





Computational tools for the spectroscopic analysis of white dwarfs

Vedant Chandra ¹★, Hsiang-Chih Hwang ¹, Nadia L. Zakamska ¹ and Tamás Budavári ^{1,2}

¹Department of Physics and Astronomy, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA

²Department of Applied Mathematics and Statistics, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA

Accepted 2020 July 21. Received 2020 July 20; in original form 2020 June 17

ABSTRACT

The spectroscopic features of white dwarfs are formed in the thin upper layer of their stellar photosphere. These features carry information about the white dwarf’s surface temperature, surface gravity, and chemical composition (hereafter ‘labels’). Existing methods to determine these labels rely on complex ab-initio theoretical models, which are not always publicly available. Here, we present two techniques to determine atmospheric labels from white dwarf spectra: a generative fitting pipeline that interpolates theoretical spectra with artificial neural networks and a random forest regression model using parameters derived from absorption line features. We test and compare our methods using a large catalogue of white dwarfs from the Sloan Digital Sky Survey (SDSS), achieving the same accuracy and negligible bias as compared to previous studies. We package our techniques into an open-source PYTHON module ‘WDTOOLS’ that provides a computationally inexpensive way to determine stellar labels from white dwarf spectra observed from any facility. We will actively develop and update our tool as more theoretical models become publicly available. We discuss applications of our tool in its present form to identify interesting outlier white dwarf systems including those with magnetic fields, helium-rich atmospheres, and double-degenerate binaries.

Key words: techniques: spectroscopic – white dwarfs.

1 INTRODUCTION

Deriving precision stellar parameters for white dwarfs, primarily their effective temperature T_{eff} and gravity $\log g$, has been subject of intense theoretical effort by multiple groups (Hubeny 1988; Fontaine, Brassard & Bergeron 2001; Koester 2008; Tremblay & Bergeron 2009; Tremblay et al. 2013). These labels are of immense interest for several areas of astrophysics. For example, a lower bound on the age of the Universe can be derived by modelling cooling tracks of the oldest white dwarfs in globular clusters (Hansen et al. 2002). Another application is the search for near- and super-Chandrasekhar close binaries (Brown et al. 2013) and for massive white dwarf merger products which can constrain the double-degenerate scenario of type Ia progenitors (Maoz, Mannucci & Nelemans 2014).

For hydrogen-rich (DA) white dwarfs, some of the most accurate determinations of white dwarf temperatures and surface gravities are made by fitting theoretical spectra to absorption lines in the Balmer series of the optical range (Fig. 1). The widths and equivalent widths of the profiles are sensitive, both to temperature and gravity. Comparisons between spectroscopic and photometric methods (Tremblay et al. 2019) indicate that for a wide range of temperatures, the spectroscopic temperature measurements are in agreement to within 2 per cent and the surface gravity measurements are in agreement within 0.2 dex. Comparing measurements made by different groups (Tremblay et al. 2010) also yields broadly consistent results.

However, the theoretical models remain uncertain for certain situations, including low temperatures (Tremblay et al. 2013; Blouin

2020), high temperatures (Levenhagen et al. 2017), and high masses (Camisassa et al. 2019). Masses are known to be consistently overestimated at the low-temperature end (Tremblay et al. 2010) and there is continued work on improving models in this regime (Tremblay et al. 2013). At these low temperatures, $\log g$ values disagree with photometric predictions from *Gaia* (Tremblay et al. 2019). An additional hurdle is that almost all the codes used to model white dwarf atmospheres are kept under restricted access by their respective groups. Rarely, limited grids of pre-computed synthetic spectra are publicly available. One exception is the freely available TLUSTY code (Bohlin, Hubeny & Rauch 2020), although their lack of a sophisticated convection model restricts their usefulness to hot white dwarfs with $T_{\text{eff}} \gtrsim 20\,000$ K.

For main-sequence stars, machine-learning based techniques have become a promising solution to predict elemental abundances. For example, the PAYNE (Ting et al. 2018) uses a neural-network to reproduce spectral models at high spectral resolution and applies them to lower resolution spectra across a wider range of metal abundances. The CANNON (Ness et al. 2015) uses a data-driven approach to assign astrophysical labels based on stellar features. ASTRONN (Leung & Bovy 2018) uses convolutional neural networks (CNNs) directly on spectral pixels to predict stellar parameters and abundances. These computational approaches are now widely used and have become a powerful technique in astrophysical research (Nataf et al. 2019; Xiang et al. 2019).

White dwarfs are a promising albeit as-yet unexplored population for data-driven techniques. In particular, 35 000 white dwarf spectra have been measured by the Sloan Digital Sky Survey (SDSS; Kepler et al. 2019) and *Gaia* DR2 revealed a large sample of 250 000 white dwarfs (Fusillo et al. 2019). Recent studies have used follow-up spectroscopy to characterize over a thousand white dwarfs within

* E-mail: vchandra@jhu.edu




Figure 1. Three un-normalized model spectra from Koester (2010) with DA (hydrogen-rich), DB (helium-rich), and DC (continuum) spectral classifications, respectively, vertically displaced by a constant for clarity. The DA spectra have their own set of synthetic models, whilst the DB and DC spectra are taken from a common grid of models. From top to bottom, these model spectra have $(T_{\text{eff}}, \log g)$ labels (16500, 6.75), (17000, 7), and (9250, 7), respectively.

100 pc of the Sun (Kilic et al. 2020; McCleery et al. 2020; Tremblay et al. 2020). Future large-scale spectroscopic surveys like the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016) and SDSS-V (Kollmeier et al. 2017) will greatly increase the population of white dwarf spectra.

Therefore, there is an urgent need to develop tools to recover the astrophysical parameters of white dwarfs from spectroscopic data and to identify interesting white dwarfs that are inconsistent with theoretical models. Narayan et al. (2019) previously released their fitting routine for DA white dwarfs that uses the high-temperature synthetic spectra from Bohlin et al. (2020) to establish faint WD spectro-photometric standards. Their method includes a sophisticated treatment of extinction, which serves their goal of sub per cent photometric accuracy. Our work uses a markedly different methodology for model interpolation and inference, and aims to be more user-friendly and computationally efficient for application to large-scale spectroscopic surveys.

In this paper, we present an open-source software package in PYTHON called WDTOOLS that is meant to be flexible across a wide variety of use-cases, from quick predictions on large data sets to detailed statistical inference for single targets. WDTOOLS is hosted online (Zenodo: Chandra 2020) with bug-reporting capabilities from GitHub,¹ along with examples and full code documentation.²

Although the model evaluation in this paper is focused on SDSS spectra with known labels, we also show that WDTOOLS can be used on spectra obtained from any observatory. The model spectra and SDSS spectra are computed with vacuum wavelengths, and we use that convention in this work. Our tool includes the ability to switch to the air wavelength convention as well. We run all our tests on a MacOS laptop computer with a 3.1-GHz quad-core processor and provide computation time estimates for this configuration.

We define our test sample of SDSS white dwarf spectra in Section 2. We describe our generative fitting pipeline in Section 3.1

and our random forest regression model in Section 3.2. We test the accuracy of our tools on a variety of spectra in Section 4 and present applications of our tools to find interesting spectra. We summarize our results and discuss the strengths and limitations of WDTOOLS in Section 5.

2 DATA AND FEATURE EXTRACTION

2.1 Data selection

We use white dwarf spectra from the SDSS to develop and test our methods. We use stellar parameters derived by Tremblay et al. (2019) as ‘ground truth’ labels as a point of comparison. These were derived using traditional χ^2 fitting with theoretical models to determine temperature and surface gravity, and likely represent the best labels for these stars so far. Alternatives include the catalogue of Fusillo et al. (2019), but we elect to forego those labels since they are derived from *Gaia* photometric observations rather than spectroscopy. Since there is an inherent systematic scatter between photometric and spectroscopic labels (Tremblay et al. 2019), we elect to compare our spectroscopic results to labels derived from spectroscopy as well.

For this study, we focus on stars classified by Tremblay et al. (2019) as DA (hydrogen-rich) white dwarfs, which make up to 70 per cent of the SDSS sample. We also extend our method of fitting model atmospheres (Section 3.1) to DB (helium-rich) white dwarfs, which constitute a further 10 per cent of the SDSS sample. For the development of our machine learning tools (Section 3.2), we exclude those stars with deviant spectral features like magnetism (H), metallic lines (Z), and pure-continuum spectra (DC). We discuss the possibility of using our fitting tool to identify such outlier spectra in Section 4.4.

2.2 Continuum normalization

Due to the lack of absorption lines other than the hydrogen Balmer series, it is relatively straightforward to identify the continuum of a DA white dwarf spectrum. However, to prevent introducing any systematic errors into spectral fitting models, a consistent continuum-normalization technique is vital. We build upon the continuum-normalization technique of Bergeron, Liebert & Fulbright (1995) and subsequent works, which individually normalize the continuum around each Balmer absorption line with summed Gaussian profiles.

Selecting a region around each Balmer line from H α to H8 on the un-normalized spectrum, we fit a composite model of a Voigt profile added to a linear function of wavelength. The Voigt profile is the convolution of a Gaussian and Lorentzian distribution, and represents well the pressure-broadened wings of the Balmer lines (Tremblay & Bergeron 2009). We divide the region around each Balmer line by the linear component of this composite fit to derive the continuum normalized spectrum for that line.

Our entire continuum-normalized spectrum is therefore discontinuous and made up of six distinct regions with some gaps in between. These regions are selected by cropping the spectrum at a fixed distance away from each Balmer line – 300 Å for H α , 200 Å for H β , 120 Å for H γ , 70 Å for H δ , 50 Å for H ϵ , and 25 Å for H8. Higher order Balmer lines are too weak to be detected for large parts of our label space and are hence omitted from our main fitting tool. The regions are selected to be large enough to include a sufficient portion of the surrounding continuum, whilst excluding nearby Balmer lines and strong telluric features. The resulting continuum-normalized spectrum is therefore unbiased relative to the

¹<https://github.com/vedantchandra/wdtools>

²<https://wdtools.readthedocs.io>

Table 1. Summary of methods presented in this work, along with the traditional approach for comparison.

| Method | Training data | Input | Output | Error analysis |
|-----------------------------|--------------------------------|-------------------------|--|----------------------|
| Generative fitting pipeline | Grid of synthetic spectra | Theoretical models | Interpolated model spectra & label predictions | Bayesian MCMC |
| Parametric random forest | SDSS spectra with known labels | Line profile parameters | Label predictions | Bootstrap resampling |
| Traditional approach | N/A | Theoretical models | Label predictions | χ^2 statistics |

depth and width of the absorption lines, and is also independent of the shape of the overall underlying blackbody continuum (Bergeron et al. 1995). Normalizing each line independently also accommodates cases where clean spectral data are only available for some subset of the Balmer lines.

For helium-rich DB spectra, there are far more features on the spectrum and the above approach is not feasible. Instead, we normalize the entire continuum at once. We pre-define a mask that excludes all known hydrogen and helium absorption lines on the optical spectrum, and fit a smoothed spline to the remaining pixels. We divide the observed spectrum by the fitted spline curve to obtain the continuum-normalized spectrum. This method is quite effective at normalizing the entire continuum to unity, albeit it suffers slightly on either end of the spectrum. If the spectrum is known to be DA, we prefer the above approach of normalizing each Balmer line individually.

All subsequent analysis in this work is performed on spectra continuum-normalized using the methods detailed above. It is also possible to independently constrain T_{eff} from the blackbody shape of the thermal continuum. Whilst directly deriving this constraint is outside the scope of this work due to numerous confounding factors like spectro-photometric calibration, the methods in Section 3.1 are designed to incorporate such prior knowledge of T_{eff} if available.

3 METHODS

In this section, we summarize the analysis techniques implemented in WDTOOLS to measure the astrophysical properties of white dwarfs from their spectra. We also tabulate the inputs, outputs, and error analysis of each method in Table 1.

3.1 Generative fitting pipeline

In our first method, we fit state-of-the-art atmospheric models and the resulting synthetic spectra to SDSS spectroscopic data. State-of-the-art white dwarf synthetic spectra are based on theoretical models for pressure and thermal broadening to accurately reproduce the hydrogen Balmer lines as a function of the stellar labels. The codes to generate these spectra are not publicly available, but the methodology of one of the most popular codes is described in Koester (2010). These models incorporate the hydrogen Stark broadening with non-ideal effects from Tremblay & Bergeron (2009) and assume local thermodynamic equilibrium.

Around 1274 DA (pure-hydrogen atmosphere) synthetic spectra are publicly available³ that span 6000–40 000 K in T_{eff} and 6.5–9.5 dex in $\log g$, where g is in units of cm s^{-2} . Additionally, 625 DB (pure-helium) synthetic spectra were provided to us (Koester, private communication) spanning 6000–40 000 K in T_{eff} and 7–9 dex in $\log g$. In this work, we test and discuss both the DA and DB models for completeness. However, the DB models are under restricted access and we therefore do not include them in our accompanying software

package. In future, if these DB synthetic spectra become publicly available, we will add them to our software.

The usual fitting methodology is to fit the labelled model spectra to continuum-normalized absorption-line data and report the labels with the lowest χ^2 error (Kepler et al. 2019). However, this approach has some drawbacks. The published model grid is sampled sparsely, requiring interpolation between label values and multiple function evaluations. As mentioned above, the codes to generate atmospheric models directly are not publicly available. Even if they were, modelling stellar atmospheres is quite computationally expensive and requires generating model grids in advance before the fitting procedure (Koester 2008). Ting et al. (2018) propose neural networks as speedy, reliable interpolators that require few ab-initio models to generate high-quality synthetic spectra for main-sequence stars. We present an adaptation of this approach using white dwarf theoretical models to speedily generate synthetic spectra in a way that can be incorporated into the fitting routine itself.

We begin with the assumption that the stellar labels – T_{eff} and $\log g$ – uniquely determine the pixels on a continuum-normalized spectrum (Ness et al. 2015). We could include radial velocity as a label here, but instead choose to leave it as a free nuisance parameter during the final fitting procedure. We pre-process the models by using linear interpolation to bring all model spectra on to an identical wavelength range in 0.5 Å intervals and then continuum-normalize them with the procedures described in Section 2.2. The model spectra are later brought down to the resolution of the observed spectrum using Gaussian convolution during the fitting step and we therefore sample the ab-initio models finely for use with higher resolution spectroscopy. We are happy to provide pre-trained neural networks that interpolate an even finer grid of wavelengths upon request.

The generative neural network takes in stellar labels and creates continuum-normalized synthetic spectra from them. The network’s input layer consists of 2 neurons, corresponding to the 2 labels: T_{eff} and $\log g$. The output layer has thousands of neurons, each corresponding to one pixel on the continuum-normalized spectrum. The aim of training the network is to predict an accurate normalized flux at every pixel given only the 2 labels at the start. The bulk of the network’s thinking occurs in the so-called ‘hidden’ layers between the input and the output. The goal of the training step is to find the optimal combination of weights and biases w , b in the hidden layers to minimize the mean squared error between the flux predicted by the neural network and the true flux from the synthetic models. Through experimentation, we find that using 256 neurons in each of 3 hidden layers is optimal, allowing the neural network to fully learn the non-linear relationships between the labels and spectrum whilst minimizing interpolation artefacts.

The neural network is relatively inexpensive to train, taking a few hours of CPU time on a regular laptop to train itself on the continuum-normalized model spectra. Once trained, the network generates synthetic spectra from labels extremely quickly on the order of 500 spectra per second. With any machine learning technique that trains itself on data, it is good practice to test the model on data outside the original training set. This ensures that the model has correctly learned the underlying physics rather than ‘overfit’

³Courtesy the Spanish Virtual Observatory.




Figure 2. The neural networks accurately reproduce synthetic DA (top panel) and DB (bottom panel) spectra from the validation set of labels unseen during training. In red is the spectrum generated by the neural network and in black is the true synthetic spectrum for those labels. For the DA stars, the continuum-normalized Balmer lines are vertically displaced for clarity and are arranged from H α at the bottom to H 8 at the top. For DB stars, the entire continuum-normalized optical spectrum is interpolated.

the training data set. We therefore randomly reserve 1 per cent of the model spectra for validation and keep them out of the training procedure. Testing the interpolation error on these validation labels (unseen during training), we obtain a relative error of under 0.01 in normalized flux units across all pixels (Fig. 2).

We differ from the architecture of the PAYNE described in Ting et al. (2018) by having a broader neural network with two hidden layers of 256 neurons each, rather than 10 neurons each. Additionally, Ting et al. (2018) trained a different neural network for each spectral pixel, but we train a single network for the entire spectrum. This incorporates the information and mapping between adjacent pixels. Whilst our neural network has more free parameters, it is capable of learning more non-linear relationships, resulting in a reduction in per cent flux error of ~ 10 per cent compared to if we used Ting et al. (2018) architecture instead. Additionally, since we use a single neural network, our training time is also much shorter (Leung & Bovy 2018). We find that using a broader neural network architecture does not produce overfitting artefacts – the interpolation is smooth across both stellar labels (Fig. 3). Newer versions of the PAYNE developed by Ting et al. (2018) also incorporate similar improvements in their architecture.

Armed with the ability to generate our own high-fidelity synthetic spectra very quickly, we are not limited to the original grid of model spectra. We can now create synthetic spectra ‘on the fly’ during the




Figure 3. The interpolation accuracy of our DA generative neural network in label space for two random pixels on the continuum-normalized spectrum. The neural network reproduces the expected non-linear behaviour without any signs of interpolation artefacts or overfitting. The interpolation quality is similarly good across all pixels.

fitting process. With unlimited synthetic spectra in hand, we proceed to construct a fitting pipeline that utilizes the neural network. Our goal is to fit for T_{eff} and $\log g$ (hereafter T and l for brevity) on the continuum-normalized Balmer absorption lines. In the Bayesian framework, the posterior distribution of the stellar labels $f(T, l|X)$ of a single star is given by

$$f(T, l|X) = \frac{f_p(T, l) \times f(X|T, l)}{f(X)}, \quad (1)$$

where X is the data (continuum-normalized spectrum) and we define $f_p(T, l)$ as our prior knowledge of the stellar labels. By default, we use a uniform prior

$$f_p(T, l) = \mathbb{1}_{[6000, 80000]}(T) \times \mathbb{1}_{[6.5, 9.5]}(l), \quad (2)$$

where $\mathbb{1}$ denotes the indicator function on the stated range. This range is chosen to match the label range of the model spectra used to train the neural network, to prevent extrapolation errors. We include the ability to replace the uniform prior on T_{eff} with a Gaussian prior if another constraint on T_{eff} from photometry or blackbody continuum is available.

The purpose of the main fitting step is to determine the empirical distribution $\frac{f(X|T, l)}{f(X)}$. Since our fitting approach relies on the discrete sampling (in label space) of synthetic spectra, it is especially suited to fitting via Markov chain Monte Carlo (MCMC) algorithms. We utilize the EMCEE sampler for our MCMC procedure (Foreman-Mackey et al. 2013). EMCEE samples the posterior distribution of the stellar labels using an affine-invariant version of the Metropolis–Hastings algorithm (Goodman & Weare 2010) on a defined likelihood. We assume a chi-square empirical likelihood

$$\frac{f(X|T, l)}{f(X)} \sim \chi_n^2, \quad (3)$$

which is standard for curve fitting in astrophysics (Bevington 1969). The number of degrees of freedom n of χ_n^2 is nominally defined as the number of pixels used minus the number of parameters in the fit.

We include the radial velocity of the star as a ‘nuisance parameter’ in the MCMC fit, marginalizing over it. This is preferred over correcting the spectrum to some fixed rest frame with a single radial velocity measurement, as these measurements are susceptible to significant biases due to temperature and surface gravity (Halenka et al. 2015). Additionally for DB stars, it is non-trivial to measure an accurate radial velocity due to the lack of unblended symmetric absorption lines, so leaving it as a free parameter is preferred. However, if the sole goal is to derive high-precision radial velocity measurements from spectra, a simpler and model-independent approach like fitting Gaussian profiles to a small region around the $H\alpha$ absorption line may be preferred (Joyce et al. 2018). This prevents any systematic shortcomings in the atmospheric models from biasing the radial velocity measurement.

We begin our fitting procedure for each star with two initial fitting steps to evaluate the ‘hot’ and ‘cold’ solutions for the star. For the cold solution, we restrict $6000 < T_{\text{eff}} \leq 15000$ and for the hot solution we restrict $15000 < T_{\text{eff}} < 40000$, allowing all other parameters to vary freely. We use the differential evolution optimization algorithm (Storn & Price 1997) in LMFIT to minimize equation (3) and select the solution with the highest likelihood (lowest χ^2). For the main MCMC fitting procedure, we initialize independent ‘walkers’ in a tight Gaussian ball around the maximum-likelihood region of parameter space derived by the initial fit. The walkers explore the parameter space for a burn-in period, during which they spread out and settle into uncorrelated samples from the posterior distribution. Finally we reset the sampler and sample the posterior label distribution $f(T, l|X)$ for a number of steps, from which we compute the maximum-likelihood (minimum χ^2) stellar labels as well as label uncertainties using the standard deviation of each posterior distribution. This entire procedure entails upwards of 100 000 likelihood calls depending on the number of posterior samples required, yet takes less than a minute of CPU time per star on a regular laptop due to the speed of the generative neural network.

In this section, we have demonstrated how even in the absence of access to a finely spaced grid of models, we can use machine learning techniques to interpolate high-quality models for both DA and DB white dwarfs. We discuss the quality and limitations of this technique further in Section 4.

3.2 Parametric random forest

In our second technique, we seek to improve traditional methods of parametrically fitting absorption line profiles to derive atmospheric parameters. DA white dwarfs are characterized by strong hydrogen absorption lines. A standard approach to analyse spectral absorption lines is the use of line profiles to quantify line parameters like centroid, width, and scale. White dwarf absorption lines are formed over a range of depths and pressures, and can be modelled by the convolution between a symmetric Voigt profile and an asymmetric Stark profile (Tremblay & Bergeron 2009; Tremblay et al. 2019). We find that including asymmetric information adds complexity without significantly improving the accuracy of our model and we therefore approximate the white dwarf line profiles using the symmetric Voigt distribution.

We use PYTHON’s LMFIT (Newville et al. 2014) to fit Voigt profiles to the $H\alpha$, $H\beta$, $H\gamma$, and $H\delta$ absorption lines, performing a χ^2 minimization each time. We leave out higher order lines since they are not well-resolved for a lot of stars in our sample and we want to ensure




Figure 4. Visualizing the ‘phase space’ of the $H\gamma$ absorption line using two parameters from our fitted profiles, the full width at half-maximum and the line amplitude. There are intricate two-dimensional structures and trends in this phase space as a function of the stellar labels T_{eff} and $\log g$.

accurate profile fits. Our fitting procedure yields two parameters for each line: full width at half-maximum (FWHM) and amplitude. The line centroid is allowed to vary freely during the fit to account for radial velocity and is discarded as a nuisance parameter. These line profiles do not perfectly fit the entire absorption line – they are designed to fit the wings well, at the cost of the absorption line core. This is because most of the physical information about the stellar labels is encoded in the wings, not the core (Halenka et al. 2015). Additionally, the line core is usually resolution-limited, having a width of $\sigma_v \simeq 30 \text{ km s}^{-1}$ compared to the SDSS resolution of $\sigma_v \simeq 60 \text{ km s}^{-1}$ at $H\alpha$.

Absorption lines are known to predictably vary with changes in T_{eff} and $\log g$, due to effects like pressure broadening. Fig. 4 demonstrates the relation between our derived width and depth for the $H\gamma$ line coloured by known spectroscopic T_{eff} and $\log g$ labels from Tremblay et al. (2019). We select $H\gamma$ for this figure since it is the most sensitive to these labels – $H\alpha$, $H\beta$, and $H\delta$ exhibit similar relations. It is clear from Fig. 4 that it is intractable to distinguish T_{eff} and $\log g$ from any single line parameter; a particular value of FWHM, for example, would correspond to multiple sets of stellar labels. Additionally, the absorption line width is more sensitive to $\log g$ at higher temperatures than at temperatures below $\sim 12000 \text{ K}$ (Fig. 5). The goal of our multiparameter approach is to combine the information of the parameters across all lines to distinguish and fit for T_{eff} and $\log g$.

We use a multivariate random forest regression model to map line parameters to stellar labels. Random forests (Breiman 2001) have become a popular tool in astronomy to learn non-linear relations between features, in both supervised (Richards et al. 2011) and unsupervised (Reis, Baron & Shahaf 2019) learning environments. A random forest is composed of logical structures called decision trees. Each decision tree maps the input feature space (in our case the eight absorption line summary statistics) to continuous values in the target variable space (in our case the two stellar labels). A tree




Figure 5. The width of $H\gamma$ is more positively correlated with $\log g$ for stars with $T_{\text{eff}} > 12,000$ K and therefore carries more information about $\log g$ at high temperatures. Overlaid in dashed black is a linear regression fit, and ρ denotes the Pearson correlation coefficient between $\log g$ and the $H\gamma$ FWHM.

begins with a root node comprising of the entire input feature space – in our case the feature space is composed of two parameters from each of four Balmer lines, for a total of eight features. It then splits into several nodes, each of which represents a different subset of the feature space. It continues this recursive partitioning until all leaves originating from a particular node share (within a tolerance) the same value of the target variable.

Whilst decision trees on their own are powerful tools for regression, their predictive power is hampered by their tendency to overfit training data. This results in a poorly generalized model which fails to make meaningful predictions outside the data set it was trained on. For some applications, it is possible to fix this problem by combining an ensemble of decision trees through bootstrap aggregation (‘bagging’; Breiman 2001). Each decision tree is trained on a bootstrap sample of the training data rather than the entire data set, and the results of all trees in the ensemble are averaged to produce the final regression output. Bagging limits overfitting and reduces the variance of the regression.

One challenge with random forests is that they do not natively quantify uncertainties or errors. We implement our own method to produce prediction uncertainties based on Coulston et al. (2016). We construct 100 bootstrap resamples of the stars in our training sample, sampling with replacement. We train an ensemble of 100 random forest models, one on each of the above bootstrapped samples. We make predictions using this entire ensemble, and use the variance of predictions across all component decision trees to define a test statistic as defined in Coulston et al. (2016), and subsequently obtain 1σ uncertainties for our fitted labels. Since each random forest in the ensemble is trained on an overlapping but different set of stars, the aggregate predictions of the ensemble capture the predictive uncertainty of the model (Coulston et al. 2016). Additionally, the variance across all the individual decision trees can capture the epistemic and noise uncertainty in the data, since each tree can be thought of as a ‘weak learner’. When testing our predictions on a validation set of stars, we find that ~ 60 per cent of our predicted labels lie within 1σ of the ground truth values obtained by prior studies and ~ 85 per cent lie within 2σ . This indicates that our uncertainty estimation procedure slightly underestimates the errors, but is otherwise in good agreement with Gaussian statistics.

Another possible challenge is the non-uniform distribution of our training sample in label space. In particular, the $\log g$ distribution is sharply peaked around $\log g \sim 8$ dex. Since most training stars have labels around this value, the random forest algorithm will tend towards predicting values closer to the mean – this worsens the prediction accuracy near the tails of the distribution. We combat this problem during the ensemble resampling step. We fit a two-dimensional Gaussian to the density distribution of the objects in the training label space (over T_{eff} and $\log g$) to obtain the kernel density estimate (KDE) for the sample. We then assign each star a weight equivalent to the reciprocal of the KDE’s probability density value for that star’s labels. We sample our stars with replacement using these weights during the ensemble resampling. Therefore, stars near the tails of the label distributions are favoured and the distribution of labels used to train each random forest is more uniform.

We use the SCIKIT-LEARN library in PYTHON to construct our random forest models with 25 decision trees each (Pedregosa et al. 2011). There are a total of eight features derived from each stellar spectrum – amplitude and FWHM for $H\alpha$, $H\beta$, $H\gamma$, and $H\delta$, respectively. These features are used to predict the two stellar labels (T_{eff} and $\log g$) simultaneously. Our ensemble model is fast to train, on the order of a few seconds for 5000 stars; predicting labels for thousands of stars at once likewise takes a couple of seconds on a regular laptop computer.

In this section, we have described a way to combine parametric information from all the absorption lines on a white dwarf spectrum to make inferences about stellar labels. This should result in a less noisy model compared to direct pixel methods, since the parametric line fit incorporates much of the noise in the spectrum flux itself. Given a large sample of spectra of which some subset has accurate labels, this method is powerful for inferring labels for the entire sample. This is a form of ‘transfer learning’ and we have previously used it to derive spectroscopic $\log g$ for a large sample of SDSS white dwarfs (Chandra et al. 2020). We evaluate this method on SDSS spectra in Section 4.

3.3 Convolutional neural networks (CNNs)

In recent years, interest in CNNs has grown, primarily in the fields of image recognition and visual artificial intelligence. Leung & Bovy (2018) describe CNNs with intrinsic uncertainty estimation to infer abundances from APOGEE spectra. We adapt their architecture to develop a Bayesian CNN that is trained to predict T_{eff} and $\log g$ from SDSS white dwarf spectra directly. The ‘true’ labels used to train this model are taken from the spectroscopic fits of Tremblay et al. (2019). We use our CNN to simultaneously incorporate all continuum-normalized spectral pixels into a regression model to predict stellar labels.

In practice, the CNN achieves label prediction accuracy similar to the random forest method described in Section 3.2. However, this kind of ‘black-box’ prediction algorithm suffers badly from a lack of interpretability. For example, since the training set of stars are all pure-hydrogen DA stars, the neural network makes unrealistic predictions when faced with a non-DA star. We find that whilst the data-driven neural network prediction models work well for applications like Leung & Bovy (2018) where several element abundances must be constrained from a complex and blended spectrum, the simpler and more interpretable methods of Section 3.1 and 3.2 work better for our data. We therefore include this discussion of CNNs here for completeness, but do not consider them further in our results or software package.

4 RESULTS

4.1 DA white dwarfs from the Sloan Digital Sky Survey

Tremblay et al. (2019) derived spectroscopic T_{eff} and $\log g$ for a sample of 5327 DA white dwarfs using their latest atmospheric model codes. These models have been shown to be consistent with results obtained by other groups and photometry, with typical root mean square (rms) discrepancy between stellar labels derived by different methods being on the order of 500 K in T_{eff} and 0.1 dex in $\log g$. We use the stellar labels from Tremblay et al. (2019) as a ‘ground truth’ sample to test the methods described in this work. For consistent model comparison with the random forest regression, we divide the sample into a training set of 4011 stars (75 per cent) and a testing set of 1316 stars (25 per cent).

We run our generative fitting pipeline on all 1316 spectra in the testing set, obtaining posterior label distributions in a minute per star on average. These fits only consider the information on the spectrum. In practice, it may be possible and desirable to combine spectroscopic data with photometry and obtain a prior on the temperature, which may improve the spectroscopic fit. Our MCMC fitting approach almost always finds the global χ^2 minimum, since a choice between hot and cold solutions on the basis of χ^2 is made during our pre-burn-in phase where we select the highest probability region of the parameter space. Regardless, we recommend always using photometry, if available, as an additional constraint on T_{eff} . We present our fitted labels compared to Tremblay et al. (2019) in Fig. 6 (top panels).

We train the bootstrap random forest ensemble on the training set of 4011 stars and make predictions with uncertainties on the test set of 1316 spectra. Our set of training features consists of line summaries from H α to H δ , since these lines are very well resolved on all stars in our training set of stars. We compute uncertainties as described in Section 3.2 and present our predictions compared to Tremblay et al. (2019) in Fig. 6 (bottom).

We can use the uncertainties reported by our methods to remove grossly under-determined labels – we reject labels with uncertainty in $T_{\text{eff}} \geq 2500$ K and uncertainty in $\log g \geq 0.5$ dex. These stars also have systematically higher χ_r^2 , although we find that χ_r^2 alone is usually not sufficient to reject spurious fits (see Section 4.4). Upon inspecting these stars further, we find that they usually have some deviation from the model assumptions, like weak magnetism or a very low signal-to-noise ratio. We compare our derived labels to those from Tremblay et al. (2019) in Fig. 6 with the minimal selection cuts on reported error described above. Both techniques recover T_{eff} within ~ 1000 K for all stars, and $\log g$ within ~ 0.12 dex.

4.2 DB white dwarfs from the Sloan Digital Sky Survey

Koester & Kepler (2015) identified 1000 SDSS spectra with a DB classification and fit them with their atmospheric models. We fit these spectra with our generative fitting pipeline. The $\log g$ constraint is not particularly specific since the atmospheric models are not very sensitive to variations in $\log g$ across all temperatures. Furthermore, Koester & Kepler (2015) noted that $\log g$ is overestimated for stars with $T_{\text{eff}} \lesssim 15000$ K, which we find with our labels as well. This effect represents a well-known systematic shortcoming in our current understanding of white dwarf atmospheres. We therefore ignore $\log g$ labels for stars with $T_{\text{eff}} < 15000$ K.

We recover T_{eff} within ~ 1200 K across the entire range of temperatures (6000 K–40000 K) and $\log g$ within 0.25 dex for stars with $T_{\text{eff}} > 15000$ K. The rms discrepancy between our temperatures

and the prior work is 632 K if we only consider stars with $T_{\text{eff}} < 22000$ K (90 per cent of the sample). As with Koester & Kepler (2015), we find a gap in the fitted temperatures between 22000 and 24000 K and a subsequent degradation of the quality of fits for higher temperatures. Overall, our label predictions are consistent with the fundamental parameter accuracy from Tremblay et al. (2019), who directly compared the Koester & Kepler (2015) spectroscopic labels to photometric labels from *Gaia*.

As mentioned above, the DB models provided to us are under restricted access and hence omitted from our accompanying package until they become publicly available. However, we have demonstrated that given a grid of DB models, our methods produce results consistent with a sample of SDSS DB stars from Koester & Kepler (2015).

4.3 DA white dwarfs from VLT-UVES

As a part of the ESO Supernova Ia Progenitor Survey (SPY), Koester et al. (2009) obtained high-resolution spectra of 615 DA white dwarfs with the Ultraviolet and Visual Echelle Spectrograph (UVES; Dekker et al. 2000) of the Very Large Telescope (VLT). The spectral resolution at H α is around $R = 18500$ with per pixel signal-to-noise ratio ≥ 15 for most spectra. Koester et al. (2009) derived atmospheric parameters for these white dwarfs using their models and χ^2 minimization. Since then their models have changed considerably, especially with the inclusion of effects from Tremblay & Bergeron (2009). Our generative neural network is trained on synthetic spectra from the latest models that include these additional effects. We therefore revisit the spectroscopic catalogue of Koester et al. (2009) using our generative fitting pipeline.

We continuum-normalize the H γ , H δ , H ϵ , and H δ spectra from cross-disperser #2 of UVES. We do not include the H α and H β lines from cross-disperser #3 due to their lower signal-to-noise ratio and the presence of noisy artefacts. We use the generative fitting pipeline described in Section 3.1 to infer T_{eff} and $\log g$ from the continuum-normalized spectra.

Comparing fitted labels from our generative fitting pipeline to prior labels from Koester et al. (2009), we find a rms discrepancy of 1500 K in T_{eff} and 0.2 dex in $\log g$. For stars above 20000 K, there is an interesting systematic bias of ~ 1000 K between our labels and those from Koester et al. (2009), albeit with low variance. Overall, we predict systematically lower T_{eff} by 1000 K and systematically higher $\log g$ by 0.2 dex. We attribute this bias to differences in the models used by our software and those used by Koester et al. (2009) over a decade ago, since the newer models include significant new physics from Tremblay & Bergeron (2009) and other theoretical studies.

We also use the random forest method of Section 3.2, pre-trained on line summaries from the H δ and H γ lines of 5327 DA white dwarfs from the SDSS as described in Section 4.1. Our software tool includes the capability to easily switch the model to any subset of the first four Balmer lines and we test its applications beyond SDSS spectra here. Fitting the SPY sample with our random forest model, we recover the labels from Koester et al. (2009) to within 1000 K in T_{eff} and 0.25 dex in $\log g$ (RMSE). This method systematically overestimates T_{eff} by 200 K and $\log g$ by 0.1 dex, an improvement over the generative fitting pipeline. Our interpretation of this improvement is that since the random forest relies on gross line summaries (FWHM and amplitude) of the absorption lines, it is less sensitive to the intricate differences in physics between the older atmospheric models and the newer ones.




Figure 6. Summary of the prediction accuracy of the generative fitting pipeline (top panels) and random forest model (bottom panels) for both T_{eff} (left-hand panels) and $\log g$ (right-hand panels), comparing our labels for a validation set of 1316 DA SDSS spectra to ‘true’ labels derived from an earlier spectroscopic study (Tremblay et al. 2019). One sigma error bars as reported by the respective methods are overlaid in cyan.

4.4 Finding outlier spectra

Apart from inferring stellar labels, the generative fitting pipeline (Section 3.1) can be used to find interesting spectra that violate the assumptions of the atmospheric model. For DA white dwarfs, our assumption is a pure hydrogen atmosphere and no magnetic field, with no metal pollution. Kepler et al. (2019) provide a catalogue of over 35 000 SDSS white dwarfs with visually confirmed spectroscopic classifications. We fit pure-hydrogen models to a random subset of the spectra in this catalogue and compute the reduced χ_r^2 goodness-of-fit statistic (χ^2 per degrees of freedom; Bevington 1969). We find that for all spectra classified as pure-hydrogen (DA) by Kepler et al. (2019), the $\chi_r^2 \sim 1$ as expected, usually slightly lower. However, for all other spectral types, the χ_r^2 values are systematically higher, including for DAH (magnetic) white dwarfs.

However, the χ_r^2 statistic is not sensitive enough for complete outlier removal by itself. That is, several spectra that are not pure-hydrogen also have $\chi_r^2 \sim 1$ due to overestimated noise or other effects. The 99th percentile χ_r^2 among the pure-DA stars is 1.25. By applying a selection cut of $\chi_r^2 < 1.25$ to the fitted spectra, we retain completeness of 99 per cent the true-positive DA stars and get rid of ~ 50 per cent of the contaminant spectra. The remaining contaminant spectra cannot be differentiated on the basis of χ_r^2 alone.

There are several ways to identify non-DA spectra that are more specific than the χ_r^2 method, for example, searching for particular absorption lines. Calcium lines are strong indicators of metal-

polluted white dwarfs (Farihi et al. 2010), with the Ca I resonance line being especially prominent in cool white dwarf spectra (Blouin et al. 2019). The traditional way to identify magnetized white dwarfs is visual inspection of the Balmer lines to search for Zeeman splitting, although this is only resolved for fields $\gtrsim 2$ MG on SDSS-resolution spectra (Kepler et al. 2013). Additionally, contaminant A/F type main-sequence stars can be removed from samples of white dwarf spectra on the basis of their narrower Balmer lines (Kepler et al. 2019). However, methods that rely on visual inspection are outright infeasible for larger data sets such as those from upcoming spectroscopic surveys. Automating the detection of potentially interesting and outlier white dwarf spectra is the subject of ongoing research in our group.

Fitting the spectrum with a pure-DA model and inspecting the quality of fit and returned stellar labels from different observables can also be useful to identify outlier spectra (Genest-Beaulieu & Bergeron 2019). Significant disagreement between spectroscopic and photometric labels can be used to identify double-degenerates like DA+DA and DA+DC binaries. We have also verified, using simulated spectra of binaries, that increasing the Doppler shift between the two stars in a pair results in a gradual increase in the reported χ_r^2 as expected. Our tool can be used in conjunction with the ‘WD-models’ package (Cheng 2020) in PYTHON to compare spectroscopic and photometric labels. Future iterations of our tool will include improved integration with ‘WD-models’, facilitating




Figure 7. Example output of our generative fitting pipeline for a pure-DA spectrum (top panel) as well as two outlier white dwarfs (middle and bottom panels) from the SDSS catalogue (Kepler et al. 2019). Balmer lines are arranged from H α to H 8 from bottom to top. The DAH is a magnetic white dwarf with symmetric Zeeman splitting of the Balmer absorption lines. DAO spectra have He II absorption lines and characteristically flattened Balmer lines. The latter spectrum shows very weak Helium absorption lines upon careful inspection, but was initially identified by the fact that our fitted spectroscopic stellar labels disagree significantly with its photometric labels (Dufour et al. 2017).

the use of photometric priors for spectroscopic fits. Single lined DA+DA binaries can look exactly like single DA spectra, and therefore using a photometric label constraint is powerful for finding inconsistent spectra. We show some examples of outlier SDSS spectra we found using our tool in Fig. 7. Genest-Beaulieu & Bergeron (2019) provide a comprehensive summary of how to use spectroscopic fitting techniques to find interesting white dwarf systems.

As a real-world example of these applications, this year we obtained spectra for five high-mass DA candidates with the Double Imaging Spectrometer (DIS) at the Apache Point Observatory (APO). We based our selection on SDSS and *Gaia* data, looking for unusual high-mass white dwarfs with high velocities and/or photometric variability. We ran our generative fitting pipeline on these spectra and derived effective surface temperatures and surface gravities. Our labels for two targets were inconsistent with photometry. One of these (SDSS J105449.74+552307.9) was revealed by follow-up Gemini spectroscopy (program GN-2018B-FT-115, PI: Hwang) to be a weakly magnetic white dwarf (DAH). Another candidate exhibits interesting radial velocity variations that warrant further follow-up observations. This is an example of ways we have used our tool to find interesting white dwarf candidates.

5 DISCUSSION

In this paper, we have developed and summarized two methods to derive stellar labels T_{eff} and $\log g$ from spectroscopic observations of hydrogen-rich white dwarfs. Each has its advantages and disadvantages for different applications. The generative fitting pipeline (Section 3.1) is a similar methodology to traditional fitting with synthetic models, albeit made much faster by the use of neural networks to generate synthetic spectra from a grid of pre-computed model spectra. This approach is best suited for single observations where an accurate characterization of the label posterior distributions is desired. This is certainly the most interpretable technique since it provides a visual goodness-of-fit to model spectra, χ^2 statistics, and statistically precise labels and uncertainties consistent with past studies (Tremblay et al. 2019).

Whilst the generative fitting pipeline is comparatively fast for single stars, it can require significant computational resources to fit a large data set of stars. The parametric random forest (Section 3.2) is a much faster method suitable for transfer learning, to derive labels for a large data set when accurate labels are known for some subset of the data set. However, some information may be lost when reducing the dimensionality from the spectrum to fitted line parameters.

One crucial limitation of the methods described here is that they rely on the accuracy of the underlying theoretical models. The generative fitting pipeline is limited by the systematic uncertainties in theoretical synthetic spectra and the random forests are additionally limited by any systematics in the training labels derived by other groups. Current theoretical limitations likely imply that spectroscopic $\log g$ cannot be constrained better than within 0.1 dex. An additional unresolved uncertainty is the ‘log g upturn’ that causes low-temperature stars (T_{eff}) to have systematically overestimated $\log g$ measurements.

Our software tool is not intended to replace careful and intensive spectroscopic analysis with full-featured atmospheric models. Many scientific cases require more advanced models, for example, those studying metal polluted or very cool white dwarfs (Blouin 2020). However, the proprietary and restricted nature of nearly all current white dwarf models necessitated a tool like ours to derive white dwarf atmospheric parameters. We will expand our tool’s functionality as more white dwarf models become publicly available in the future. Our generative fitting pipeline can be trained on new model spectra that span a wider label space and can be generalized to additional labels like magnetic field and metal abundances. Additionally, our random forest tool can be re-trained when larger spectroscopic data sets are fit by groups with proprietary model atmospheres.

In this work, we have summarized the methodology and accuracy of WDTOOLS and summarized several possible applications. We

utilized WDTOOLS in our study of white dwarf gravitational redshifts (Chandra et al. 2020), using the generative fitting pipeline of Section 3.1 to derive $\log g$ for over 3000 DA white dwarfs. When combined with gravitational redshifts, this enabled us to empirically measure the mass–radius relationship of white dwarfs across a wide range of masses, the first measurement of its kind for such a large sample of stars. We hope the open-source availability of our tool, as well as future public releases of white dwarf atmospheric models, will enable more statistical studies of the ever-growing sample of known white dwarfs in our Galaxy.

ACKNOWLEDGEMENTS

We thank the referee, Pier-Emmanuel Tremblay, for constructive comments and suggestions that significantly improved our manuscript. VC thanks Yuan-Sen Ting, JJ Hermes, Kirsten Hall, and Sihao Cheng for productive conversations and feedback on the analysis. We thank Detlev Koester for providing us with a grid of DB synthetic spectra to validate our software. VC is supported by the JHU PURA and DURA. VC and HCH are supported by Space@Hopkins. VC, HCH, and NLZ were supported in part by NASA-ADAP 80NSSC19K0581.

Based in part on observations from the Sloan Digital Sky Survey. Funding for the Sloan Digital Sky Survey IV (Blanton et al. 2017) has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (Gaia Collaboration et al. 2016, 2018), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC). Part of this research project was conducted using computational resources at the Blue Crab cluster of the Maryland Advanced Research Computing Center (MARCC).

Based in part on observations collected at the European Southern Observatory, as well as on observations obtained with the Apache Point Observatory 3.5-meter telescope, which is owned and operated by the Astrophysical Research Consortium. Based on observations obtained at the international Gemini Observatory, a program of NSF's NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation on behalf of the Gemini Observatory partnership: the National Science Foundation (United States), National Research Council (Canada), Agencia Nacional de Investigación y Desarrollo (Chile), Ministerio de Ciencia, Tecnología e Innovación (Argentina), Ministério da Ciência, Tecnologia, Inovações e Comunicações (Brazil), and Korea Astronomy and Space Science Institute (Republic of Korea). This research has made use of the Spanish Virtual Observatory supported from the Spanish MINECO/FEDER through grant AyA2017-84089.

DATA AVAILABILITY

The *Gaia*, SDSS, and ESO UVES data used in this work to validate our software are freely available from their respective data archives (<https://gea.esac.esa.int/archive/>, <http://skyserver.sdss.org/dr16/en/home.aspx> and http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form respectively). The Gemini spectra we observed during program GN-2018B-FT-115 are now publicly available from their archive (<https://archive.gemini.edu/searchform>). The software tool we present here is open-source and hosted on GitHub (footnote 1) and Zenodo (<http://doi.org/10.5281/zenodo.3828686>). The DA synthetic spectra we use to train our model are hosted online (footnote 3) and the DB synthetic spectra provided to us by Koester (private communication)

are currently under restricted access. The PYTHON notebooks used to generate the figures in this paper are available from the corresponding author upon request.

REFERENCES

- Bergeron P., Liebert J., Fulbright M. S., 1995, *ApJ*, 444, 810
 Bevington P. R., 1969, *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill, New York
 Blanton M. R. et al., 2017, *AJ*, 154, 28
 Blouin S., 2020, *MNRAS*, 496, 1881
 Blouin S., Allard N. F., Leininger T., Gad ea F. X., Dufour P., 2019, *ApJ*, 875, 137
 Bohlin R. C., Hubeny I., Rauch T., 2020, *AJ*, 160, 21
 Breiman L., 2001, *Mach. learn.*, 45, 5
 Brown W. R., Kilic M., Allende Prieto C., Gianninas A., Kenyon S. J., 2013, *ApJ*, 769, 11
 Camisassa M. E. et al., 2019, *A&A*, 625, A87
 Chandra V., 2020, *wdtools: Computational Toolbox for the Spectroscopic Analysis of White Dwarfs*
 Chandra V., Hwang H.-C., Zakamska N. L., Cheng S., 2020, preprint ([arXiv:2007.14517](https://arxiv.org/abs/2007.14517))
 Cheng S., 2020, WD-Models, https://github.com/SihaoCheng/WD_models
 Coulston J. W., Blinn C. E., Thomas V. A., Wynne R. H., 2016, *Photogramm. Eng. Remote Sens.*, 82, 189
 Dekker H., D'Odorico S., Kaufer A., Delabre B., Kotzłowski H., 2000, in Iye M., Moorwood A. F., eds, *Proc. SPIE Conf. Ser. Vol. 4008, Optical and IR Telescope Instrumentation and Detectors*. SPIE, Bellingham, p. 534
 DESI Collaboration et al., 2016, preprint ([arXiv:1611.00036](https://arxiv.org/abs/1611.00036))
 Dufour P., Blouin S., Coutu S., Fortin-Archambault M., Thibeault C., Bergeron P., Fontaine G., 2017, in Tremblay P.-E., G ansicke B., Marsh T., eds, *ASP Conf. Ser. Vol. 509, The Montreal White Dwarf Database: A Tool for the Community*. Astron. Soc. Pac., San Francisco, p. 6
 Farihi J., Barstow M. A., Redfield S., Dufour P., Hambly N. C., 2010, *MNRAS*, 404, 2123
 Fontaine G., Brassard P., Bergeron P., 2001, *PASP*, 113, 409
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
 Fusillo N. P. G. et al., 2019, *MNRAS*, 482, 4570
 Gaia Collaboration et al., 2016, *A&A*, 595, A1
 Gaia Collaboration et al., 2018, *A&A*, 616, A1
 Genest-Beaulieu C., Bergeron P., 2019, *ApJ*, 871, 169
 Goodman J., Weare J., 2010, *Commun. Appl. Math. Comput. Sci.*, 5, 65
 Halenka J., Olchawa W., Madej J., Grabowski B., 2015, *ApJ*, 808, 131
 Hansen B. M. S. et al., 2002, *ApJ*, 574, L155
 Hubeny I., 1988, *Comput. Phys. Commun.*, 52, 103
 Joyce S. R. G., Barstow M. A., Holberg J. B., Bond H. E., Casewell S. L., Burleigh M. R., 2018, *MNRAS*, 481, 2361
 Kepler S. O. et al., 2013, *MNRAS*, 429, 2934
 Kepler S. O. et al., 2019, *MNRAS*, 486, 2169
 Kilic M., Bergeron P., Kosakowski A., Brown W. R., Agueros M. A., Blouin S., 2020, *ApJ*, 898
 Koester D., 2008, preprint ([arXiv:0812.0482](https://arxiv.org/abs/0812.0482))
 Koester D., 2010, *Mem. Soc. Astron. Ital.*, 81, 921
 Koester D., Kepler S. O., 2015, *A&A*, 583, A86
 Koester D., Voss B., Napiwotzki R., Christlieb N., Homeier D., Lisker T., Reimers D., Heber U., 2009, *A&A*, 505, 441
 Kollmeier J. A. et al., 2017, preprint ([arXiv:1711.03234](https://arxiv.org/abs/1711.03234))
 Leung H. W., Bovy J., 2018, *MNRAS*, 3277, 3255
 Levenhagen R. S., Diaz M. P., Coelho P. R. T., Hubeny I., 2017, *ApJS*, 231, 1
 Maoz D., Mannucci F., Nelemans G., 2014, *ARA&A*, 52, 107
 McCleery J. et al., 2020, *MNRAS*, 20
 Narayan G. et al., 2019, *ApJS*, 241, 20
 Nataf D. M. et al., 2019, *AJ*, 158, 14
 Ness M., Hogg D. W., Rix H. W., Ho A. Y., Zasowski G., 2015, *ApJ*, 808, 16
 Newville M., Stensitzki T., Allen D. B., Ingargiola A., 2014, *Astrophysics Source Code Library*, record ascl:1606.014

- Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
Reis I., Baron D., Shahaf S., 2019, *AJ*, 157, 16
Richards J. W. et al., 2011, *ApJ*, 733, 10
Storn R., Price K., 1997, *J. Glob. Optim.*, 11, 341
Ting Y.-S., Conroy C., Rix H.-W., Cargile P., 2018, *ApJ*, 879, 22
Tremblay P. E., Bergeron P., 2009, *ApJ*, 696, 1755
Tremblay P. E., Bergeron P., Kalirai J. S., Gianninas A., 2010, *ApJ*, 712, 1345
Tremblay P.-E., Ludwig H.-G., Steffen M., Freytag B., 2013, *A&A*, 559, A104
Tremblay P. E., Cukanovaitė E., Gentile Fusillo N. P., Cunningham T., Hollands M. A., 2019, *MNRAS*, 482, 5222
Tremblay P. E. et al., 2020, *MNRAS*, 497, 130
Xiang M. et al., 2019, *ApJS*, 245, 34

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