the 3rd Int. Conf. on Learning Representations, ICLR 2015 https://dblp.org/db/conf/iclr/iclr2015.html

Kochanek, C. S., Shappee, B. J., Stanek, K. Z., et al. 2017, PASP, 129, 104502

Lightkurve Collaboration 2020, How to remove scattered light from TESS data using the `RegressionCorrector`?https://docs.lightkurve.org/tutorials/2-creating-light-curves/2-3-removing-scattered-light-using-regressioncorrector.html

Lightkurve Collaboration, Cardoso, J. V. D. M., Hedges, C., et al. 2018, Lightkurve: Kepler and TESS time series analysis in Python, Astrophysics Source Code Library, http://ascl.net/1812.013

Liu, Y., San Liang, X., & Weisberg, R. H. 2007, JAtOT, 24, 2093

Llama, J., Jardine, M., Mackay, D. H., & Fares, R. 2012, MNRAS, 422, L72

Lomb, N. R. 1976, Ap&SS, 39, 447

Lu, Y. L., Angus, R., Agüeros, M. A., et al. 2020, AJ, 160, 168

Mackay, D. H., Jardine, M., Collier Cameron, A., Donati, J. F., & Hussain, G. A. J. 2004, MNRAS, 354, 737

Marilli, E., Frasca, A., Covino, E., et al. 2007, A&A, 463, 1081

McKinney, W. 2010, in Proc. of the 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 56

Mathur, S., García, R. A., Régulo, C., et al. 2010, A&A, 511, A46

McQuillan, A., Aigrain, S., & Mazeh, T. 2013, MNRAS, 432, 1203

McQuillan, A., Mazeh, T., & Aigrain, S. 2014, ApJS, 211, 24

Netto, Y., & Valio, A. 2020, A&A, 635, A78

Newton, E. R., Mondrik, N., Irwin, J., Winters, J. G., & Charbonneau, D. 2018, AJ, 156, 217

Nielsen, M. B., Gizon, L., Cameron, R. H., & Miesch, M. 2019, A&A, 622, A85

Oelkers, R. J., & Stassun, K. G. 2018, AJ, 156, 132

Oelkers, R. J., Rodriguez, J. E., Stassun, K. G., et al. 2018, AJ, 155, 39

Paszke, A., Gross, S., Massa, F., et al. 2019, in Advances in Neural Information Processing Systems 32, ed. H. Wallach et al. (Red Hook, NY: Curran Associates Inc.), 8024

Pepper, J., Pogge, R. W., DePoy, D. L., et al. 2007, PASP, 119, 923

Perez, F., & Granger, B. E. 2007, CSE, 9, 21

Reinhold, T., Bell, K. J., Kuszlewicz, J., Hekker, S., & Shapiro, A. I. 2019, A&A, 621, A21

Reinhold, T., & Hekker, S. 2020, A&A, 635, A43

Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, JATIS, 1, 014003

Rüdiger, G., Küker, M., Käpylä, P. J., & Strassmeier, K. G. 2019, A&A, 630, A109

Santos, A. R. G., García, R. A., Mathur, S., et al. 2019, ApJS, 244, 21

Scargle, J. D. 1982, ApJ, 263, 835

Schrijver, C. J., & Harvey, K. L. 1994, SoPh, 150, 1

Shappee, B. J., Prieto, J. L., Grupe, D., et al. 2014, ApJ, 788, 48

The pandas development team 2020, pandas-dev/pandas: Pandas, Zenodo, doi:10.5281/zenodo.3509134

Thomas, A. E. L., Chaplin, W. J., Davies, G. R., et al. 2019, MNRAS, 485, 3857

Torrence, C., & Compo, G. 1998, BAMS, 79, 61

van Ballegooijen, A. A., Cartledge, N. P., & Priest, E. R. 1998, ApJ, 501, 866

van Saders, J. L., Ceillier, T., Metcalfe, T. S., et al. 2016, Natur, 529, 181

van Saders, J. L., Pinsonneault, M. H., & Barbieri, M. 2019, ApJ, 872, 128

Vanderspek, R., Doty, J. P., Fausnaugh, M., et al. 2018, TESS Instrument Handbook, (Baltimore, MD: MAST) https://archive.stsci.edu/missions/tess/doc/TESS_Instrument_Handbook_v0.1.pdf

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261

Zhao, Q., & Zhang, L. 2005, in Int. Conf. on Neural Networks and Brain (New York: IEEE), 1089