

Analysis of medium resolution spectra by automated methods – Application to M 55 and ω Centauri^{*}

P. G. Willemsen¹, M. Hilker¹, A. Kayser^{1,2}, and C. A. L. Bailer-Jones³

¹ Sternwarte der Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany
e-mail: willemse@astro.uni-bonn.de

² Astronomisches Institut der Universität Basel, Venusstrasse 7, 4102 Binningen, Switzerland

³ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

Received 8 September 2004 / Accepted 8 February 2005

Abstract. We have employed feedforward neural networks trained on synthetic spectra in the range 3800 to 5600 Å with resolutions of $\sim 2\text{--}3$ Å to determine metallicities from spectra of about 1000 main-sequence turn-off, subgiant and red giant stars in the globular clusters M 55 and ω Cen. The overall metallicity accuracies are of the order of 0.15 to 0.2 dex. In addition, we tested how well the stellar parameters $\log g$ and T_{eff} can be retrieved from such data without additional colour or photometric information. We find overall uncertainties of 0.3 to 0.4 dex for $\log g$ and 140 to 190 K for T_{eff} . In order to obtain some measure of uncertainty for the determined values of $[\text{Fe}/\text{H}]$, $\log g$ and T_{eff} , we applied the bootstrap method for the first time to neural networks for this kind of parametrization problem. The distribution of metallicities for stars in ω Cen clearly shows a large spread in agreement with the well known multiple stellar populations in this cluster.

Key words. globular clusters: individual: ω Centauri – globular clusters: individual: M 55 – stars: fundamental parameters – methods: data analysis – techniques: spectroscopic – stars: abundances

1. Introduction

The advance in modern spectroscopic techniques coupled with 8-meter class telescopes makes it possible to simultaneously obtain spectra for a large number of stars in clusters. Although such spectra as analysed in this work do not have high enough resolution to make detailed line analyses, they offer the possibility to measure abundances with good precision (~ 0.2 dex). This permits, for example, statistically significant analyses of stellar populations in dwarf spheroidal galaxies or studies of the intra- and intercluster abundance variations of globular clusters.

While the data acquisition and subsequent reduction are themselves challenging, the analysis of such a large amount of data can be efficiently handled by the introduction of automated methods. This can be as accurate as spectral analyses (using line indices etc.) carried out by a human expert, as proven in several earlier works (see e.g. Snider et al. 2001; Soubiran et al. 1998). Several of these works show that the accuracy of the results was limited by the quality of the data rather than by the automated algorithm used to analyse them.

Recent studies using automated methods to retrieve stellar parameters T_{eff} , $\log g$ and $[\text{Fe}/\text{H}]$ from observed data were limited in various ways. (In the following we call this process

“parametrization”.) For example, Bonifacio & Caffau (2003) employed a minimum distance-based classifier to high resolution spectra of stars in the Sagittarius Dwarf Spheroidal but were restricted to certain luminosities and temperatures. Earlier efforts, especially those using neural networks, were also limited in their analysis of errors of the determined stellar parameters. Most often, the precision with which stellar parameters can be retrieved is evaluated by using a validation set with known target values. While this is correct to get some estimate of the overall uncertainty, this technique cannot be applied to previously unseen data, where such additional information is of course not available. However, any parameter estimate is only meaningful if some direct measure of uncertainty can be applied. To overcome this, we employed the bootstrap method to neural networks which allows us to assign individual error bars to the estimated stellar parameters.

Snider et al. (2001), using neural networks on spectral data with resolutions similar to those in this work, could determine all three stellar parameters (T_{eff} , $\log g$ and $[\text{Fe}/\text{H}]$) well. Indeed, they showed that the precisions are similar to those that can be obtained from fine analyses of high resolution spectra. While their neural network training and validation sets contained spectra from spectroscopic observations, our work relies entirely on synthetic spectra as templates. Of course, “real world” spectroscopic data with stellar parameters assigned from spectral synthesis analyses are the ideal case but are expensive in

^{*} Based on observations obtained at the European Southern Observatory, Chile (Obs. Prog. 67.D-0300 and 69.D-0172).

terms of observing time and effort. In this work we demonstrate that synthetic spectra in combination with automated methods can be used to determine reliable stellar parameters from spectroscopic observations.

We examine spectra from stars in the globular clusters M 55 and ω Centauri. These clusters serve as a demonstration of our methods. We have recently obtained a much larger database of stellar spectra in several globular clusters, and these will be subject to an extensive analysis in the future. M 55 has a metallicity of about -1.80 dex (Zinn & West 1984; Harris 1996; Richter et al. 1999) and is used primarily to validate our method. ω Cen, in contrast, is known to harbour several stellar populations with different kinematical characteristics and a metallicity range from -2.5 to -0.5 dex (see e.g. Norris & Da Costa 1995; Suntzeff & Kraft 1996), suggesting an entirely different formation history to that of normal globular clusters. Indeed, several earlier spectroscopic and photometric studies found a prolonged formation period of about 2 to 5 Gyr, suggesting that this object is more likely the nucleus of an accreted dwarf galaxy than a globular cluster (Lee et al. 1999; Hilker & Richtler 2000; Hughes et al. 2004; and especially Hilker et al. 2004). In addition, it was found that there is a clear trend of the α element abundances with the overall metallicity (see e.g. Pancino et al. 2002 or Origlia et al. 2003). The populations in ω Cen can be subdivided into a dominant metal-poor component with a metallicity peak at -1.7 dex which accounts for about 70% of the stars. Some 25% of the stars belong to the intermediate metallicity population with $[\text{Fe}/\text{H}] \sim -1.2$ dex while the last 5% are metal-rich stars with $[\text{Fe}/\text{H}] \sim -0.6$ dex (see e.g. Smith 2004).

The purpose of this work is to demonstrate the ability of automated methods (here neural networks) to derive accurate metallicities as well as surface temperatures and gravities from spectra with medium resolutions of $\sim 2\text{--}3$ Å. Sections 2 and 3 give a summary of the acquisition, reduction and selection of the observations, while Sect. 4 describes the set of synthetic spectra. The training procedure of the neural networks and the application of the bootstrap method are described in Sect. 5. In Sect. 6 the networks are validated on (observed) spectra of stars in M 55 (with known metallicity and $\log g$ and T_{eff} estimated from isochrones) as well as specific stars with known stellar parameters. The same networks are then applied to stellar spectra in ω Cen in order to derive preliminary metallicity values. Since this cluster is indeed very peculiar, we do not attempt to interpret these results in great detail. Rather, we want to show that the networks can yield reliable results across a large range of abundances for stars in different evolutionary states.

2. Observations and reductions

2.1. Data acquisition

The data were obtained on the nights 7–12 May 2002 at the VLT at ESO/Paranal (Chile) in visitor mode. The instrument was the FORS2 low dispersion spectrograph in combination with the mask exchange unit (MXU) and a 4×4 k MIT CCD with a field of view of 6.8 square arcminutes. On average, about 60 stars were observed per mask, with the

candidates chosen from Strömgren photometry (Hilker & Richtler 2000). The slit width was fixed to $1''$ and the slit length ranged from 5 to $7''$. The seeing was always better than $1''$. While Hilker et al. (2004) and Kayser et al. (2005) examined the main-sequence turn-off and SGB stars, the present work includes observations of the red giant and upper main sequence stars in both ω Cen and M 55. Each mask was observed through a “blue” grism (ESO 600I+25, second order) and a “red” grism (ESO 1400V+18) with maximum wavelength coverages of 3690 to 4888 Å and 4560 to 5860 Å. The corresponding dispersions are 0.58 Å pix^{-1} and 0.62 Å pix^{-1} , respectively. The exposure times were chosen in such a way as to ensure that the minimum signal to noise ratio (S/N) per pixel was $\sim 40\text{--}50$ for the majority of objects. For giant stars, the average S/N was slightly larger, about $\sim 70\text{--}80$. In the same way, the observations for M 55 were carried out with different masks for main-sequence, subgiant and giant stars. In order to have reference spectra from stars with known stellar parameters, long slit spectra with the same grism/filter combinations were obtained for eleven bright stars with a typical S/N of about 100–150. These stars are referred to as “standard stars” in the following. In addition to the science frames, bias, flat field and wavelength calibration images were obtained during daytime. The log of the observations reduced for this work is shown in Table 1. A summary and outline of the reductions and analysis based on line indices for the main sequence turn-off stars is found in Hilker et al. (2004) and Kayser et al. (2005). It should be noted that the sample of observed stars is not complete. The focus of these observations was to get a large sample of main-sequence turn-off and subgiant stars which are best suited to break the age-metallicity degeneracy. The upper main-sequence and red giant stars primarily serve to provide additional constraints for isochrone fitting of the different (age/metallicity) populations in ω Cen.

2.2. Data reduction

The data reduction of the observations was done by various tools within IRAF¹. The images of each mask were bias corrected, (dome) flat fielded and cleaned of cosmic rays. The image was then corrected for pixel-to-pixel variations and in a next step the individual spectra were extracted, both being performed by the “apall” task. The sky was generally subtracted using adjacent regions in the same slit, but in a few cases it was necessary to take the background information from other slit regions in the mask. In these cases, the sky correction was performed by the task “skytweak”. Although normally satisfactory, this procedure did result in erroneous sky subtractions in some cases. We therefore excluded spectra reduced with “skytweak” from the analysis. Each stellar spectrum was wavelength calibrated by 10 to 15 lines of He, Ne, Hg and Cd and rebinned to a final dispersion of 1 Å pix^{-1} , yielding an effective resolution of 2 to 2.5 Å.

¹ IRAF is distributed by the National Optical Astronomy Observatories, which are operated by the Association of Universities for Research in Astronomy, INC., under cooperative agreement with the National Science Foundation.

Table 1. Table of the observations reduced and analysed for this work in addition to those reported by Kayser et al. (2005). The different field names in ω Cen are abbreviated by their location in the cluster (e.g om-*e* corresponds to *east*). The ending “rgb” refers to red giant stars, “to” to main-sequence turn-off stars, “ms” correspondingly to upper main-sequence stars. The field of view of a MXU mask is $6'.8 \times 6'.8$.

Date	Mask	α_{2000}	δ_{2000}	Exp. time (red) [s]	Exp. time (blue) [s]
2002/05/10	om-n-rgb	13:26:46.7	-47:18:43.3	$1 \times 120, 1 \times 240$	2×300
2002/05/11	om-e-rgb	13:27:48.1	-47:27:27.4	2×180	2×240
2002/05/11	om-s1-rgb	13:26:57.4	-47:37:43.1	2×180	2×240
2002/05/11	om-n-ms	13:26:46.7	-47:18:43.2	3×1080	3×1200
2002/05/12	om-s2-rgb	13:26:57.3	-47:43:42.9	2×180	2×240
2002/05/12	om-s2-ms	13:26:57.3	-47:43:40.0	3×960	3×1200
2002/05/12	om-s1-ms	13:26:57.8	-47:37:43.7	3×900	3×1200
2002/05/12	om-e-ms	13:27:48.1	-47:27:30.8	3×900	3×1200
2002/05/10	M 55-to	19:39:59.1	-30:53:01.2	3×600	3×720
2002/05/11	M 55-rgb	19:39:58.9	-30:52:56.4	2×180	2×240
2002/05/12	M 55-ms	19:39:58.9	-30:52:59.0	2×900	2×900




Fig. 1. Examples of observed spectra. The left panels show the spectra as obtained by the blue grism, the right column that of the red grism. The upper panel shows spectra of a red giant star, the lower for a main sequence star in M 55, i.e. $[\text{Fe}/\text{H}] = -1.8$ dex. Note the larger number of absorption lines in the spectrum of the cool giant. In the lower plot the wavelength ranges used for the neural network analysis are indicated.

The continuum normalization of the spectra is crucial, since it can directly influence the parametrization performance. We tried several techniques to divide out the continuum. These included different combinations of line exclusion, median filtering techniques and continuum fitting functions. It was found that excluding the regions adjacent to the strongest absorption lines in combination with the IRAF task “continuum” and careful adjustment of the various parameter settings therein yielded the best results. But we note that a good continuum fit is naturally harder to find for metal-rich stars ($[\text{Fe}/\text{H}] \sim -0.5$ dex) due to heavy line blending. This should be kept in mind when interpreting the parametrization results for the more metal-rich stars. Furthermore, the S/N in some cases is not the same for the two grisms for the same object. We did not explicitly test how this influences the overall parametrization but the results suggest that this effect is of little importance.

Since the slits cover a range of positions on the mask, the wavelength coverage is not the same for all stars. In order to produce a homogeneous grid of all spectra, we restricted the wavelength range from 3850 to 5700 Å with an overlap of the red and blue grisms’ contributions at 4700 Å. Moreover, for the analysis only certain wavelength ranges in the spectra were used, the regions being defined after extensive simulations on synthetic spectra. These are 3900–4400 Å, 4820–5000 Å and 5155–5350 Å. As expected, these regions, found by comparing the parametrization results for networks trained on different ranges, contain the most prominent metal and hydrogen lines. The restriction to certain wavelength intervals and the exclusion of (hopefully) unimportant ranges minimizes the number of inputs to the network, thus increasing the ratio of training templates to free parameters in the network (see Sect. 5). Examples of typical spectra are shown in Fig. 1.

Table 2. Summary of the stars used for validating the network model. The $[\text{Fe}/\text{H}]$, $\log g$ and T_{eff} values given are the averages from the works of several authors found in the literature (see Cayrel de Strobel et al. 2001). The uncertainties are either the standard deviations about the mean values (in the case that multiple values are listed) or the errors stated by the author (when there is only a single measurement). The last column shows the number of separate determinations in each case.

ID	$\overline{[\text{Fe}/\text{H}]}$ [dex]	$\sigma([\text{Fe}/\text{H}])$ [dex]	$\overline{\log g}$ [dex]	$\sigma(\log g)$ [dex]	$\overline{T_{\text{eff}}}$ [K]	$\sigma(T_{\text{eff}})$ [K]	num
BD-032525*	-1.90	0.08	4.00	0.57	5739	40	2
BD-052678	-1.89	–	4.43	–	5570	–	1
BD-084501	-1.59	0.27	4.00	–	5750	210	2
BD-133442	-2.95	0.21	3.92	0.25	6240	75	3
HD089499**	-2.14	0.08	2.92	0.88	4875	100	3
HD097916	-1.05	0.19	4.09	0.10	6222	140	5
HD179626	-1.26	0.04	3.70	0.06	5625	40	1
HD192718	-0.74	0.04	4.00	0.06	5650	40	1
HD193901	-1.08	0.12	4.38	0.32	5713	70	8
HD196892	-1.02	0.10	4.04	0.27	5900	100	4
HD205650	-1.12	0.13	3.93	0.64	5729	90	3

* Spectroscopic binary.

** Star in binary system.

3. Grid of observed spectra

We determined stellar radial velocities by cross-correlating the spectra with suitably chosen template spectra. With these data at hand, we corrected the spectra of each star individually in each grism (blue and red). The radial velocities were also used as an extra indicator for cluster membership (in addition to the position of the star in the colour–magnitude diagram) by comparing them to the velocities of the clusters. These are 175 km s^{-1} and 238 km s^{-1} for M 55 and ω Cen, respectively (Harris 1996). In order to account for systematic effects not corrected in this data reduction (such as small mismatches between the wavelength calibration and science images), we chose rather large velocity intervals for the cluster membership decision, namely $[100:200] \text{ km s}^{-1}$ for M 55 and $[150:350] \text{ km s}^{-1}$ for ω Cen. In this way about 6% of all stars were excluded as being nonmembers.

The distributions of the velocity errors from the red and blue grism are shown in Fig. 2. Each error combines the individual cross-correlating error for a given spectrum (as determined from the IRAF routine “fxcor”) and that of the corresponding template stars.

Some additional spectra were eliminated because of the presence of light polluting bright stars next to the slit. However, two effects coming from bright stars that are close to the observed field can cause erroneous sky subtraction. First, there may be reflections that cause gradients in the background, and second, there may be a contamination in the wings of the PSF. Finally, some low quality spectra were also eliminated.

We split our observed stars into two luminosity class categories based on the available Johnson photometry: MSSGB (comprising all main-sequence and subgiant stars) and RGB (all red giant stars). This cut is sensible since the stellar parameters $\log g$ and T_{eff} of these two groups are significantly different. Moreover, in order to improve the neural network performance, it is generally recommended to simplify the problem as much as possible (see e.g. Haykin 1999). More explicitly, by splitting the regression problem (the mapping from inputs




Fig. 2. Distributions of the velocity errors of the red grism (left column) and blue grism (right column) for the spectra of M 55 (top panels) and ω Cen (bottom panels). The errors in the red grism are somewhat larger because of a smaller number of prominent absorption lines available for the cross-correlation procedure.

to outputs) into two parts (in terms of stellar parameters), we hope that the overall approximation of the underlying function by the network is improved by lifting possible ambiguities in the mapping function. Tests showed that the results are indeed better when considering two object categories instead of one. This is certainly also due to the different S/N of the MSSGB and RGB objects. As described in Sect. 2, the giants have generally higher S/N than the subgiants and main sequence stars. It was therefore necessary to train networks on different ranges in S/N (see below). For M 55 there are 130 stars in the MSSGB and 35 stars in the RGB sample, while for ω Cen the numbers are 800 and 66, respectively.

Since we have synthetic training samples only for stars up to the helium flash, we removed all horizontal branch and suspected asymptotic giant branch stars in both our M 55 and ω Cen samples. We will consider the parametrization of these in a future study, which will also include an analysis of α and CN abundances.

4. Grid of synthetic spectra

We have calculated synthetic (template) spectra using the model atmospheres from Castelli et al. (1997) in combination with the latest version of SPECTRUM (Gray & Corbally 1994) along with its line list. Since our study focuses on globular clusters, the metallicity range of the calculated spectra was limited to $-2.5 \text{ dex} \leq [\text{Fe}/\text{H}] \leq -0.5 \text{ dex}$, while temperature and surface gravity are limited to $3500 \text{ K} \leq T_{\text{eff}} \leq 8500 \text{ K}$ and $0 \text{ dex} \leq \log g \leq 5.0 \text{ dex}$ respectively. Note that the grid is not complete, i.e. not all parameter combinations exist. The spectra were calculated with a step size of 0.02 \AA in the wavelength range 3850 to 5700 \AA . The microturbulence was fixed at 2 km s^{-1} . Similarly to the procedure of Bailer-Jones et al. (1997), the spectra were dispersion-corrected and binned to have the same resolution and wavelength range as the observed data.

As will be described in Sect. 5, we used spectra with different S/N to train the network models. To do so, we added the corresponding noise for a given S/N before normalizing the spectra. The wavelength of scaling the S/N was chosen at 4500 \AA which is in the overlap region of the two grisms.

For the present tests we considered solar scaled α abundances, something which is empirically found to be only the case for halo objects with metallicities $> -1 \text{ dex}$ (see e.g. McWilliam 1997). However, this does not pose a problem since the parametrizer is explicitly trained only on metallicity as expressed by $[\text{Fe}/\text{H}]$ (and not on α abundances). The regression model (the neural network) is therefore expected to focus on the changes as caused by different $[\text{Fe}/\text{H}]$ values, thus ignoring spectral dependencies caused by other specific elements.

5. Artificial neural networks

Neural networks have proven useful in many scientific disciplines for analysing data by providing a nonlinear mapping between an input domain (in this case the spectra) and an output domain (the stellar parameters). For details on this subject with special emphasis on stellar parametrization see, for example, Willemsen et al. (2003), Bailer-Jones et al. (1997), Bailer-Jones (2000), Snider et al. (2001), and especially Bailer-Jones (2002) for a general overview. The software used in this work is that of Bailer-Jones (1998).

A neural network can be visualized by several layers, where each layer is made up of several so-called nodes or neurons. In general, for the type of network used in this work, one distinguishes an input-layer, one or two hidden-layers, and an output-layer. The nodes in each layer are connected to all the nodes in the preceding and/or following layer via some free parameters (“weights”). While the nodes in the input layer simply hold the inputs to the network (i.e. the fluxes in the wavelength bins of the spectrum), each node in the hidden layer

performs a weighted summation of all its inputs. This sum is then passed through a nonlinear “transfer function” (here the \tanh function) and the result is then passed further to a node in the next layer. The output nodes do the same as the nodes in the hidden layer, except that the transfer function is here chosen to be linear. Before the network can be applied to predict the stellar parameters for previously unseen (observed) spectra, it must be trained on an existing set of (here synthetic) spectra with known stellar parameters. During this training phase, the weights of the network are iteratively adjusted in order to accurately predict the outputs (stellar parameters), based on the inputs (spectral fluxes). This weight adjustment is done according to some measure of error (“cost function”). Here we use the sum-of-squares error calculated from the networks’ outputs and the corresponding true stellar parameters for a given synthetic spectrum. Once this error is sufficiently small, the training is stopped and the weights are fixed. The network can now be used to predict the stellar parameters of a star based on the fluxes in a stellar spectrum.

There are several parameters in a neural network which must be adjusted, often by careful experiments. In the model used here these are essentially a weight decay term (which inhibits overfitting by a too complex network), a weighting term for each stellar parameter in the cost function, the number of hidden layers, as well as the number of weights in each hidden layer and the number of training iterations. After extensive tests we found that a network with two hidden layers each having 13 neurons provided the best results. More neurons did not significantly improve the results but only increased the training time, while a smaller number of neurons resulted in a poorer model with larger errors.

In principle, one could train three independent network models, each with a single output, for each of the three stellar parameters (T_{eff} , $\log g$ and $[\text{Fe}/\text{H}]$). It is well known, however, that changes in the values of these stellar parameters may have similar effects on a stellar spectrum (both lines and continuum). A special example is the helium abundance, where an increase of this element changes the line shapes in the same way as does surface gravity (Gray 1992). To account for such degeneracies we therefore trained the networks on all three stellar parameters simultaneously. As expected from the above, but contrary to what Snider et al. (2001) found, a network with multiple outputs performed better than one trained only on one single parameter. A possible reason for this might be that we used two hidden layers. It is generally known that a second hidden layer is important to model the “subtle” information in the data (here metallicity and surface gravity), see e.g. Bishop (1995). Networks with two hidden layers might therefore be better suited for this kind of parametrization problem.

5.1. Network training and regularization

The observed spectra cover a range of different S/N . The neural network model (or the training data) should account for this. Indeed, we found that networks trained on templates with high S/N and then validated on spectra with a much lower S/N increased the bias and (to a smaller degree) also the variance

Table 3. Summary of the regularization results from neural networks with different training sets and weight decay term γ . Increasing γ results in smaller weights of the network and thus a more regularized solution. However, too large a value of γ will again result in larger deviations, i.e. there is a trade-off in setting this parameter. The metallicity deviations for the standard stars are given in terms of the median values of the difference (computed value – literature value). It can be seen that training on noise-free data and validating on noisy data systematically underestimated metallicities. The results demonstrate that noise in the network inputs can help improve the regularization.

Noise in training set	γ	[Fe/H] offset
no	0.0001	-0.16
no	0.001	-0.11
no	0.01	-0.07
yes	0.0001	0.18
yes	0.001	0.02
yes	0.01	0.05

of the distribution of residuals. This behaviour was also found in tests by Snider et al. (2001). To investigate this further, we trained networks on different training sets with multiple noise versions of a spectrum (only one noisy spectrum at each S/N) for a given stellar parameter combination. In addition, we tried different values for the weight decay (regularization) parameter γ . We found that a network trained only on high S/N templates ($S/N = 100$) can reproduce the stellar parameters for spectra with different S/N reasonably well, but only if the weight decay term is chosen to be rather large. This shows that this is a regularization problem, i.e. a network trained on high S/N spectra will overfit the data unless “restrained”. When using multiple noisy versions of a training template the results improved slightly, which is understandable since additional noise in the training data serves as an additional regularization mechanism (noise being hard to fit). The results of these tests are summarized in Table 3.

We therefore trained the networks on sets of noisy versions of the training templates (synthetic spectra). A training sample of size N is a collection of several stellar parameter combinations expressed by \mathbf{y}_i ($i \in [1, N]$) and associated spectral flux vectors \mathbf{x}_i , where each of these is represented P times with different S/N ratios, i.e.

$$\begin{pmatrix} (\mathbf{x}_1; \mathbf{y}_1) \\ (\mathbf{x}_2; \mathbf{y}_2) \\ \vdots \\ (\mathbf{x}_N; \mathbf{y}_N) \end{pmatrix} \text{ where } \mathbf{x}_i = \begin{pmatrix} \mathbf{x}_i^{S/N_1} \\ \mathbf{x}_i^{S/N_2} \\ \vdots \\ \mathbf{x}_i^{S/N_P} \end{pmatrix}. \quad (1)$$

For the MSSGB case (main sequence and subgiants) we used synthetic spectra with $3.5 \text{ dex} \leq \log g \leq 5.0 \text{ dex}$, $4000 \text{ K} \leq T_{\text{eff}} \leq 8750 \text{ K}$, $-2.5 \text{ dex} \leq [\text{Fe}/\text{H}] \leq -0.5 \text{ dex}$ and scaled to S/N ratios between 10 and 100, in steps of 10. For the RGB stars, the stellar parameters were $0.0 \text{ dex} \leq \log g \leq 3.5 \text{ dex}$, $3500 \text{ K} \leq T_{\text{eff}} \leq 6000 \text{ K}$ and the same metallicity range but with S/N ratios in the range from 50 to 150, in steps of 10. Note that there is an overlap in $\log g$ and T_{eff} for the two samples which is necessary to properly test those stars with parameters at the edge of the grids.

As mentioned above, the networks can also be tuned by a weighting term for each stellar parameter in the cost function. We used this term to set up two network models as follows.

modell is a network tuned to yield good metallicity parametrization results. This network was trained on all three stellar parameters, i.e. $[\text{Fe}/\text{H}]$, $\log g$ and T_{eff} , but with special weight on metallicity.

model2 is a network tuned to yield good $\log g$ and T_{eff} performance. As *modell*, this network was trained on all stellar parameters, but with a higher weight on surface gravity and temperature.

Note that each model was trained on each of the two training sets (MSSGB and RGB) separately. Since the minimization process can become trapped in local minima during learning, we set up a committee of 10 networks (with different initial weight settings) for each of the models.

As described in Sect. 2.2, we used three disjoint wavelength ranges from the spectrum, yielding a total of 878 wavelength bins (network inputs). Including the “bias” weights in the input and hidden layers, the total number of weights is therefore 11 648. This is large and it might be surprising that there are meaningful results given that there are only 4990 and 4818 training templates for the MSSGB and RGB cases, respectively. However, the effective number of free parameters in the network is certainly much lower, because much of the spectrum contains redundant information. Furthermore, the smoothness of the input-output function which is to be approximated by the network will also lower the effective number of free parameters.

5.2. Error estimates from Bootstrapping neural networks

In most applications of neural networks to determine stellar parameters, uncertainties (error bars) on the estimated parameters are derived from the distribution of the parameter residuals of some validation set with known target values. While this method is correct in order to get an overall estimate of the precision to which parameters can be determined, it does not yield individual uncertainties for each determined parameter. This is crucial, however, since objects that lie close to the boundary of the training grid tend to have higher uncertainties, given that the performance of neural networks (or indeed any regression model) is generally weak at the grid boundaries (see e.g. Bishop 1995). The concept of using the bootstrap method for the determination of error bars in the framework of neural networks is discussed in several articles (e.g. Heskes 1997; Dybowski & Roberts 2000). For the special case of stellar parametrization see Willemsen et al. (2004).

Since a neural network is a regression model that maps inputs onto outputs, we can use the concepts of standard errors and confidence intervals also for this kind of parametrizer. There are basically two sources of uncertainty arising with neural networks. The first stems from the fact that the training data are noisy and incomplete, i.e. the construction of a training grid by (randomly) sampling templates is already prone to sampling variations. The second source of uncertainty is given by the

model limitations which arise from the optimization failing to find the global optimum. Another source of error is an inappropriate (e.g. inflexible) parametrization model (e.g. linear). Since we assume that our network model approximates the underlying mapping from the inputs to the outputs reasonably well, the latter source of uncertainty is not considered further.

The bootstrap was introduced by Efron (1979) for estimating various sample properties such as bias, variance and confidence intervals on any population parameter estimate. Given a (random) training sample S comprising pairs of inputs x and corresponding outputs y , i.e. $S = (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$, drawn from a population F , we want to estimate some parameter (e.g. the network weights) $\theta = f(F)$. This could be done by calculating $\hat{\theta} = g(S)$ based on S . The bootstrap is a data-based method for statistical inference which allows us to estimate the error in the network outputs from different values of $\hat{\theta}$. To determine this bootstrap standard error, one needs to build bootstrap samples. A bootstrap sample, S^* , is a random sample of the same size N as the original sample which is created by randomly resampling S with replacement (i.e. $S^* \subseteq S \subset F$). In this way, one obtains B bootstrap samples, the b th bootstrap sample given by $(x_1^{*b}, y_1^{*b}), (x_2^{*b}, y_2^{*b}), \dots, (x_N^{*b}, y_N^{*b})$. For each sample, we minimize $\sum_{n=1}^N [y_n^{*b} - y(\theta; x_n^*)]^2$ (i.e. we derive B regression functions by training a network on each of the B bootstrap samples) to yield $\hat{\theta}^{*b}$ in each case. The (nonparametric) estimate of the bootstrap standard error, BSE, is the square root of the variance of the distribution, which for the n th predicted value y_n is

$$\text{BSE}(x) = \sqrt{\frac{1}{B-1} \sum_{b=1}^B [y(x_n; \hat{\theta}^{*b}) - y_{\text{boot}}(x_n; \cdot)]^2} \quad (2)$$

where $y_{\text{boot}}(x_n; \cdot)$ is shorthand for the bootstrap committee's prediction given by $\frac{1}{B} \sum_{b=1}^B y(x_n; \hat{\theta}^{*b})$.

Note that the random resampling is done over the stellar parameters and not over the noise versions.

It is generally recommended that the number of bootstrap samples be in the range from 25 to 200 (Efron & Tibshirani 1993). For this work, we chose $B = 30$ to save computational power. Limited tests showed that larger numbers did not significantly change the results. The random resampling was done with a uniform random number generator as given in Press et al. (1992).

6. Validation of the network model

In the following, the results of the network validation are presented. As mentioned before, the metallicity results were obtained from network *modell* which was tuned to yield accurate results for this parameter. The results for surface gravity $\log g$ and temperature T_{eff} were found from network *modell2*, i.e. a network trained on all three stellar parameters as outputs but with a higher weight on $\log g$ and T_{eff} (see Sect. 5.1).

Unless stated otherwise, we summarize the overall parametrization results in terms of the average absolute deviation (over some set of spectra), i.e.

$$A = \frac{1}{N} \cdot \sum_{p=1}^N |C(p) - T(p)| \quad (3)$$

where p denotes the p^{th} spectrum and T is the target (or “true”) value for this parameter. The quantity $C(p)$ is the parametrization output averaged over a committee of 10 networks (see Sect. 5).

6.1. Results for M 55

In order to compare the networks' outputs of surface gravity and temperature based on the observed spectra, we have to get an estimate of the “true” values for these parameters. For this, we did not use the pre-imaging *BV* Johnson photometry (because of missing photometric standard star observations and saturated giant stars), but rather decided to use more precise measurements of M 55 obtained in the Strömgren system. These data were obtained (in addition to observations in other filters) at the Danish 1.5 m telescope at La Silla on the nights 12–15 July 2001. The instrument's field of view is $6'.3 \times 6'.3$, the seeing was typically about $1''.5$. A summary of the data reduction is given in the next section.

6.1.1. Strömgren photometry of M 55

Initially the photometric calibration, relating instrumental to standard magnitudes, was determined. For this we used observations of 12 E region stars (Jönch-Sørensen 1993) which were reduced in a standard way (bias subtracted and flat fielded) and the instrumental magnitudes determined via aperture photometry. The extracted magnitudes were then combined in one calibration equation that takes the airmass of the images into account. The science images for a given filter were bias subtracted and flatfielded in a standard way. The corrected science frames were then averaged to increase the signal to noise ratio. Instrumental magnitudes were derived via PSF photometry with the aid of the packages DAOPHOT and ALLSTAR (Stetson 1987). A summary of the observations is given in Table 4. The calibrated data were finally extinction corrected using the transformation $E(b - y) = 0.7 \cdot E(B - V)$ (Crawford & Barnes 1970) and $E(v - y) = 1.68 \cdot E(b - y)$ as calculated from a synthetic extinction curve taken from Fitzpatrick (1999) in combination with the Strömgren filter transmission curves. Upper limits for the photometric errors are 0.02 mag for the y and b bands and 0.03 mag for the v band. Since the fields of the MXU and the Danish 1.5 m observations only overlapped in certain parts, photometric information was not available for all objects.

6.1.2. Surface gravities and temperatures for stars in M 55

To determine $\log g$ and T_{eff} from the Strömgren photometry, we measured the distances in the three-dimensional v, b, y space between the observed magnitudes and the corresponding filter magnitudes of an appropriate isochrone. For this we used the isochrone taken from Bergbusch & Vandenberg (1992) (transformed to Strömgren colours by Roberts & Grebel 1995) with an age of 16 Gyr and $[\text{Fe}/\text{H}] = -1.80$ dex with an assumed cluster distance of 5.3 kpc and a reddening of

Table 4. Table of Strömgren observations of M 55. The columns show the central coordinates of the observed fields and the exposure times in the different filters. Certain parts of these fields overlap with those of the FORS2 spectroscopic observations.

Date	Field	α_{2000}	δ_{2000}	Exp. time(v) [s]	Exp. time(b) [s]	Exp. time(y) [s]
2001/07/12	M 55-1	19:40:19.0	-30:59:30.0	$2 \times 30, 3 \times 240$	$1 \times 20, 2 \times 90$	$3 \times 10, 3 \times 60$
2001/07/13	M 55-2	19:40:19.0	-30:53:30.0	$1 \times 30, 2 \times 240$	$1 \times 20, 2 \times 90$	$1 \times 10, 2 \times 60$
2001/07/14	M 55-3	19:39:47.1	-30:56:01.1	$1 \times 30, 2 \times 240$	$1 \times 20, 2 \times 90$	$1 \times 10, 2 \times 60$


$E(B - V) = 0.08$ mag (Harris 1996). We found that these parameters could best fit the $(b - y):y$ and $(v - y):y$ colour magnitude diagrams. The reason why we used Strömgren isochrones which are still on the old age scale is that we had no access to newly calculated models in this photometric system.

We then “projected” each observed star onto the isochrone by calculating the distances between the tabulated values of v , b and y of the isochrone and the corresponding observed magnitudes. A star was then assigned the $\log g$ and T_{eff} value of the isochrone for which the calculated distance in this three-dimensional magnitude space is minimum. Stellar parameter uncertainties $\Delta(\log g)$ and $\Delta(T_{\text{eff}})$ are found by defining a (three-dimensional) error ellipse from the photometric errors $\Delta(v)$, $\Delta(b)$ and $\Delta(y)$ and by calculating the distance to the isochrone for equally spaced steps on this ellipsoid. The errors of $\log g$ and T_{eff} are then found from the resulting maximum and minimum values for each parameter. This procedure is visualized in Fig. 4. Note that because of the slope of the isochrone and the position of a particular star in the colour magnitude diagram, these errors are not symmetric.

Note that the age value of the Strömgren isochrone is offset by 2.5 Gyr from the 13.5 Gyr as determined by the best fitting (α -enhanced) isochrone for Johnson photometry (Fig. 3). To check for consistency, we did tests with the available pre-imaging Johnson photometry and the $[\alpha/\text{Fe}] = 0.3$ dex Johnson isochrone (taken from Kim et al. 2002). We found that there were significant differences (of about ~ 80 K) in the T_{eff} values only for the upper main sequence and turn-off region stars, while the $\log g$ differences were of the order of 0.1 to 0.2 dex. Since we only want to demonstrate the principal capability of the networks to derive surface gravity and temperature estimates, we did not investigate this further.

The parametrization performances for $\log g$ and T_{eff} for stars in M 55 for which Strömgren photometry was available are presented in Fig. 5. Note that these results were obtained from network *model2*, i.e. networks trained with a higher weight on $\log g$ and T_{eff} . In order to appreciate the results one should keep in mind that the networks were trained on synthetic spectra with a single microturbulence value of $\xi = 2 \text{ km s}^{-1}$. However, as Gray et al. (2001) showed, a range of different microturbulences is important in order to correctly estimate surface gravity values. It should also be noted that the T_{eff} and $\log g$ values as estimated from the isochrone have uncertainties of their own. This is especially the case for those stars that lie close to the main-sequence turn-off, where a perpendicular projection of a star onto the isochrone (see above) gives rise to larger uncertainties.

From the plots, it can be seen that the parametrization yields reasonable results. The scatter is somewhat larger for

**Fig. 3.** Photometry in M 55. The upper plot is the pre-imaging Johnson BV photometry for a selected range of luminosities. Two isochrones both with 13.5 Gyr and $[\text{Fe}/\text{H}] = -1.8$ dex are overplotted. The left one has solar scaled α abundances, the right one has $[\alpha/\text{Fe}] = 0.3$ dex. The lower plot shows $v - y$ Strömgren photometry and an isochrone with $[\text{Fe}/\text{H}] = -1.8$ dex and 16 Gyr (the old age scale). The heavy points show the stars for which spectra have been obtained.

high $\log g$ and T_{eff} values, which is most likely due to the fact that the values determined from the isochrone are uncertain in these regimes. For T_{eff} we observe a small systematic offset of about 80 K. Given the intrinsic uncertainties of the target values and the fact that these are only limited tests (i.e. we did not put too much effort into finding the optimum network configurations), we conclude that the results are close to what can be obtained from such data in general (see Sect. 6.3).

6.1.3. Metallicity results for M 55

The metallicity values determined from the neural networks for the combined MSSGB and RGB sample are shown in the left panel in Fig. 6. These results were obtained from network *model1*, i.e. networks trained on all stellar parameters but with a higher weight on metallicity, see Sect. 5.1. The overall




Fig. 4. A sketch to help visualize the procedure to find $\log g$ and T_{eff} values from the Strömgren photometry. The solid line is an isochrone in the three-dimensional magnitude space, while the ellipsoid is a three-dimensional error box about the measured point. The size of the error ellipsoid has been magnified by a factor 15 and displaced for clarity. The surface gravity and temperature for the measured point is found from the shortest distance between this point and the isochrone. The errors for these two stellar parameters are found by projecting the error ellipsoid onto the isochrone and taking the difference between the extreme values of T_{eff} and of $\log g$. This is shown in the two-dimensional projection on the b, v plane.

parametrization performance is good. The mean value of the metallicity distribution is -1.86 dex compared with a literature value of -1.80 dex. The overall error as defined in Eq. (3) is 0.15 dex.

6.2. Model validation for standard stars

We compared the parameter determinations of the neural networks on standard stars with those found from other methods and listed in the catalogue of Cayrel de Strobel et al. (2001). In cases of multiple observations (from different authors) in this catalogue, we calculated the (unweighted) mean and the corresponding standard error for those values which were determined from high resolution spectroscopy. A summary of the values is given in Table 2. The network models were then validated on the spectra of these stars (see Sect. 2). We only used the networks trained for the MSSGB sample for this test (see Sect. 3). The results for $\log g$ and T_{eff} were obtained from network *model2* with a higher weight (during minimization) on surface gravity and temperature, while the results for $[\text{Fe}/\text{H}]$ are from network *model1* with a higher weight on metallicity, see Sect. 5.1. The parametrization results are summarized in Fig. 7.

From the plots we see that the metallicity results are quite good, with an overall error of less than 0.15 dex. Interestingly,

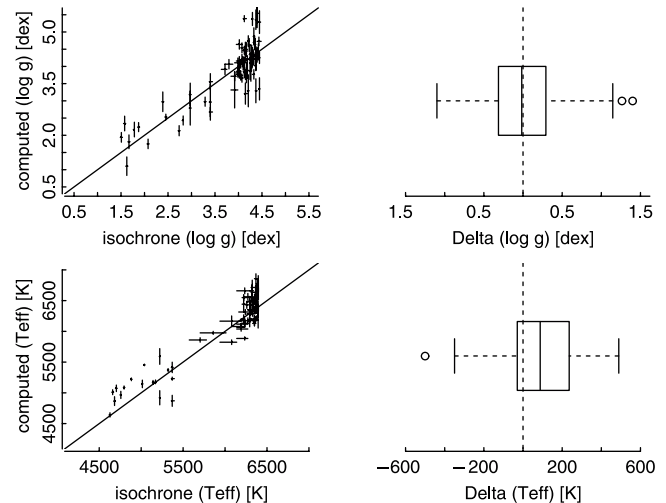


Fig. 5. Parametrization results for $\log g$ (upper line) and T_{eff} (lower line) for stars in M55. The left panels shows the parameter values as determined from the isochrone versus the corresponding computed value from the network for the two parameters. The solid line is the identity. The horizontal and vertical error bars for each point are calculated from the photometric error ellipse (see Fig. 4) and from 30 bootstrap networks, respectively. The right panel shows the boxplots of the residuals for each parameter, $\Delta = (\text{computed value} - \text{isochrone value})$. The solid line in the box shows the median value of the residuals while the borders of the box are the first and third quartiles. Single points which have a larger distance to the median than 1.5 times the boxlength and are identified as outliers. The overall parametrization errors as defined in Eq. (3) are 0.32 dex and 190 K for the two parameters.




Fig. 6. Distribution of the metallicities for stars (combined MSSGB and RGB sample) in M55 (left panel) and ω Cen (right panel) as determined by our method. The dashed line is centred at -1.80 dex (the literature value of $[\text{Fe}/\text{H}]$ for M55) while the solid line denotes the mean of the distribution. There is only a small offset of 0.06 dex. The overall error for the determined metallicities in M55 as defined in Eq. (3) is 0.15 dex.

the overall parametrization performance for $\log g$ is not as good as for the stars in M55, despite the fact that the S/N of the standard stars is much higher (but there are only 11 stars so the statistics is rather poor). Given that the stellar parameters of the standard stars are similar to those of the main-sequence stars in M55, and from the fact that the input spectral ranges

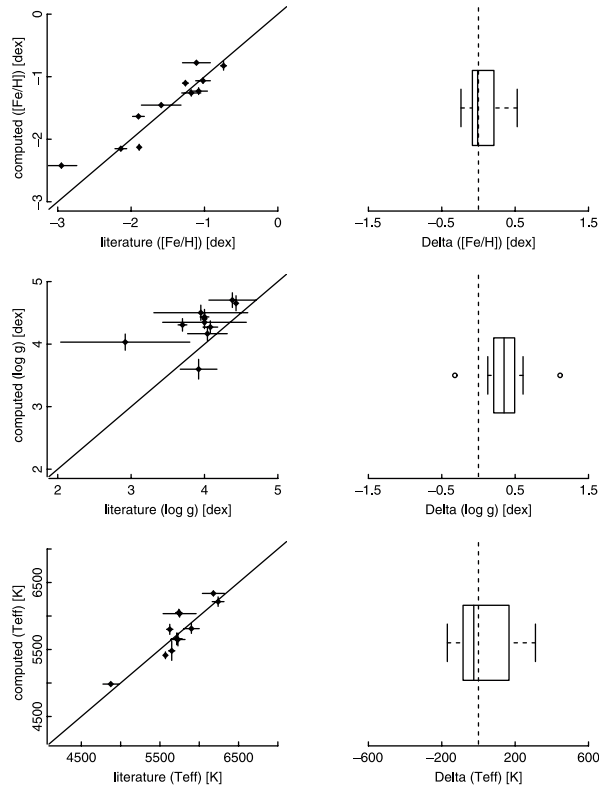


Fig. 7. As in Fig. 5 but for the “standard” stars showing the network outputs compared to the averaged literature values for each of the parameters $[\text{Fe}/\text{H}]$, $\log g$ and T_{eff} (from top to bottom). The errors for the literature values are found from the scatter about the unweighted mean value for the different stated values (see Table 2), while the computed errors are found from 30 bootstrap networks. Note that the metallicity range on which the networks were trained was limited to values larger than -2.5 dex, so we cannot expect good parametrization performances for metallicities below this value. Excluding the single point with the literature value $[\text{Fe}/\text{H}] = -2.95$ dex, we find an overall average error in $[\text{Fe}/\text{H}]$ as defined in Eq. (3) of 0.15 dex, while that for $\log g$ and T_{eff} are 0.42 dex and 145 K respectively.

to the networks are identical for both types of stars, we conclude that these larger deviations are, paradoxically, related to the larger S/N . As outlined above, the S/N of the template spectra with which the networks were trained influences the degree of regularization, thus effecting the extent of systematic offsets (see Sect. 5). Indeed, the S/N of the template spectra was in the range from 10 to 100, i.e. lower than that of the standard stars. To investigate this further, we performed limited tests with networks trained on spectra with S/N values in the range from 100 to 150 while keeping the weight decay parameter fixed. We found that the systematic offset in $\log g$ for the standard stars could indeed be decreased in this way (albeit not completely removed), while the variance remained almost the same. While these results support the idea of insufficient regularization, such a strong sensitivity of the network training to the S/N of the template spectra is surprising and needs to be investigated in more detail in future studies. An explanation of why there are no systematic trends for $[\text{Fe}/\text{H}]$ is possibly given by the fact that the networks for this parameter were highly tuned, in order to eliminate inaccuracies.




Fig. 8. The spectroscopically observed stars in ω Cen are marked as heavy dots in the Strömgren colour magnitude diagram from Hilker & Richtler (2000). Based on the radial velocities, filled circles were identified as members, while open circles denote non-members.

Taking the stellar parameters predicted by the network for a given observed spectrum, we can generate the corresponding synthetic spectrum. Figure 9 shows a comparison between two (observed) standard star spectra and the appropriate synthetic spectrum produced by interpolating the training sample using the stellar parameters as predicted by the neural networks. It can be seen that the model spectra represent the observed spectra rather well, especially in the case that the network’s predictions of the stellar parameters are close to the literature values.

6.3. Discussion of the parametrization results

The overall parametrization performance can be compared with that of Snider et al. (2001). They trained networks on real spectral data with similar resolutions and for stars with similar parameter ranges as in this project. They reported precisions of the parameter determinations for $\log g$ of 0.25 to 0.30 dex and 135–150 K for T_{eff} while that for $[\text{Fe}/\text{H}]$ was 0.15 to 0.20 dex. Our results are in absolute agreement with these values, even though our networks were trained on synthetic data.

6.4. Metallicity results for ω Cen

We applied the neural networks to the MSSGB and RGB samples of ω Cen to determine metallicities for these stars. We stress, however, that the metallicities found are somewhat preliminary and we do not attempt to give a more detailed analysis. The complex formation and evolution history of ω Cen will ultimately require a much more sophisticated approach in which synthetic spectra with different helium and/or α element abundances are taken into account. The results were obtained from network *model1* which was tuned to improve metallicity performance. A plot of all stars observed is shown in Fig. 8.




Fig. 9. A comparison of two observed standard star spectra (solid lines) with the corresponding synthetic spectra (dotted lines) as predicted by the neural network. The network predicts only the stellar parameters, but from these the synthetic spectra were generated through interpolation of the training template grid of synthetic spectra. The upper plot is for a star with stellar parameter residuals of $\Delta([\text{Fe}/\text{H}]) = -0.04$ dex, $\Delta(\log g) = 0.12$ dex and $\Delta(T_{\text{eff}}) = -89$ K, where Δ is the difference (network output – literature value). The corresponding network outputs are $[\text{Fe}/\text{H}] = -1.06$ dex, $\log g = 4.16$ dex and $T_{\text{eff}} = 5811$ K. The lower plot shows an example where the difference between the $\log g$ prediction of the network and the literature value is rather large. For this star, we find $\Delta([\text{Fe}/\text{H}]) = -0.01$ dex, $\Delta(\log g) = 1.11$ dex and $\Delta(T_{\text{eff}}) = 109$ K and networks outputs of $[\text{Fe}/\text{H}] = -2.15$ dex, $\log g = 4.03$ dex and $T_{\text{eff}} = 4984$ K. The S/N of the synthetic spectra was decreased to 150, i.e. similar to that of the observed spectra.

Note that we only considered the stars up to the top of the giant branch in this study (see Sect. 3), i.e. there are no HB or AGB stars. These will be examined in a future study.

The overall distribution of the metallicities found by our neural network analysis for ω Cen is shown in the right panel of Fig. 6. When compared to the distribution of metallicities found for M 55, we note a significantly larger spread. *From the bootstrap errors for the metallicity values plus the smaller overall metallicity spread found for the stars in M 55 (Fig. 6), we must conclude that this large spread does not indicate an erroneous parametrization but rather a real effect of the multiple stellar populations in ω Cen.*

7. Concluding remarks

We have shown that the metallicity information in medium resolution spectra can be readily assessed by neural networks trained on synthetic spectra and without additional photometric information. The neural network approach works in an objective way and is, ultimately, time efficient, relevant when large sets of spectra of stars in different objects (clusters) are to be analysed. In order to obtain some measure of uncertainty on the stellar parameters determined, we applied the bootstrap method to neural networks for the first time for this kind of parametrization problem.

The estimated metallicity accuracies are of the order of ~ 0.2 dex which is close to what one can expect from such data with any method. We additionally performed limited tests on the retrieval of $\log g$ and T_{eff} and found that the overall accuracies for these parameters are ~ 0.3 to 0.4 dex and ~ 150

to 190 K. These results could probably be further improved once a suitable grid of synthetic spectra with a range of different microturbulent velocities is available. In addition, the effects of α abundances should be taken into account to test how well this additional parameter can be determined and whether neural networks are able to discriminate between the effects of $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ from normalized spectra alone. Moreover, and especially with regard to ω Cen, the recent suggestions of enhanced helium abundances for the intermediate metallicity population (Norris 2004) should be taken into account in future simulations of the synthetic templates. As mentioned above, an increased helium abundance results in an overall increased molecular weight which affects the absorption lines in similar ways to $\log g$.

Future tests could also include photometric data (e.g. from the pre-imaging observations). Since temperature strongly affects the continuum of a stellar energy distribution, any colour information should improve the performance for this parameter. Once a larger set of observed spectra for stars in normal globular clusters (i.e. with a well defined metallicity) is available, we will train networks on these data. Although synthetic spectra can reproduce observed stellar energy distributions rather well, we believe that a network trained on real data with similar noise and detector characteristics will result in better metallicity estimates for previously unseen data.

Acknowledgements. P.G.W. is grateful to the DLR for support under the grant 50QD0103. We thank R.O. Gray for distributing the SPECTRUM code. We also thank K. S. de Boer and T. A. Kaempfer for helpful discussions, O. Cordes and O. Marggraf for assistance with the computer system level support and the referee for helpful comments.

References

- Bailer-Jones, C. A. L. 1998, Statnet – a feedforward interpolation neural network, Tech. rep., <http://www.mpia-hd.mpg.de/homes/calj/statnet.html>
- Bailer-Jones, C. A. L. 2000, *A&A*, 357, 197
- Bailer-Jones, C. A. L. 2002, in *Automated Data Analysis in Astronomy*, ed. R. Gupta, H. Singh, & C. A. L. Bailer-Jones (New Delhi, India: Narosa Publishing House), 51
- Bailer-Jones, C. A. L., Irwin, M., Gilmore, G., & von Hippel, T. 1997, *MNRAS*, 292, 157
- Bergbusch, P. A., & Vandenberg, D. A. 1992, *ApJS*, 81, 163
- Bishop, C. 1995, *Neural Networks for Pattern Recognition* (Oxford University Press)
- Bonifacio, P., & Caffau, E. 2003, *A&A*, 399, 1183
- Castelli, F., Gratton, R. G., & Kurucz, R. L. 1997, *A&A*, 318, 841
- Cayrel de Strobel, G., Soubiran, C., & Ralite, N. 2001, *A&A*, 373, 159
- Crawford, D. L., & Barnes, J. V. 1970, *AJ*, 75, 978
- Dybowski, R., & Roberts, S. J. 2000, in *Clinical Applications of Artificial Neural Networks*, ed. R. Dybowski, & V. Gant (Cambridge University Press)
- Efron, B. 1979, *Ann. Statist.*, 7, 1
- Efron, B., & Tibshirani, R. 1993, *An Introduction to the Bootstrap* (New York: Chapman and Hall)
- Fitzpatrick, E. L. 1999, *PASP*, 111, 63
- Gray, D. F. 1992, *The observation and analysis of stellar photospheres* (Cambridge Astrophysics Series)
- Gray, R. O., & Corbally, C. J. 1994, *AJ*, 107, 742
- Gray, R. O., Napier, M. G., & Winkler, L. I. 2001, *ApJ*, 121, 2148
- Harris, W. E. 1996, *AJ*, 112, 1487
- Haykin, S. 1999, *Neural Networks – A Comprehensive Foundation* (Upper Saddle River, New Jersey: Prentice Hall)
- Heskes, T. 1997, in *Advances in Neural Information Processing Systems*, ed. M. C. Mozer, M. I. Jordan, & T. Petsche (The MIT Press), 9, 176
- Hilker, M., Kayser, A., Richtler, T., & Willemsen, P. 2004, *A&A*, 422, L9
- Hilker, M., & Richtler, T. 2000, *A&A*, 362, 895
- Hughes, J., Wallerstein, G., van Leeuwen, F., & Hilker, M. 2004, *AJ*, 127, 980
- Jönch-Sørensen, H. 1993, *A&AS*, 102, 637
- Kayser, A., Hilker, M., Richtler, T., & Willemsen, P. G. 2005, *A&A*, in prep.
- Kim, Y., Demarque, P., Yi, S. K., & Alexander, D. R. 2002, *ApJS*, 143, 499
- Lee, Y.-W., Joo, J.-M., Sohn, Y.-J., et al. 1999, *Nature*, 402, 55
- McWilliam, A. 1997, *ARA&A*, 35, 503
- Norris, J. E. 2004, *ApJ*, 612, L25
- Norris, J. E., & Da Costa, G. S. 1995, *ApJ*, 447, 680
- Origlia, L., Ferraro, F. R., Bellazzini, M., & Pancino, E. 2003, *ApJ*, 591, 916
- Pancino, E., Pasquini, L., Hill, V., Ferraro, F. R., & Bellazzini, M. 2002, *ApJ*, 568, L101
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C* (Cambridge University Press)
- Richter, P., Hilker, M., & Richtler, T. 1999, *A&A*, 350, 476
- Roberts, W. J., & Grebel, E. K. 1995, *Bull. Amer. Astron. Soc.*, 27, 816
- Smith, V. V. 2004, in *Origin and Evolution of the Elements*, ed. A. McWilliam, & M. Rauch (Carnegie Observatories Astrophysics Series), 188
- Snider, S., Allende Prieto, C., von Hippel, T., et al. 2001, *ApJ*, 562, 528
- Soubiran, C., Katz, D., & Cayrel, R. 1998, *A&AS*, 133, 221
- Stetson, P. B. 1987, *PASP*, 99, 191
- Suntzeff, N. B., & Kraft, R. P. 1996, *AJ*, 111, 1913
- Willemsen, P. G., Bailer-Jones, C. A. L., & Kaempf, T. A. 2004, *Analysis of Stellar Parameter Uncertainty Estimates from Bootstrapping Neural Networks*, Tech. rep., Gaia-ICAP, ICAP-PW-004
- Willemsen, P. G., Bailer-Jones, C. A. L., Kaempf, T. A., & de Boer, K. S. 2003, *A&A*, 401, 1203
- Zinn, R., & West, M. J. 1984, *ApJS*, 55, 45