AARHUS UNIVERSITY

Coversheet

This is the publisher's PDF (Version of Record) of the article.

This is the final published version of the article.

How to cite this publication:

Marc Hon, Dennis Stello, Jie Yu (2017). "Deep learning classification in asteroseismology." *Monthly Notices of the Royal Astronomical Society*, 469 (4), 4578–4583.

<https://doi.org/10.1093/mnras/stx1174>

Publication metadata

Title: *Deep learning classification in asteroseismology*
Author(s): *Marc Hon, Dennis Stello, Jie Yu*
Journal: *Monthly Notices of the Royal Astronomical Society*
DOI/Link: <https://doi.org/10.1093/mnras/stx1174>
Document version: *Publisher's PDF (Version of Record)*

This article has been accepted for publication in Monthly Notices of the Royal Astronomical Society ©: 2016 The Authors Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

Redistribution subject to OUP license or copyright; see Terms of Use at https://academic.oup.com/journals/pages/access_purchase/rights_and_permissions/self_archiving_policy_p Use of this article is subject to OUP terms of use and conditions.

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

If the document is published under a Creative Commons license, this applies instead of the general rights.

Deep learning classification in asteroseismology

Marc Hon,^{1★} Dennis Stello^{1,2,3} and Jie Yu²

¹*School of Physics, The University of New South Wales, Sydney, NSW 2052, Australia*

²*Sydney Institute for Astronomy (SfA), School of Physics, University of Sydney, NSW 2006, Australia*

³*Department of Physics and Astronomy, Stellar Astrophysics Centre, Aarhus University, Ny Munkegade 120, DK-8000 Aarhus C, Denmark*

Accepted 2017 May 10. Received 2017 April 28; in original form 2017 March 22

ABSTRACT

In the power spectra of oscillating red giants, there are visually distinct features defining stars ascending the red giant branch from those that have commenced helium core burning. We train a 1D convolutional neural network by supervised learning to automatically learn these visual features from images of folded oscillation spectra. By training and testing on *Kepler* red giants, we achieve an accuracy of up to 99 per cent in separating helium-burning red giants from those ascending the red giant branch. The convolutional neural network additionally shows capability in accurately predicting the evolutionary states of 5379 previously unclassified *Kepler* red giants, by which we now have greatly increased the number of classified stars.

Key words: asteroseismology – methods: data analysis – techniques: image processing – stars: oscillations – stars: statistics.

1 INTRODUCTION

A key concept in determining stellar ages of red giants is distinguishing the evolutionary state, i.e. classifying between hydrogen-shell burning stars ascending the red giant branch (RGB) and those that have commenced core helium burning (HeB). Space missions such as *Kepler* (Borucki et al. 2010) have provided an enormous quantity of red giant oscillation spectra. The upcoming *TESS* mission (Ricker et al. 2014) will further increase the amount of data, making manual or semi-automatic classification of the population class of each star infeasible. Automated methods do exist, however considerable effort is required into defining and acquiring features such as the observed period spacing ΔP (Bedding et al. 2011; Stello et al. 2013), the asymptotic period spacing $\Delta \Pi_1$ (Bedding et al. 2011; Mosser et al. 2014; Vrad, Mosser & Samadi 2016) or the structure of mixed modes (Elsworth et al. 2016) in order to separate the populations. Furthermore, these methods require relatively high signal-to-noise data.

In this paper, we present a deep learning method that allows spectral features to be learnt with the help of the machine using convolutional neural networks. These are machine learning methods that mimic biological neuron structures, aimed towards feature detection in data (Fukushima 1980). They have achieved significant success over the past few years in computer vision methods such as image recognition (Krizhevsky, Sutskever & Hinton 2012) and even facial recognition (Garcia & Delakis 2004).

We introduce the concept of representing the oscillation frequency spectra of stars as ‘images’. While spectra are represented as a 1D array of values and are not images in the conventional 2D

sense, we interpret the spectra in an image-like fashion in such a way that we wish to identify the spatial structure of the spectra in frequency as visual features, as seen by ‘an expert eye’. Hence, these images are simple representations of the power excess, which contain visual features for RGB–HeB classification that the neural network can ‘see’ and hence learn from. By learning from an existing set of classified red giants based on $\Delta \Pi_1$ measurements, we use 1D convolutional neural networks as a form of supervised machine learning aimed to automatically learn features separating RGB from HeB stars in order to make fast yet accurate classification predictions on vast amounts of unclassified stars.

2 METHODS

Here we describe the preparation of the image representation known as a *folded spectrum*, along with an overview of convolutional neural networks, and the construction of a deep learning classifier to classify the image representation.

2.1 Data

We obtain the evolutionary state classifications of 5673 *Kepler* stars based on automated asymptotic period spacing measurements by Vrad et al. (2016, hereafter V16), and add 335 stars from the classification by Mosser et al. (2014, hereafter M14) that are not already in V16’s sample, to a total of 6008 stars. We then assign RGB stars with the binary class 0 and HeB stars with class 1. About 30 per cent of stars in the data set are RGB stars. We randomly choose 1008 stars as test data, with the remaining 5000 stars for training. Additionally, we have an *unclassified set* comprising 8794 *Kepler* red giants that are known to oscillate but have not been given classifications by V16 or M14. We want to predict the population

* E-mail: mtyh555@uowmail.edu.au

Figure 1. Comparisons between an RGB star KIC 11293804 (top), with an HeB star KIC 5810333 (bottom), both having large frequency spacing $\Delta\nu \simeq 3.92 \mu\text{Hz}$. Panels (a) and (b) are the original power spectra, while (c) and (d) are *folded spectrum* image representations. The oscillation modes are labelled by their degree l , while the frequency of maximum oscillation power, ν_{max} , is indicated by the dashed vertical line.

labels of all stars in our unclassified set using our trained neural network.

2.2 Image representation

As our image representation, we define the *folded spectrum* as the $4\Delta\nu$ -wide power spectrum segment centred at ν_{max} , folded by a length of $\Delta\nu$ (see Figs 1(a) and (c)). The spectra and values for $\Delta\nu$ and ν_{max} were derived from end-of-mission *Kepler* data using the SYD pipeline (Huber et al. 2009, Yu et al., in preparation). Because the neural network requires a fixed input array length, we bin each folded spectrum into 1000 bins.

A comparison of spectral image representations between RGB and HeB stars is shown in Figs 1(c) and (d). RGB stars clearly exhibit acoustic modes that are highly localized (Figs 1(a) and (c)), while HeB stars show broader mode distributions particularly for non-radial modes because of the stronger coupling between core and envelope (Figs 1(b) and (d)) (Dupret et al. 2009; Grosjean et al. 2014). With acoustic resonances less localized, HeB spectral representations notably have greater visual complexity as compared to RGB spectra. Besides the structure of modes, the location of the $l = 0$ mode, represented by ϵ , can be a strong indicator in distinguishing population classes (Kallinger et al. 2012). However, ϵ is not the sole feature that is used to recognize population classes from an image. The lack of a clear boundary separating the two evolutionary states shown by the observed spread in ϵ (Kallinger et al. 2012) and from theoretical studies (Christensen-Dalsgaard et al. 2014) makes ϵ unsuitable as the sole selection criterion. However, information about ϵ complements features extracted from mixed modes in the image.

As image pre-processing, we normalize each spectrum by its max power value. Then, to avoid edge effects, we append the image with a copy of itself and apply a super-Gaussian (higher-order Gaussian) weight function as shown in Fig. 2.

2.3 Convolutional neural networks

An artificial neural network is a mathematical representation of a biological neuron network (Fig. 3a). Mathematical neurons contain real numbers and connect to other neurons in subsequent levels

Figure 2. (a) Appended and normalized folded spectrum of KIC 10790301 with mode identification. s is the shape parameter or order ($=2$ for standard Gaussian), while σ is the standard deviation of the super-Gaussian weight. The solid vertical line separates the original image from its appended copy. (b) Resulting image from application of the super-Gaussian weight to the spectrum in (a).

(network layers) by mathematical operations in order to form a network capable of computing solutions to complex problems. The total input from one layer to a neuron in the next network layer is given by $\mathbf{w} \cdot \mathbf{x}$, where $\mathbf{x} = (x_0, x_1, x_2, x_3, \dots, x_n)$ is an input vector with n number of features from the input layer (represented by the number of neurons in the input layer in Fig. 3a), with $x_0 = 1$. In our study, $\mathbf{x}$ is the power in each frequency bin of the image at the very first input layer. In subsequent layers, $\mathbf{x}$ will become manipulated representations of the original input image. The weight vector, $\mathbf{w} = (w_0, w_1, w_2, \dots, w_n)$, links each input to a neuron in the subsequent layer, with w_0 known as the input bias, b , which is analogous to the intercept in a linear regression.

The total input is linear, however it is passed through a non-linear *activation function* (Rosenblatt 1962), f , such that the network becomes capable of approximating complex non-linear representations. In this study, we use the *rectified linear unit* activation function $f(x) = \max(0, x)$ for every neural network layer except the output layer. This function is suited for feature learning in neural networks (Nair & Hinton 2010). A common design in neural networks is to stack multiple layers and have the input pass through consecutive intermediate or *hidden* layers to reach the output layer (Fig. 3b). Such a design is known as a feedforward neural network, because inputs are computed and fed forward through the network to the output layer. A simple feedforward neural network has fully connected layers, such that each neuron in a layer is fully connected to the neurons in the following layer. Each neuron connection for a fully connected layer is permitted to have distinct weights.

Convolutional neural networks (LeCun & Bengio 1998) are a variant of feedforward neural networks in which the layer connections are constrained. This constraint comes in the form of *weight sharing*, where weights across neurons within a layer are constrained to only a fixed set of values, known as a *filter*. The content of the neurons in a convolutional layer is computed by sliding this filter across neurons in that layer (Fig. 3c). Hence, the filter is analogous to a kernel convolution. By using a fixed-length filter as weights instead of allowing each neuron connection to have their own distinct weights, features across a local ‘patch’ of data, for instance, the shape of an $l = 1$ mixed mode in Fig. 2b, can be learned to be detected by the network. Detected features across a layer are computed and then stored in a *feature map* (Rumelhart, Hinton & Williams 1988). Feature maps act as the hidden layer for convolutional layers, with the exception that whereas fully

Figure 3. A general schematic of neural network layers. (a) An elementary input–output connection in a neural network layer. Each neuron (circle) holds a real number. The activation function, f , maps the dot product of the weight vector, $\mathbf{w}$, and the input vector, $\mathbf{x}$, into a non-linear output. (b) Three fully connected feedforward neural network layers with one hidden layer and two output neurons. (c) Convolutional filtering applied to a 1D input. A kernel with weights (w_1 , w_2 , w_3) is moved over input neurons, three at a time. The weights across inputs here are fixed as a generic set of values for the feature map, whereas in (a), each weight may be unique.

Figure 4. General schematic of a convolutional neural network. The length of each array in this general schematic is arbitrary.

connected layers usually have a single array for its hidden layer, convolutional layers typically have more than one feature map because the detection of multiple features requires multiple feature maps. A feature map is given by:

$$\mathbf{g}^{(l)} = f\left(\sum_{i=0}^m \mathbf{w}^{(l)} \circ \mathbf{x}_i\right), \quad (1)$$

with m denoting the number of stars in the data set and l being the feature map index. Each feature map can be said to represent a detected local spatial feature of the data. However, in image recognition, as a greater number of convolutional layers are added, the outputs of deeper layers often become increasingly difficult to interpret from a human visual perspective (Zeiler & Fergus 2014).

A pooling layer is commonly applied after convolutional layers. Pooling reduces the length of convolutional layer outputs by clustering adjacent neurons into a fewer number of neurons over a specific function that compares the neuron values. To illustrate this, in our study, we use pooling layers that apply a 4-neuron *max-pooling*, which reduces a collection of four adjacent neurons into one neuron by selecting the neuron with the maximum value amongst them. In principle, this achieves a form of local spatial invariance within layers (Bengio 2013), while reducing the complexity of the neural

network at the same time because fewer weights are required for a smaller layer size. Following feature detection by multiple iterations of convolution and pooling layers, it is common for a fully connected layer to be connected for a final layer of computation before the output layer. A general structure of a convolutional neural network is illustrated in Fig. 4.

2.4 Learning and optimization

The objective of the convolutional neural network is to learn a particular set of weights from a training set that minimizes the error in approximating a ground truth y with a predicted output $\hat{y}$. In binary classification, y are binary values (associated with the binary classes; 0 for RGB and 1 for HeB) while $\hat{y}$ are two-element vectors with output scores for each class. We use the normalized exponential function known as the *softmax* function at the output layer, such that each output neuron contains the value:

$$p(y = j|\mathbf{x}) = \frac{e^{\mathbf{x} \cdot \mathbf{w}_j}}{\sum_{k=1}^K e^{\mathbf{x} \cdot \mathbf{w}_k}}, \quad (2)$$

which defines the score of class j out of $K = 2$ classes. These scores are similar to probabilities, however, probability calibration is usually required to express these scores in terms of actual population

probabilities. Since such work is beyond the scope of this study, it is sufficient to interpret the output scores as the prediction *likelihood* of a certain population class for a star. A score close to 1 indicates a high predicted likelihood for an HeB star, while a score close to 0 implies a high likelihood for an RGB star. We use a score threshold of 0.5 to assign the dominant population label of a target, such that a predicted score close to 0.5 implies a lack of classifier confidence in identifying the star’s population label as both populations ‘appear’ to be equally likely to the classifier. This score threshold has been found to provide a good separation of populations in our classifier.

A suitable error function to minimize is the *cross-entropy* or log-loss (Murphy 2012), related to the difference between a true value, y , and a predicted value, $\hat{y}$. Thus, the log-loss is the measure of similarity between y and $\hat{y}$. We compute the error, E , of the network by taking the average of all cross-entropies of the network’s predictions on m stars, such that

$$E(\mathbf{y}, \hat{\mathbf{y}}) = -\frac{1}{m} \sum_{i=1}^m \left[y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i) \right]. \quad (3)$$

The error minimization uses a gradient descent algorithm, where the error derivatives with respect to a layer’s weights are calculated from the end output and backpropagated to previous layers (Rumelhart, Hinton & Williams 1986). Layer weights are then updated sequentially over small subsets (mini-batches) of the training examples, which speeds up learning compared to updating over the full training set at once because it approximates the gradient and the error surface curvature (LeCun et al. 1998). We train our network with mini-batches of 128 stars. During the training process, multiple passes of feedforward and backpropagation are iterated until the error converges to a minimum. By training a convolutional neural network with a suitable optimization objective, we obtain a *classifier* (a fixed network with a trained set of fixed weights) that takes in our image as input and outputs a two-element vector $\hat{\mathbf{y}}$ with the vector elements as the RGB and HeB likelihoods, respectively.

2.5 Classifier structure and hyperparameters

A deep learning classifier usually has multiple stacks of neural network layers on top of one another as its structure. This structure contains a vast combination of free parameters (*hyperparameters*), which have to be empirically determined. Our classifier structure (Fig. 5) uses two parallel yet identical convolutional layer stacks such that both stacks see the same input. The sequence of Convolution-Pooling-Dropout (see explanation of dropout below) is carried out three times in succession before merging into a fully connected layer. We use a filter size of 32 in each convolutional layer with 2/4/8 feature maps for iterations 1/2/3. This structure is defined by choosing the simplest, computationally cheap combination of structure and hyperparameters that has the best metric performance on a *hold-out validation set*. This set is part of the initial training data that is now left out in training for classifier structure selection.

To prevent overfitting on training data, we use *dropout* layers (Hinton et al. 2012), which randomly sets layer outputs to zero with probability p_{drop} . This prevents layer outputs from memorizing the training data. We performed dropouts with $p_{\text{drop}} = 0.6$ after each convolutional layer and constrained the norm of the layer’s weight vector to a value of 2. This combination has been shown to be effective at preventing overfitting (Srivastava et al. 2014). We constructed the classifiers with the Keras library (Chollet 2015) built on top of Theano (Theano Development Team 2016). With a

Figure 5. Structure of our convolutional neural network. The number in brackets for fully connected layers is the input array length. The structure outlined by the dashed box is run three times in succession before merging the layer outputs by tensor concatenation.

Quadro K620 GPU, training with 200 training iterations requires approximately 15 min, while predicting on thousands of new stars takes only a few seconds.

3 RESULTS

3.1 Classifier performance

We report metrics on the mean of 10 separate hold-out validation sets, with each set having one-tenth of the training set in size (10-fold cross-validation), and on the test set. The metrics used to describe classifier performance are defined as follows:

Accuracy: The number of correct predictions out of all predictions.

Precision(P): For a class, the ratio of correct predictions to *all made predictions* towards that same class. Here it is the classifier’s ability not to label an HeB star as an RGB star.

Recall(R): For a class, the ratio of correct predictions to *all stars* truly in that same class. Here it is the classifier’s ability to find all HeB stars.

F1 Score: The harmonic mean of precision and recall, defined by $2 \frac{P \times R}{P + R}$, with 1 as a perfect score.

ROC AUC: Receiver Operating Characteristic’s Area Under Curve, which measures the classifier’s average performance across all possible score thresholds. Has a value of 1 for a perfect classifier (Swets, Dawes & Monahan 2000).

Log Loss: Negative logarithm of prediction scores i.e. the *cross entropy*. Measures how well prediction scores are calibrated with an ideal value of 0.

Precision, recall and the F1 scores complement accuracy in cases like ours where the population ratio is far from 50:50, while ROC AUC evaluates the overall classifier performance. Log loss reports the performance of class score outputs, with confident predictions rewarded low error when correct while penalized heavily otherwise. As observed in Table 1, the accuracy of the classifier on the cross-validation sets is generally above 98 per cent. On the test set where the classifier benefits from training on the entire training set, the classifier is capable of classifying with a 99 per cent accuracy and suffers a lower log loss. Having high values of precision, recall and

Table 1. Metrics over the mean of 10-fold cross-validation (CV) and over the test set.

Data set	CV (± 1 SD)	Test
Accuracy	0.982 ± 0.005	0.990
Precision	0.982 ± 0.005	0.990
Recall	0.982 ± 0.005	0.991
F1 Score	0.982 ± 0.005	0.991
ROC AUC	0.998 ± 0.002	0.996
Log Loss	0.055 ± 0.020	0.044

Figure 6. (a) $\epsilon - \Delta\nu$ diagram and (b) $\Delta P - \Delta\nu$ diagram of test set predictions. The colour corresponds to the score values of predictions, with deeper colours corresponding to a greater confidence towards a particular class (0 for RGB and 1 for HeB). Our classifications are available at: <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stx1174#supplementary-data>.

F1 score also indicates that the classifier is not heavily biased in predicting a particular population class that would not reflect the true population ratio.

Fig. 6 shows the test set results in $\epsilon - \Delta\nu$ (Huber et al. 2010; White et al. 2011) and $\Delta P - \Delta\nu$ (Bedding et al. 2011) diagrams. We derived the ϵ values using the method described in Stello et al. (2016a,b). One can see that ‘disputed’ predictions, namely predictions that are not in agreement with the ‘truth’ labels from V16 or M14, are more concentrated towards the low- $\Delta\nu$ regions. The classifier appears to be confident in most of its predictions (deep red and deep blue symbols), while most of the uncertain predictions are disputed. Upon inspection of the spectra of the 10 disputed stars, we visually verify that four of them, all with $2.9 \mu\text{Hz} < \Delta\nu < 5.2 \mu\text{Hz}$, had incorrect ground truth labels. Another four are confirmed to be due to the classifiers inaccuracy. These stars have $\Delta\nu > 7.0 \mu\text{Hz}$ in the diagrams. The final two stars are ‘high’ luminosity red giants with $\Delta\nu < 2.9 \mu\text{Hz}$. Visual inspection was inconclusive as the spectrum of one had suppressed dipole modes with a moderate level of noise, while the other appeared much like an RGB star but was previously given a late HeB classification as its ground truth. From theory, we do not expect to see a clear difference between RGB and late HeB stars because of the lack of coupling between core and envelope in such stars (Stello et al. 2013, fig. 4b).

Figure 7. $\epsilon - \Delta\nu$ diagram of the unclassified set for stars with $\Delta\nu \geq 2.8 \mu\text{Hz}$. The colour scheme is similar to that of Fig. 6. The region of the secondary clump stars is also indicated. The classifications for this plot are available at: <https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stx1174#supplementary-data>.

3.2 Classifying the unclassified set

We now use our trained classifier to predict the evolutionary state of the unclassified set (Fig. 7). It can be seen that the predictions reflect the $\Delta\nu - \epsilon$ relation of RGB stars (Kallinger et al. 2012) well for the entire $\Delta\nu$ range spanned by the training set ($2.8 \mu\text{Hz} \leq \Delta\nu \leq 18 \mu\text{Hz}$). The secondary clump of HeB stars is seen in the diagram with $\Delta\nu \simeq 6-9 \mu\text{Hz}$ below the RGB $\Delta\nu - \epsilon$ relation. In addition, the predictions also clearly show the HeB population at $\Delta\nu \simeq 3-4 \mu\text{Hz}$, with ϵ values mostly ranging about 0.7–1.0. In Figs 6a and 7, stars with $\Delta\nu \simeq 4 \mu\text{Hz}$ and $\epsilon \gtrsim 1.2$ are low ϵ stars that have ‘wrapped around’ vertically in the diagram.

The classifier has mostly very confident predictions, shown by the vast majority of predictions in Fig. 7 having deep shades of red and blue for RGB and HeB predictions, respectively. The exception to this is at the region of intersection between RGB and HeB populations at $2.8 \mu\text{Hz} \leq \Delta\nu \leq 4.0 \mu\text{Hz}$ and $0.9 \leq \epsilon \leq 1.0$. The paleness of the symbol colours in this region indicate a prediction uncertainty, which we attribute to the scarcity of stars in our training set with $\Delta\nu$ and ϵ values within this range.

Despite the predictions capturing the general distribution of red giant populations within the $\Delta\nu - \epsilon$ diagram, the classifier has its limitations in classifying based on image representation alone. For instance, it does not explicitly discriminate between frequency spacings, such that it can erroneously predict HeB stars at high $\Delta\nu$ ($\Delta\nu \gtrsim 9 \mu\text{Hz}$), where no HeB stars exist. However, only a very small fraction of predictions are subject to this inaccuracy. Another important limitation of these predictions is imposed by the parameter range of the training data. Our training data only includes RGB stars down to $\Delta\nu \simeq 2.8 \mu\text{Hz}$, hence we infer that the reliability of the classifier predictions also holds to a similar $\Delta\nu$ threshold. Due to this, we do not provide classifications for 1139 red giants with $\Delta\nu < 2.8 \mu\text{Hz}$ in our unclassified set, leaving behind 7655 red giants with $\Delta\nu \geq 2.8 \mu\text{Hz}$ from the initial 8794 in the unclassified set. Out of the remaining 7655 red giants, we find that 517 stars have been previously classified by Stello et al. (2013), from which 43 of their predictions do not agree with ours. We visually verify that they had incorrectly labelled approximately half of the disputed stars as HeB, with a majority of them having $\Delta\nu \simeq 5.5 \mu\text{Hz}$. In addition, 1991 stars with $\Delta\nu \geq 2.8 \mu\text{Hz}$ have been classified by

Elsworth et al. (2016), from which 232 stars are overlapping with the stars classified by Stello et al. (2013). Our classifier produces disputing predictions for 195 out of the 1991 stars classified by Elsworth et al. (2016). These disputes are mostly for those stars that they have predicted as RGB in the range of $2.8 \mu\text{Hz} \lesssim \Delta\nu \lesssim 5 \mu\text{Hz}$, which we find are split in roughly equal numbers into those that our classifier predicts correctly, those that our classifier predicts incorrectly and those where we are uncertain of the true population by visual inspection. In the end, we produce new classifications for 5379 stars not classified before by any other method. In future work, we will develop methods to improve the generalization of classifier predictions across a greater range of asteroseismic parameters. In addition, we will also look into spectral features of high $\Delta\nu$ RGB stars, which can deceive the classifier to incorrectly predict the star as HeB.

4 CONCLUSIONS

We have developed a convolutional neural network to perform fast and efficient classification of red giant stars into those ascending the RGB and those that have commenced core HeB. We presented folded oscillation spectra as images, which contain visual features that are learned by the convolutional neural network. Training and testing on *Kepler* data yielded a 98 per cent cross-validation accuracy and a 99 per cent test set accuracy, benchmarked against classifications based on asymptotic period spacing measurements. Out of the predictions that were in conflict with the ‘ground truth’, most scenarios of classifier inaccuracy were limited to the intermediate to high $\Delta\nu$ range, whereas for several low $\Delta\nu$ disputed cases, the input population labels were either incorrect or ambiguous based on visual inspection.

We also made predictions on 7655 *Kepler* red giants that do not have their asymptotic period spacing measured, from which 5379 have not been previously classified by any other means. We observed good agreement with the expected distribution of red giant populations in $\epsilon - \Delta\nu$ space for $\Delta\nu > 2.8 \mu\text{Hz}$. Despite being currently limited to predicting within the asteroseismic parameter ranges of the training set, this new, simple and effective method of classifying oscillation spectra seems promising for further future classifications on large data sets in asteroseismology.

ACKNOWLEDGEMENTS

Funding for this Discovery mission is provided by NASA’s Science Mission Directorate. We thank the entire *Kepler* team without whom this investigation would not be possible. DS is the recipient of an Australian Research Council Future Fellowship (project number FT1400147). We would also like to thank Timothy Bedding, Daniel Huber and the group at The University of Sydney for fruitful discussions.

REFERENCES

- Bedding T. R. et al., 2011, *Nature*, 471, 608
 Bengio Y., 2013, in Dediu A. H., Martin-Vide C., Mitkov R., Truthe B., eds, *Statistical Language and Speech Processing*. Springer, Berlin-Heidelberg, p. 1
 Borucki W. J. et al., 2010, *Science*, 327, 977
 Chollet F., 2015, Keras. Available at: <https://github.com/fchollet/keras>
 Christensen-Dalsgaard J., Silva Aguirre V., Elsworth Y., Hekker S., 2014, *MNRAS*, 445, 3685
 Dupret M.-A. et al., 2009, *A&A*, 506, 57
 Elsworth Y., Hekker S., Basu S., Davies G. R., 2016, *MNRAS*, 466, 3344

- Fukushima K., 1980, *Biol. Cybern.*, 36, 193
 Garcia C., Delakis M., 2004, *IEEE Trans. Pattern Anal. Mach. Intell.*, 26, 1408
 Grosjean M., Dupret M.-A., Belkacem K., Montalbán J., Samadi R., Mosser B., 2014, *A&A*, 572, A11
 Hinton G. E., Srivastava N., Krizhevsky A., Sutskever I., Salakhutdinov R. R., 2012, preprint ([arXiv:1207.0580](https://arxiv.org/abs/1207.0580))
 Huber D., Stello D., Bedding T. R., Chaplin W. J., Arentoft T., Quirion P.-O., Kjeldsen H., 2009, *Communications in Asteroseismology*, 160, 74
 Huber D. et al., 2010, *ApJ*, 723, 1607
 Kallinger T. et al., 2012, *A&A*, 541, A51
 Krizhevsky A., Sutskever I., Hinton G. E., 2012, in Pereira F., Burges C. J. C., Bottou L., Weinberger K. Q., eds, *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc., NY, p. 1097
 LeCun Y., Bengio Y., 1998, in Arbib M. A., ed., *Convolutional Networks for Images, Speech, and Time Series*. MIT Press, Cambridge, MA, p. 255
 LeCun Y., Bottou L., Orr G. B., Müller K.-R., 1998, in Grégoire M., Geneviève O., Klaus-Robert M., eds, *Efficient BackProp*. Springer-Verlag, London, p. 9
 Mosser B. et al., 2014, *A&A*, 572, L5 (M14)
 Murphy K. P., 2012, *Machine Learning: A Probabilistic Perspective*. MIT Press Ltd., MA
 Nair V., Hinton G. E., 2010, in Fürnkranz J., Joachims T., eds, *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, Rectified Linear Units Improve Restricted Boltzmann Machines. Omnipress, WI, p. 807
 Ricker G. R. et al., 2014, *J. Astron. Telescopes, Instrum. Syst.*, 1, 014003
 Rosenblatt F., 1962, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Report (Cornell Aeronautical Laboratory), Spartan Books
 Rumelhart D. E., Hinton G. E., Williams R. J., 1986, *Nature*, 323, 533
 Rumelhart D. E., Hinton G. E., Williams R. J., 1988, in Anderson J. A., Rosenfeld E., eds, *Learning Internal Representations by Error Propagation*, Neurocomputing: Foundations of Research, MIT Press, Cambridge, MA, p. 673
 Srivastava N., Hinton G., Krizhevsky A., Sutskever I., Salakhutdinov R., 2014, *J. Mach. Learn. Res.*, 15, 1929
 Stello D. et al., 2013, *ApJ*, 765, L41
 Stello D., Cantiello M., Fuller J., García R. A., Huber D., 2016a, *PASA*, 33, e011
 Stello D., Cantiello M., Fuller J., Huber D., García R. A., Bedding T. R., Bildsten L., Silva Aguirre V., 2016b, *Nature*, 529, 364
 Swets J. A., Dawes R. M., Monahan J., 2000, *Scientific American*, pp. 82
 Theano Development Team, 2016, preprint ([arXiv:1605.02688](https://arxiv.org/abs/1605.02688))
 Vrad M., Mosser B., Samadi R., 2016, *A&A*, 588, A87 (V16)
 White T. R., Bedding T. R., Stello D., Christensen-Dalsgaard J., Huber D., Kjeldsen H., 2011, *ApJ*, 743, 161
 Zeiler M. D., Fergus R., 2014, in Fleet D., Pajdla T., Schiele B., Tuytelaars T., eds, *Computer Vision – ECCV 2014, Visualizing and Understanding Convolutional Networks*. Springer International Publishing, NY, p. 818

SUPPORTING INFORMATION

Supplementary data are available at [MNRAS](https://www.mnras.org) online.

DeepLearningCatalogue(TestSet)

DeepLearningCatalogue(UnclassifiedSet)

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

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