

Assembling a high-precision abundance catalogue of solar twins in GALAH for phylogenetic studies

Kurt Walsen,^{1,2,3*} Paula Jofré,^{2,3} Sven Buder^{1b,4,5}, Keaghan Yaxley^{1b,6}, Payel Das^{1b,7}, Robert M. Yates^{1b,8}, Xia Hua^{1b,9}, Theosamuele Signor,^{3,10} Camilla Eldridge,^{2,3} Alvaro Rojas-Arriagada^{1b,2,11}, Patricia B. Tissera^{1b,2,12,13}, Evelyn Johnston^{1b,2,3}, Claudia Aguilera-Gómez,¹² Manuela Zoccali^{1b,12,14}, Gerry Gilmore^{15,16} and Robert Foley^{1b,17}

¹Departamento de Ingeniería Matemática, Facultad de Ciencias Físicas y Matemáticas, Universidad de Chile, Av. Beauchef 851, 8370458 Santiago, Chile

²Millennium Nucleus ERIIS

³Instituto de Estudios Astrofísicos, Facultad de Ingeniería y Ciencias, Universidad Diego Portales, Av. Ejército 441, 8370191 Santiago, Chile

⁴Research School of Astronomy & Astrophysics, Australian National University, Canberra ACT 2611, Australia

⁵ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

⁶Research School of Biology, Australian National University, Canberra ACT 2601, Australia

⁷Physics Department, University of Surrey, Guildford GU2 7XH, United Kingdom

⁸Centre for Astrophysics Research, University of Hertfordshire, Hatfield AL10 9AB, United Kingdom

⁹Mathematical Sciences Institute, Australian National University, Canberra ACT 0200, Australia

¹⁰Inria Chile Research Center, Av. Apoquindo 2827, piso 12, 7550312 Las Condes, Santiago, Chile

¹¹Departamento de Física, Universidad de Santiago de Chile, Av. Víctor Jara 3659, 9170124 Santiago, Chile

¹²Instituto de Astrofísica, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 782-0436 Macul, Santiago, Chile

¹³Centro de Astro-Ingeniería, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 782-0436 Macul, Santiago, Chile

¹⁴Millennium Institute of Astrophysics, Av. Vicuña Mackenna 4860, 782-0436 Macul, Santiago, Chile

¹⁵Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹⁶Institute of Astrophysics, FORTH Cert, N. Plastira 100, GR-70013 Vassilika Vouton, Crete, Greece

¹⁷Leverhulme Centre for Human Evolutionary Studies, Department for Anthropology and Archaeology, University of Cambridge, Fitzwilliam Street, Cambridge CB2 1QH, UK

Accepted 2024 January 19. Received 2023 December 18; in original form 2023 August 24

ABSTRACT

Stellar chemical abundances have proved themselves a key source of information for understanding the evolution of the Milky Way, and the scale of major stellar surveys such as GALAH have massively increased the amount of chemical data available. However, progress is hampered by the level of precision in chemical abundance data as well as the visualization methods for comparing the multidimensional outputs of chemical evolution models to stellar abundance data. Machine learning methods have greatly improved the former; while the application of tree-building or phylogenetic methods borrowed from biology are beginning to show promise with the latter. Here, we analyse a sample of GALAH solar twins to address these issues. We apply *The Cannon* algorithm to generate a catalogue of about 40 000 solar twins with 14 high precision abundances which we use to perform a phylogenetic analysis on a selection of stars that have two different ranges of eccentricities. From our analyses, we are able to find a group with mostly stars on circular orbits and some old stars with eccentric orbits whose age–[Y/Mg] relation agrees remarkably well with the chemical clocks published by previous high precision abundance studies. Our results show the power of combining survey data with machine learning and phylogenetics to reconstruct the history of the Milky Way.

Key words: methods: data analysis – techniques: spectroscopic – catalogues – stars: abundances – Galaxy: evolution.

1 INTRODUCTION

Chemical enrichment plays an important role in the formation and evolution of our Galaxy. Significant advances in our understanding of galaxy evolution have come from interpreting data from major stellar spectroscopic surveys, such as the Galactic Archaeology with

HERMES (GALAH) Survey (De Silva et al. 2015; Buder et al. 2021), the Gaia–ESO Survey (Gilmore et al. 2022; Randich et al. 2022) and the Sloan Digital Sky Survey (Abdurro’uf et al. 2022). These surveys have observed hundreds of thousands of stars for which chemical abundances of up to 30 elements have been measured (see extensive discussion in Jofré, Heiter & Soubiran 2019). However, while this increase of data is essential, it is also important to develop tools for analysing larger data sets. In this paper, we are concerned with the presentation and assessment of such methods

* E-mail: kurt.walsen.b@gmail.com

in two ways – first, the extraction of chemical information from stars via stellar spectroscopy, and secondly, the analysis of chemical abundances to retrieve an evolutionary history of Galactic chemical enrichment.

Galactic history shares some features with biological evolution. Indeed, Jofré et al. (2017) and Jackson et al. (2021) showed that phylogenetic techniques can be employed to reconstruct the evolution of star formation within the solar neighbourhood using a small but very precise sample of stellar abundances of solar twins. Interpreting their results was difficult because of selection effects in their stellar samples, which were a rather small sample of spectra obtained from public archives (Nissen 2016; Bedell et al. 2018). While major spectroscopic surveys retrieve data on millions of stars observed with a well-defined selection function, there is a trade-off between the scale of data and its resolution. This study is thus motivated by carefully creating a sample of high-precision chemical abundances from survey data. We then use a phylogenetic tree algorithm to assess the formation sequences of the Milky Way. First, though, we discuss the potential and dangers of using machine learning tools (ML) to extract higher precision chemical abundances from stellar spectra for this purpose and then we introduce phylogenetic methods.

1.1 Chemical abundances from stellar spectra

In this era of large spectroscopic surveys, ML has become a revolutionary way to both, precisely and quickly derive spectral properties (see e.g. Leung & Bovy 2019; Wheeler et al. 2020; Ambrosch et al. 2023, to name a few). The revolution became notable perhaps with the introduction of *The Cannon* (Ness et al. 2015) into the field. The method finds a polynomial function that directly relates the spectra with labels. To do so, the function is found from stellar spectra for which these labels are known. *The Cannon* is very fast, and provides more precise results than standard methods when the spectra are noisy or not of very high resolution. Since it is easy to implement, it has quickly been applied on a large variety of stellar spectra (Casey et al. 2017; Buder et al. 2018; Wheeler et al. 2020; Nandakumar et al. 2022; Manea et al. 2023). Many other spectral analyses using ML have followed (Leung & Bovy 2019; Ting et al. 2019; Guiglion et al. 2020). Nowadays, there is a large variety of chemical data products derived from neural networks or mathematical functions trained on synthetic or observed spectral grids. Since ML allows labels to be transferred from one survey to the other (Wheeler et al. 2020; Nandakumar et al. 2022), it is possible to put several surveys on the same scale provided there are stars in common between surveys (Ness 2018). Indeed, ML is so powerful that today it might seem absurd to aim performing a spectroscopic analysis on million spectra with the ‘standard’ methods (see discussion of such methods in Jofré et al. 2019).

Although methods like *The Cannon* provide many advantages over the standard methods, it relies on training sets that are calibrated with the standard analyses. Choosing the training set is not straightforward, it should sample fully and evenly the parameter space of the test set, and should have accurate and precise values for the labels that want to be determined. A grid of synthetic spectra is powerful because it ensures even sampling (Ting et al. 2019), but synthetic spectra and real spectra are different from each other leading to a level of uncertainty in the results (see discussion in O’Brian et al. 2021). ML methods are suitable when the methods used on the observed test stars are the same as for the training set stars, especially if there is information additional to the spectra (interferometry, accurate parallaxes, astroseismology, etc.). This extra information provides higher accuracy for the labels (Miglio et al. 2013; Jofré et al. 2014;

Heiter et al. 2015). The problem here, though, is the sampling. ML works better with large data sets, and so using the largest number of stars as possible is preferable. However, the full parameter space is not often available in these very large data sets, and so there is a trade-off between sample size and data completeness. It is thus worth investigating systematically how to train an ML algorithm with the more limited data available today.

1.2 Phylogenetic trees as a promising tool to trace Galactic chemical evolution

Phylogenetic trees are graphs that illustrate the shared evolutionary history among a data set, allowing us to understand the hierarchical pattern of ancestry and descent which connects all of the observations (Baum, Smith & Donovan 2005). Phylogenetic methods can reconstruct ancestral relationships as long as there is a shared history and a heritable process linking the data objects. These objects are normally individuals, species, and higher taxa in biology, where methods to analyse them have been developed (Felsenstein 1988), but they are applicable more broadly. By making the hypothesis that the stars in the Milky Way disc come from the same but evolving interstellar medium (ISM), and that the evolutionary marker (i.e. the heritable component) of the ISM is the chemical composition, we can use the chemical abundances of low-mass stars as fossil records for building phylogenetic trees (see also Freeman & Bland-Hawthorn 2002).

The hypothesis that stars in the Milky Way form from the same evolving ISM is a simplification of reality. Indeed, the Milky Way has accreted dwarf galaxies, depositing in the ISM some gas that has been enriched by a different chemical evolutionary history. An example is the interaction of the Milky Way with Sagittarius, which has affected the star formation history of the disc significantly (Ruiz-Lara et al. 2020). This process can be interpreted as ‘hybridization’, which in biology stands for the result of horizontal gene transfer, that is, genetic information passed between species and not only in an ancestor-descendant way. The impact on phylogenetic studies when hybridization is present is an active topic of research.

Another simplification of this hypothesis is the fact that chemical abundances in low-mass stars are not as constant over their lifetime. On one hand, heavy elements sink in the atmospheres due to gravitational settling (Lind et al. 2008; Souto et al. 2018), causing an effect in the measured abundances depending on the age and the mass. On the other hand, newly processed heavy elements can be dredged up to the atmosphere because of mixing processes inside stars (e.g. Aguilera-Gómez, Jones & Chanamé 2023).

Furthermore, the way the ISM mix after the ejection of new chemical elements is not homogeneous, making new stellar populations not necessarily identical in their chemical pattern. All these issues however have an impact on chemical tagging studies overall.

Phylogenetic trees have already been constructed in Jofré et al. (2017) and Jackson et al. (2021). These papers focused on solar twins for the practical reason that estimates of chemical abundances in solar twins are very precise, particularly if they are derived differentially with respect to the Sun (e.g. Nissen & Gustafsson 2018). This implies that differences in the chemical compositions are not necessarily caused by a systematic effect in spectroscopic analysis. Jofré et al. (2017) used high precision data published by Nissen (2015, 2016) and Jackson et al. (2021) the data published by Bedell et al. (2018). The trees were built using a nearest neighbourhood distance method, which essentially considers the pairwise distance in chemical abundances between stars to find the hierarchical differences, displaying them in a tree. Jofré et al. (2017) found a tree with different branches where the relationship

Table 1. Wavelength grid for the spectra used in this work.

CCD	Begin/Å	End/Å	Dispersion/Å/pixel	Nr. pixels
1	4715.94	4896.00	0.046	3915
2	5650.06	5868.25	0.055	3968
3	6480.52	6733.92	0.064	3960
4	7693.50	7875.55	0.074	2461

between branch length and age was different, suggesting that with trees it might be possible to identify stellar families (i.e. groups of stars that cluster together and so may have a shared history), but more importantly, study their different evolutionary processes such as chemical enrichment rate. Jackson et al. (2021) followed that study by enlarging the number of stars and by choosing elemental abundance ratios which evolve with time, e.g. the so-called *chemical clocks* (e.g. Nissen 2016; Casali et al. 2020; Jofré, Jackson & Tucci Maia 2020). They also found different branches that had different ages and dynamical distributions, and attributed them to different stellar groups co-existing in the solar neighbourhood. How far these stellar groups were representative of the broader stellar population remains uncertain due to the selection function.

1.3 Aim and structure of this study

In this work, we make the first study of a phylogenetic tree from survey data, for which we chose a set of solar twins from GALAH DR3 (Buder et al. 2021). We compare our tree using the published abundances from GALAH and a set of abundances obtained using ML, for which the precision is higher. The idea of this comparison is to initiate the literature regarding the precision needed in chemical abundances to perform phylogenetic studies with stars (also see recent discussions in de Brito Silva et al. 2023, using simulated data). To do so, we first apply the spectral fitting machinery of *The Cannon* to GALAH data (Buder et al. 2021) to improve the precision of the chemical abundance measurements. We systematically assess the steps in training *The Cannon* in solar twins.

GALAH has observed a very large sample of solar twins (to date about 40 000). By carefully applying *The Cannon* we provide a sample of high-precision abundances for a much larger sample of solar twins than those previously published from high-resolution data. This takes the sample size from around 500 stars (Casali et al. 2020) to two orders of magnitude. We then select a sample of this catalogue to test if the phylogenetic signal improves from GALAH to *The Cannon* abundances for the same stars.

The paper is structured as follows: In the next section, we describe the data used, and this is followed by a description of how *The Cannon* training is set up. Our new catalogue is presented in Section 3 and we apply phylogenetic techniques to analyse this catalogue in Section 4. Some general conclusions are presented in Section 5.

2 DATA AND METHODS

In this work, we use stellar spectra, parameters, and abundances published as part of the third data release of GALAH (Buder et al. 2021, hereafter, GALAH DR3). Stellar spectra were observed with the HERMES spectrograph (Sheinis et al. 2015) at the Anglo-Australian Telescope with a resolution of $\sim 28\,000$ across four wavelength regions in the optical (see Table 1) that include absorption features of more than 30 elements. For this study, we use the radial velocity corrected and normalized spectra downloaded from

the DATACENTRAL¹ web interface and interpolate them onto a common wavelength scale given in Table 1. The normalized spectra that are provided on DATACENTRAL as part of GALAH DR3 have been velocity corrected and normalized with the routines as described in section 2.3 of Buder et al. (2021) and based on the GUESS code by J. Lin as part of the reduction pipeline (Kos et al. 2017). In short, a cross-correlation with the best-fitting spectrum of the AMBRE grid (De Laverny et al. 2012) is performed to shift the observed spectra to the rest frame. Subsequently, third order (CCDs 1 and 2) and fourth order (CCDs 3 and 4) polynomial fits are performed for pre-selected continuum regions while applying sigma clipping. Flux uncertainties are derived from the provided relative error spectra and multiplied with the flux.

We use the stellar parameters (T_{eff} , $\log g$, and $[\text{Fe}/\text{H}]$) as well as logarithmic elemental abundances relative to the Sun and iron, that is, $[\text{X}/\text{Fe}]$. These were extracted from the GALAH spectra via χ^2 optimization of on-the-fly computed synthetic spectra with Spectroscopy Made Easy (SME; Valenti & Piskunov 1996; Piskunov & Valenti 2017). Synthetic spectra were computed based on 1D MARCS model atmospheres (Gustafsson et al. 2008) and non-local thermodynamic equilibrium for eleven elements (Amarsi et al. 2020) and local thermodynamic equilibrium for all other elements. As laid out in section 3.3 of Buder et al. (2021) with references, thirteen elements H, Li, C, O, Na, Mg, Al, Si, K, Ca, Mn, Fe, and Ba have been measured from GALAH spectra. The optimization included the constraint of surface gravities $\log g$ from bolometric luminosity estimates based on photometric information from the 2MASS survey (Skrutskie et al. 2006) and distances inferred from the *Gaia* satellite's second data release (Bailer-Jones et al. 2018; Gaia Collaboration et al. 2018).

The GALAH DR3 catalogue further provides age estimates, which are needed for evolutionary studies. These ages and their uncertainties are estimated via isochrone interpolation with stellar parameters through the Bayesian fitting machinery *BSTEP* (Sharma et al. 2018). Valued-added catalogues in GALAH also include information about orbital properties, such as total velocities, eccentricities and actions. For more details we refer the reader to Buder et al. (2021).

2.1 Abundance uncertainties

The total error budget in GALAH considers a combination of precision and accuracy uncertainties. Precision uncertainties are calculated from the internal covariance uncertainties of SME, that is, the uncertainties are computed from the diagonal of the covariance matrix given by SME fitting procedures, which were adjusted to be consistent with the scatter of repeat observations as a function of the signal-to-noise ratio (SNR). For this work, we consider the precision to be the maximum between both uncertainties described above. Accuracy uncertainties for the stellar parameters are derived from comparisons with reference stars, such as the *Gaia* FGK benchmark stars (Jofré et al. 2014, 2018; Heiter et al. 2015), whereas for the abundances we only have precision uncertainties reported. We thus consider for this work the GALAH DR3 precision uncertainties, which only report the final abundance uncertainty based on the maximum of the internal covariance error for the abundance fit and the SNR response to repeat observations (Buder et al. 2021).

¹<https://datacentral.org.au/services/download/>

2.2 Solar twin selection

We are interested in the solar twins of GALAH DR3. We select them from the file called `GALAH_DR3_main_allspec_v2.fits2`. This file contains measurements of astrophysical parameters and chemical abundances of 678 423 spectra from 588 571 stars derived as explained above. We consider only spectra that have a signal-to-noise ratio (SNR) above 10 in all 4 CCDs. In order to select the solar twins from that sample, we used the reported astrophysical parameters and performed the following cuts:

$$\begin{aligned} |T_{\text{eff}} - 5777| &< 200 \text{ K} \\ |\log g - 4.44| &< 0.3 \\ |[\text{Fe}/\text{H}]| &< 0.3 \end{aligned} \quad (1)$$

where we assume that the solar temperature is 5777 K, the solar surface gravity is 4.44 and the logarithmic number density of iron to hydrogen compared to the sun, $[\text{Fe}/\text{H}]$, is 0.0 (Prša et al. 2016). We note our selection is wider than Bedell et al. (2018) and Nissen (2015) in terms of parameter space but so are the errors in GALAH parameters compared to these studies. We therefore prefer to keep the parameter range wider and still call our sample solar twins.

We further used the quality flags on all the chemical abundances available in the data set, specifically `flag_sp` and `flag_X_fe`, to verify which stars had reliable astrophysical parameters. When the flag is zero, it means that the measurement of the stellar parameters and chemical abundances have no reported problem. We also considered a high-detection rate of the different elements, i.e. stars with several elements measured. Only 14 elements have a detection rate above 90 per cent in solar twins. These elements correspond to Na (100 per cent), Mg (100 per cent), Al (98 per cent), Si (100 per cent), Ca (97 per cent), Sc (100 per cent), Ti (97 per cent), Cr (100 per cent), Mn (100 per cent), Ni (94 per cent), Cu (98 per cent), Zn (99 per cent), Y (94 per cent), and Ba (100 per cent). Finally, we removed stars that had reported chemical abundance ratios with $[\text{X}/\text{Fe}] > 1$ or $[\text{X}/\text{Fe}] < -1$ to avoid outliers. All these cuts gave us 44 317 entries. A further selection of entries with an available 1D spectrum in `DATACENTRAL` with no negative fluxes gives us spectra for 39 157 unique stars for our catalogue. Some of these stars, however, have been observed more than once. We use this sample of repeat observations to assess the uncertainties in our results (see Section 3.2 for details).

2.3 Determining abundances with The Cannon

We use the method *The Cannon* to determine the new abundances. This is a data-driven approach that allows us to derive stellar labels (in our case, stellar parameters and chemical abundances) from stellar spectra. In short, the code connects the flux of a spectrum at different wavelengths with a set of labels, by constructing a polynomial model with linear, quadratic, and cross-term coefficients for the labels. It was introduced by Ness et al. (2015) and has since then been widely applied to stellar spectra (Casey et al. 2017; Buder et al. 2018; Nandakumar et al. 2022). One of its most practical strengths is that it does not use a physical but an empirical model of the spectra. This allows the method to obtain labels at comparable high precision compared with a standard derivation of abundances from e.g. synthetic spectral fitting, and it does it fast and computationally cheap.

The Cannon requires the existence of a subset of reference objects (in other words, a training set) which has well determined stellar

Table 2. Selected line masks for the elemental abundance analysis using *The Cannon*. The full table is available online as supplementary material.

Element	Wavelength (Å)	Line mask (Å)
Na	4751.8218	[4751.7968, 4751.8468]
Na	5682.6333	[5682.6083, 5682.6583]
Na	5688.2050	[5688.1800, 5688.2300]
Mg	4730.0286	[4730.0036, 4730.0536]
Mg	5711.0880	[5711.0630, 5711.1130]
Mg	7691.5500	[7691.5250, 7691.5750]
Mg	7722.5930	[7722.5680, 7722.6180]
Mg	7759.2980	[7759.2730, 7759.3230]
Al	6696.0230	[6695.9980, 6696.0480]
Al	6698.6730	[6698.6480, 6698.6980]
Al	7835.3090	[7835.2840, 7835.3340]
Al	7836.1340	[7836.1090, 7836.1590]
–	–	–
Y	4819.6383	[4819.6133, 4819.6633]
Y	4854.8611	[4854.8361, 4854.8861]
Y	4883.6821	[4883.6571, 4883.7071]
Y	5662.9241	[5662.8991, 5662.9491]
Y	5728.8865	[5728.8615, 5728.9115]
Ba	5853.6680	[5853.6430, 5853.6930]
Ba	6496.8970	[6496.8720, 6496.9220]

labels and must cover the parameter space sufficiently well and evenly. In a data set like GALAH, it was found that there are too few metal-poor stars with well-determined labels to characterize well enough the metal-poor population observed by this survey (Buder et al. 2018). In our case, by focusing on the solar twins only, we have higher confidence to build a valuable training set for *The Cannon*.

For this work, we used the open source version of *The Cannon* provided by Casey et al. (2017)³, and built a quadratic model for the solar twins of GALAH DR3 using different masks over the wavelength range for each of the considered labels (see below).

2.3.1 Masks

Buder et al. (2018) discussed that *The Cannon* had a better performance while using masks in the spectra for each label than when using the entire spectrum without filtering specific wavelength regions. This means our *The Cannon* models are not performed in every pixel for every label but only on pixels that were selected to contain information about the label. That information is known from synthetic spectra (Buder et al. 2018). Later on, these masks have been considered by SME to perform the fitting of observations with synthetic spectra to provide the abundances of all GALAH stars in GALAH DR3. To make models for the different *The Cannon* labels predictions, we have considered lines for the selected elements taken from Buder et al. (2018) and the masks listed in Table 2. No masking was considered for the stellar parameters.

2.3.2 Training set

We are interested in training our model with a set of high quality spectra with accurate label measurements. That model has to be used to generate new labels and uncertainties for the 39 157 solar twin stars. To select the best quality data, we base our criterion on the SNR of the spectra, because GALAH DR3 provides its most accurate and precise parameters and abundances for high-SNR spectra (Buder

²https://www.galah-survey.org/dr3/the_catalogues/

³<https://github.com/andycasey/AnnieLasso/tree/master/thecannon>

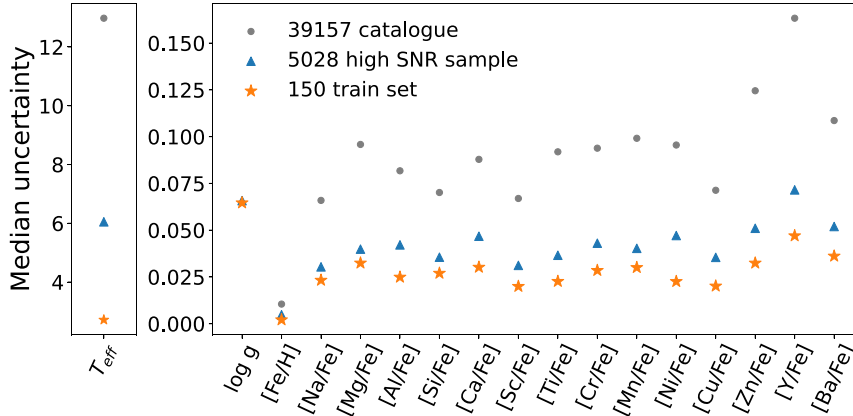


Figure 1. Median internal uncertainties per label as reported in GALAH DR3. In dots, we show all 39 157 solar twins from our selected catalogue. In triangles, we show the 5028 high SNR sample of solar twins. The stars represent the selection of 150 stars used as final training set (see discussions in Section A4). Uncertainties are smaller for the high SNR samples which are used for training.

et al. 2021). We thus consider a set of spectra with SNR above 50 across all CCDs because these are the highest quality spectra obtained in GALAH DR3 and thus provide the most precise results. Since we want to make new predictions on the labels of the whole solar twin data set, we consider the full data set of 39 157 spectra as a test set. There are 5132 high SNR spectra (hereafter, called high SNR sample) for training and 39 157 for test/prediction. These are different spectra of different stars. Among these 5132 high SNR spectra, there are still some stars which have problems of normalization or data reduction. These appeared as outliers in the $[X/Fe]$ versus $[Fe/H]$ diagrams, which we identified after visually inspecting the spectra of the outliers. We removed them from our sample reducing our training set to 5028 stars.

Fig. 1 shows the median uncertainties in all labels for different selections of our sample. In dots, we show the median internal uncertainties of the entire solar twin catalogue. In triangles, we plot the median uncertainties for all stars in the high SNR sample, and in stars, we show the median uncertainties for a selection of a training set of 150 stars (Section A4). We can see how our training set has labels that are more precise than the rest of the catalogue for all labels except for surface gravity, which was predominantly estimated from non-spectroscopic features. Therefore, its precision is not dominated by SNR.

In Fig. 2, we show the individual abundances as a function of the metallicity, defined as $[Fe/H]$, for the 39 157 solar twins as well as the 5028 high SNR solar twins subset and the 150 high SNR solar twin selected as training set (see Section A4) following the same colour scheme as Fig. 1, with the full sample as contours to observe the bulk of the distribution followed by the training set (in orange). We see that the high SNR solar twin subset does not fully cover the parameter space of the entire catalogue. The low-SNR spectra might induce a spread in the abundances that is driven by the uncertainties. However, we observe the selected training set of 150 stars has a good coverage of the metallicity and abundances when compared with the high-SNR sample.

3 THE CANNON CATALOGUE

In this section, we present our results about the performance of our new catalogue of high precision abundances of solar twins in GALAH. Our catalogue of abundances can be downloaded from VizieR. We comment that besides the cuts on flags of the abundance

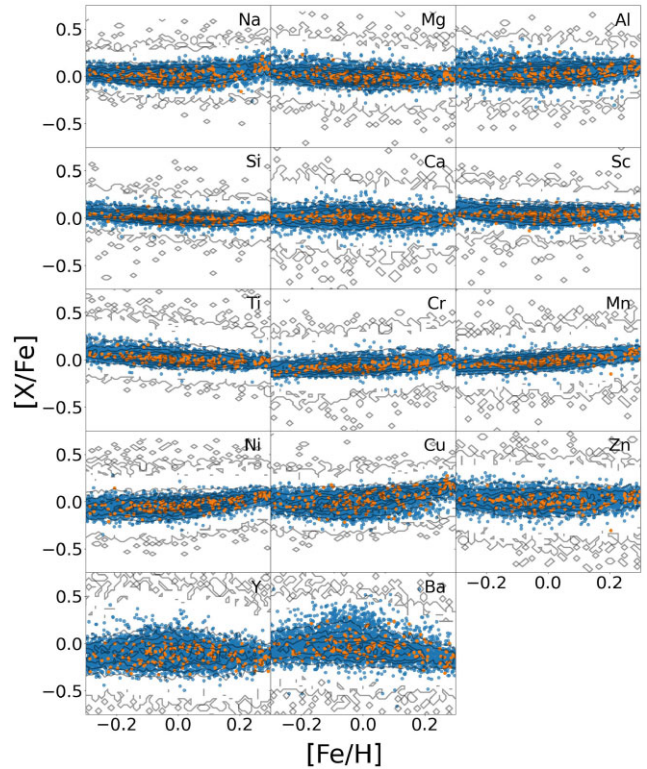


Figure 2. Panels of abundances in the format of $[X/Fe]$ versus $[Fe/H]$. The corresponding abundance X is displayed in each panel. In contours, all 39 157 solar twins data. In blue, all 5 028 high SNR solar twins data. In orange, 150 high-SNR spectra of solar twins training set selection (see Section A4). The entire data set covers a wider range in abundances, but that could be due to higher uncertainties. Both train set and high SNR set have similar range in abundances.

determinations of GALAH-DR3 and visual inspection of clear outliers in the training set, no further quality flag has been applied to this catalogue. Therefore, we recommend users to revise the source spectra in case of any unusual finding regarding abundances of solar twins derived with this method. This catalogue can be used as a complementary catalogue to the new one provided for giant stars using a very similar methodology by Manea et al. (2023).

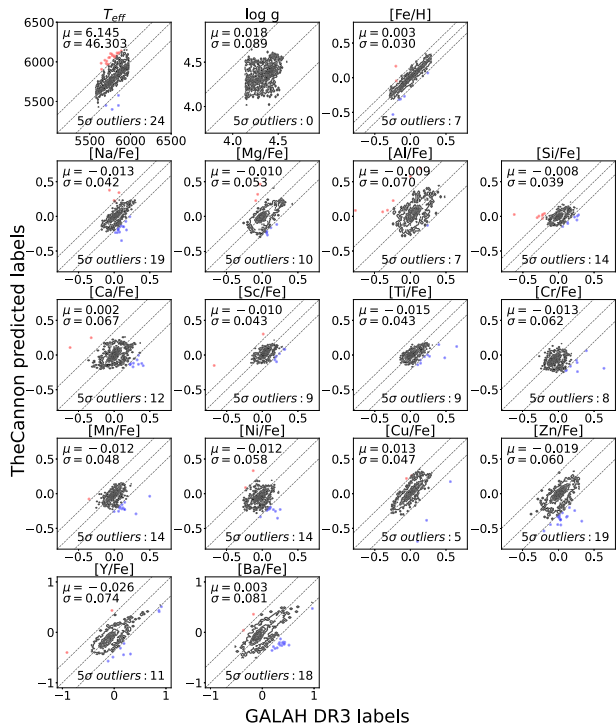


Figure 3. Comparison of labels obtained with *The Cannon* and GALAH DR3 for stars with $\min(\text{SNR}) > 50$. Each panel correspond to a different label. Outer dashed lines correspond to 5σ boundaries. Overestimates above the upper dashed line are plotted in red, underestimates below the lower dashed line are plotted in blue. The mean μ and standard deviation σ of the difference between results are specified in each panel, as well as the number of outliers found outside each boundary.

3.1 Comparison with GALAH DR3

We use the *The Cannon* Model 3, namely the model built with 150 solar twins of minimum SNR of 117 across all CCDs (See Appendix A). In order to validate our results, we look at the agreement of the predicted values over the 39 157 solar twin catalogue at high and low SNR, as well as the uncertainties reported by our *The Cannon* model and compare them to the current GALAH DR3 reported uncertainties.

In Fig. 3, we show the comparisons of our labels and the GALAH DR3 results for stars in the high-SNR sample. For stellar parameters T_{eff} and $[\text{Fe}/\text{H}]$, we have a good agreement within 50 K and 0.03 dex, respectively, with some exceptions of overestimates in the 5σ boundary for T_{eff} . However, for $\log g$, we are not able to consistently recover the labels, obtaining a scatter in the one-to-one relation of 0.1 dex which is very large considering the small range of surface gravities in our sample. We recall that GALAH DR3 does not derive surface gravities from the spectra because these spectra are not sensitive to surface gravity. It is thus not surprising that the agreement is poor. For the chemical abundances, we have a good agreement overall. However, we can notice some outliers for Al and Cr in the sense that we overestimate the abundances for a few metal-poor stars and underestimate the abundances for some other metal-richer stars.

In Fig. 4, we display the same results but for stars with low SNR (below 50). As expected, the number of outliers increases, with stars predicted to have effective temperatures outside the solar twin range. It is expected that the comparison will be worse in this case, as we

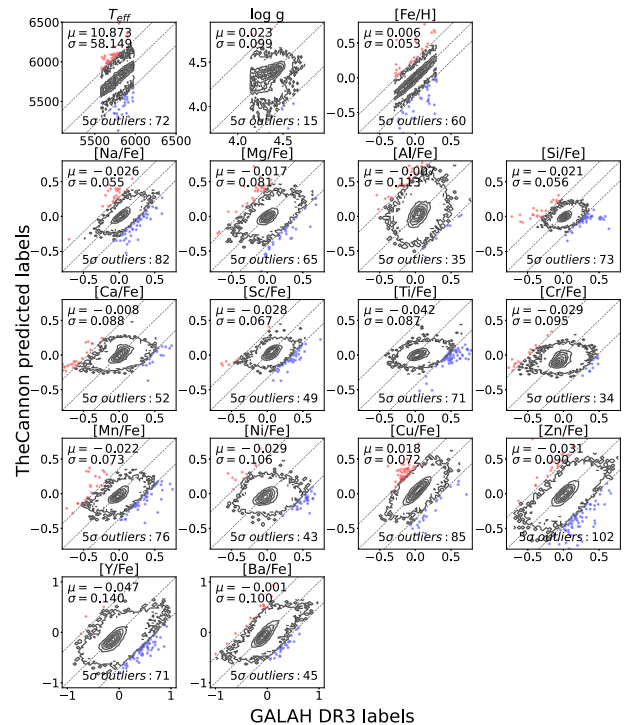


Figure 4. Same as Fig. 3 but for stars with spectra of $\min(\text{SNR}) < 50$.

Table 3. False positive proportion of solar twin stars from the original selection that are excluded from the box limited by equation (1) after adding or removing 1, 2, or 3 σ uncertainty from the new *The Cannon* parameters.

Label	σ	2σ	3σ
Solar twin	0.251117	0.489644	0.706464
T_{eff}	0.056644	0.116863	0.179610
$\log g$	0.187476	0.392216	0.614526
$[\text{Fe}/\text{H}]$	0.021835	0.045305	0.070434

know the measurements obtained by the pipeline of GALAH DR3 are more uncertain for low-SNR spectra.

We still obtain a good agreement for Na, Mg, Si, Sc, Cu, Zn, Y, Ba, although with a higher dispersion than the high-SNR stars. However for Ca, Cr, we observe two trends in our prediction. The model makes near-flat predictions for metal-poor and metal-rich stars, resulting in two groups in the comparisons, with the model underestimating the abundances for these groups. For Al, we observe a higher slope in the one-to-one relation, e.g. we overestimate this abundance. For Ti and Ni, we obtain a flat prediction, namely all the stars in the low-SNR regime have solar-like Ti abundance.

In our final catalogue, we remove all the outliers outside 5σ boundaries found in both high and low SNR. This translates into further removing 88 high-SNR stars and 749 low-SNR stars, obtaining a final catalogue of 38 320 solar twins.

An interesting value to quantify is the false positive proportion of our solar twin selection based on our definition from equation (1). In Table 3, we list the proportion of stars that are excluded from the box limited by equation (1) when varying the stellar parameters by 1, 2, or 3 σ . The first row considers the entire box, but the second, third, and last rows show the impact of each parameter separately. We can see that the largest impact is the surface gravity. This can be

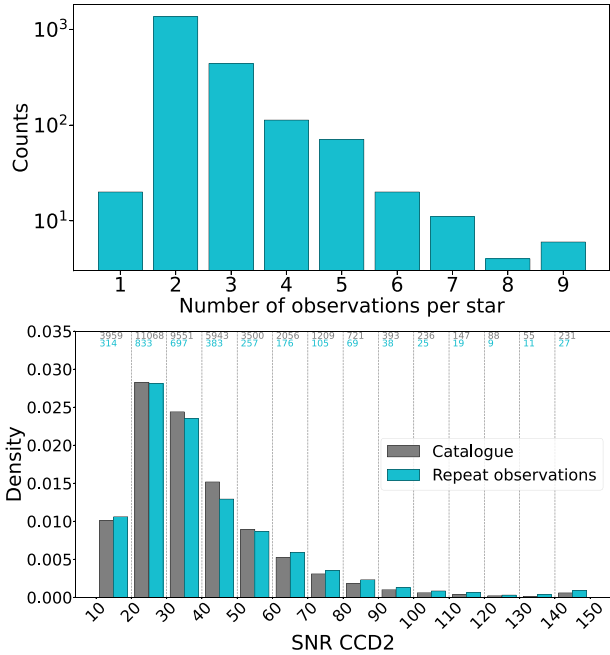


Figure 5. Top panel: Histogram of number of observations for repeated stars in the entire GALAH DR3 sample. Bottom panel: Density histogram as function of SNR in CCD2. In grey (left bins), all 39 157 solar twins catalogue. In cyan (right bins), repeat observations of solar twins. Raw counts above each bin on the top of the figure. Note for the last SNR bin, we collapsed the all catalogue data with SNR above 140.

explained from the agreement between GALAH DR3 and our *The Cannon* $\log g$, which disagree more than the other parameters (see Fig. 4).

3.2 Repeat observations

To evaluate our results, we consider repeat observations for some stars. There are a total of 5145 spectra for 1910 solar twins. The top panel of Fig. 5 shows a histogram with the number of observations for the stars in the repeat observations sample. In general, we have two or three observations for each star, about hundred stars with three and four observations and a few of stars having more than four observations.

For this sample, we compute the uncertainties as follows: for each star, we consider its observation of maximum SNR in CCD2, then for each repeat observation we compute the difference in the prediction compared to the one of maximum SNR. Finally, after doing it for all the stars in the sample, we divide the SNR range in bins of length 10 and we compute for each bin the standard deviation of the difference in predictions of the repeat observations within SNR bin to their respective maximum SNR prediction.

The bottom panel of Fig. 5 shows a normalized density histogram of the repeat observation counts (cyan right bins) as function of the CCD2 SNR. In grey, we show the observation counts for the entire solar twin catalogue. The spectra of repeat observations cover well the SNR distribution of the entire catalogue. The figure shows that our sample is a good representation of the SNR of the entire GALAH DR3 data set.

3.3 Comparison of uncertainties

We use the internal uncertainties obtained by our model and compare them to the ones already given by GALAH DR3. We also use the repeat observations sample described above, where we made predictions with the model and compare them with the uncertainties of repeat observations given by GALAH DR3. In Fig. 6, we summarize our findings, we have both internal uncertainties (in solid lines) and repeat observations uncertainties (in dashed lines) for GALAH DR3/SME in black and our *The Cannon* in red as function of SNR. We included the SME uncertainties for $\log g$ even though the estimation for this label does not rely on spectral fitting like our *The Cannon* model does. For the internal uncertainties, we highlight for each label the median of the reported uncertainty taken over each SNR bin. For the repeat observations uncertainties, we report for each label the standard deviation σ of the absolute differences (of the predicted label with the predicted label of the maximum-SNR observation) taken over each SNR bin in a similar way as it is done by GALAH DR3 (Buder et al. 2021).

In general, the internal uncertainties reported by our model are below the ones reported by GALAH DR3. For GALAH DR3, they tend to increase as SNR decreases. Our model also predicts labels more uncertain at lower SNR, but the difference in uncertainties between high and low SNR is smaller than for GALAH DR3. For the uncertainties obtained from repeat observations however, our results are comparable to GALAH DR3 for all SNR ranges. The uncertainties of repeat observations for the stellar parameters T_{eff} and $[\text{Fe}/\text{H}]$ are similar for both pipelines as function of the SNR, and are generally higher than the internal uncertainties. For the chemical abundances, we find some cases, where our model has higher uncertainties than SME, especially at high SNR. A notable example is Ca, where our *The Cannon* model reports higher repeat observations uncertainties at high SNR. We note that in GALAH DR3 a more detailed approach in masking telluric lines was made (Buder et al. 2021) but here we used a less refined mask which likely contains more telluric features.

For Cu and Zn differences are negligible, which could be due to the fact that in both procedures the masking of the spectra is the same therefore the methods consider the same information from the spectra. The opposite difference is found for Si, Sc, Ti, Cr, Ni, Y where in general our uncertainties of repeat observations are lower than GALAH DR3. In particular, for Si and Mn our masks have more pixels than SME in GALAH DR3 which included an additional filtering due to blending but were not labelled as such in the masks. Our *The Cannon* model does not seem to be affected in terms of precision suggesting that perhaps the further filtering in the GALAH DR3 masks was too strict for solar twins, removing pixels. For Ti, we also observe that *The Cannon* obtains more precise abundances. SME determines Ti separately from neutral and ionized Ti lines, whereas *The Cannon* takes all lines together. It is thus expected to have better precision in our model because more pixels are used.

In general, at lower SNR, *The Cannon* performs better since it is always using information from the whole spectrum, whereas SME applies further filtering in masking the detected lines in each spectrum (Buder et al. 2021). When there is more information in lines available in the spectra, *The Cannon* will use this information, for example for regions where pixels are neglected by SME due to possible blending. *The Cannon* can outperform the fitting of one element at a time with SME, since it can learn the variation of blended features as well.

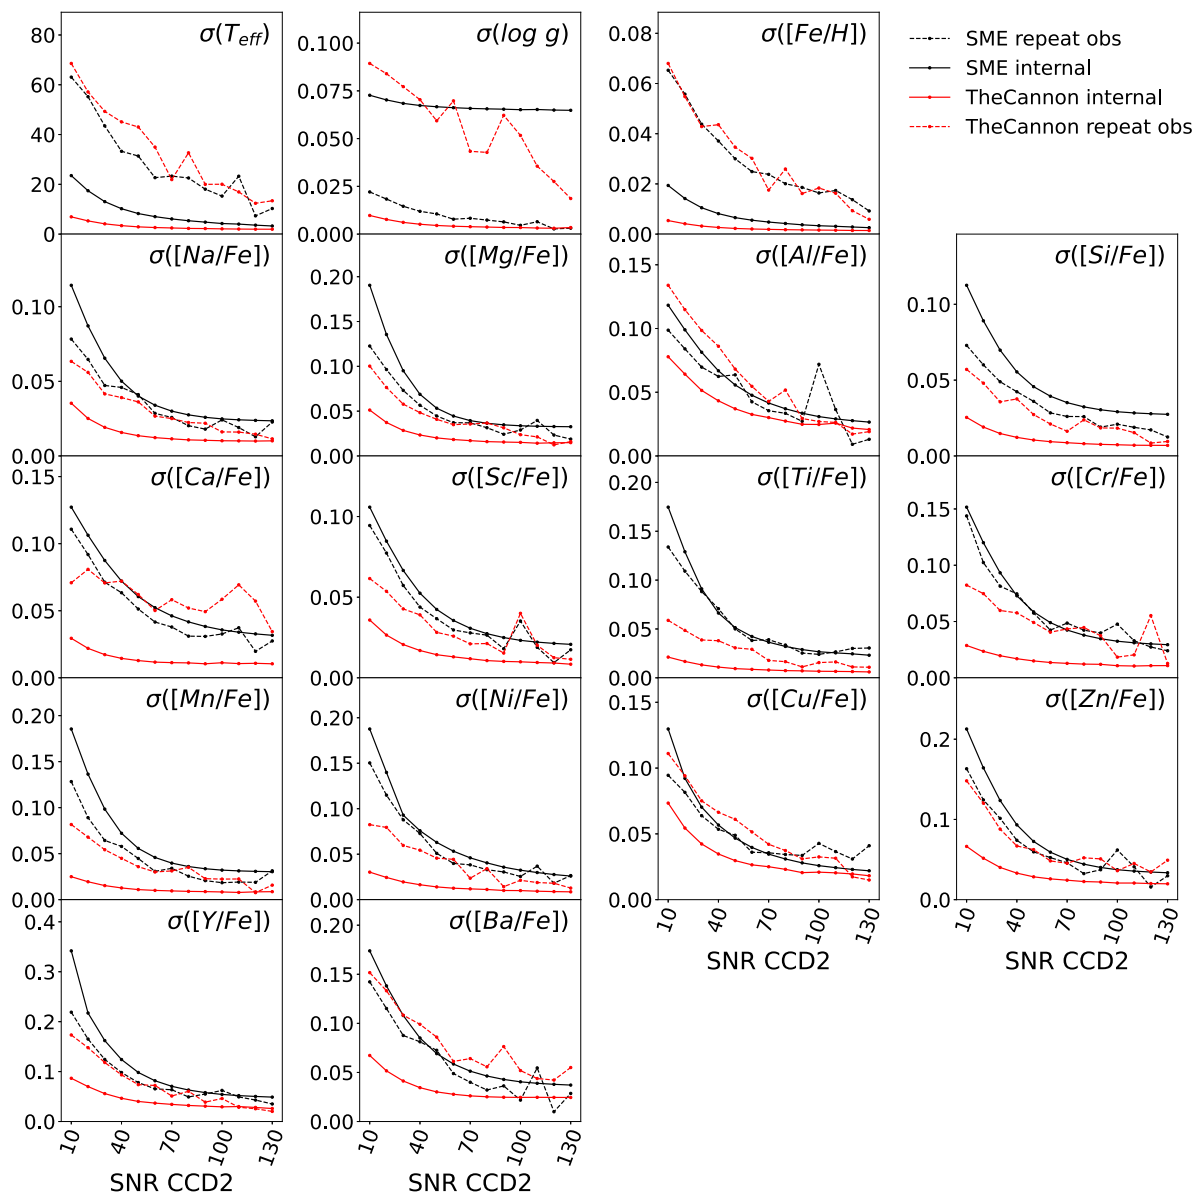


Figure 6. Standard deviation of uncertainty labels as function of SNR CCD2. In black, SME uncertainties. In red, *The Cannon* uncertainties. Solid lines represent the internal uncertainties given by the covariance matrix of the fitting by both SME and *The Cannon*. Dashed lines represent the uncertainties of repeat observations.

Furthermore, *The Cannon* is able to learn the shapes of imperfect continuum fluxes, as is the case for low-SNR spectra. If we are using the same amount of pixels as SME, the repeat observations uncertainties are comparable since they are fitting essentially the same features.

3.4 Abundance distributions

Fig. 7 shows the individual abundances as a function of metallicity for *The Cannon* values coloured as density plots and for GALAH DR3 as contours. Salmon and dark red stars displayed correspond to the lowest/highest eccentricity stars in the catalogue to be further analysed using phylogenetic trees (see Section 4). We can see how the precision in terms of the dispersion improves significantly, with exceptions for Al, Cu, Zn, Y, and Ba. Al, Cu, and Zn have very weak

lines which makes them a very difficult element to measure even for *The Cannon*. The internal uncertainties are comparable for these elements as discussed in the previous section. Y and Ba are elements that are expected to present a large scatter. Stars in binary systems that had an AGB companion might have been polluted by s-process elements such as Y and Ba that were produced by the AGB star (Escorza et al. 2019). Based on *Gaia* DR3, however, we do not find a clear signature of binarity for stars with higher [Ba/Fe] and [Y/Fe] abundance ratios in their RUWE⁴ parameter or the uncertainties in

⁴RUWE stands for Renormalized Unit Weight Error and is determined in *Gaia* to account for uncertainties in the astrometric solution. It is considered that a value above 1.4 could indicate that the uncertainty in the solution is higher due to the non-single nature of the star, which would change its position in the sky.

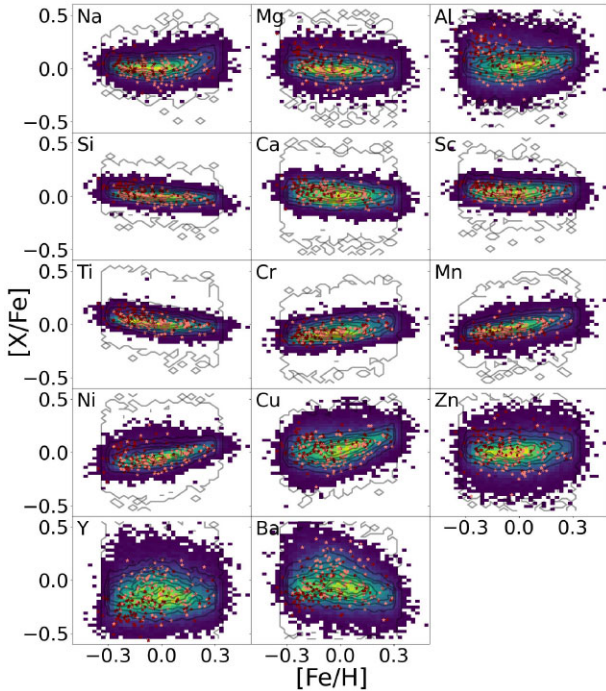


Figure 7. Individual abundances as a function of metallicity for the solar twins analysed in this work. Coloured density plots correspond to our labels as determined using *The Cannon* while contours delineate the distribution of GALAH DR3. Stars selected for further analysis in Section 4 are displayed as salmon and dark red stars for low and high eccentricity stars, respectively.

the radial velocities. If such high Y and Ba stars were in binary systems, their separations would be large and their periods long.

Elements such as Mg, Si, Ca, and Ti are formed predominantly in SN II progenitors (see Kobayashi, Karakas & Lugaro 2020, and references therein). In this sample, we do not find stars that might have formed from an α -enhanced gas such as the thick disc, deducing that this sample is composed predominantly by thin disc stars. Iron-peak elements such as Sc, Cr, Mn, and Ni follow the trends observed in other higher resolution and higher precision studies, such as Adibekyan et al. (2012), Bensby, Feltzing & Oey (2014), and Battistini & Bensby (2015).

Fig. 8 illustrates our values as a function of age. Here, the age corresponds to the values reported by GALAH DR3. We note these ages are not fully consistent with our new parameters since they are estimated with the GALAH DR3 parameters. A new derivation of ages is beyond the scope of this paper. Here, we aim to provide an illustration of the type of studies that could be performed with the entire usage of our catalogue.

As above, the coloured density plots represent our values determined with *The Cannon*. The black line corresponds to the linear regression fits of the abundance-age trend determined by Spina et al. (2016) and Bedell et al. (2018) who performed a high precision spectroscopic analysis of solar twins using high resolution. The trends found using our data are displayed as red lines. These trends were found with linear regression fits considering the uncertainties in both axes (age and $[X/Fe]$), using a reduced representation of the catalogue. We divided the age (x -axis) range in 100 bins of equal width, then for each bin, if the subsample size is greater than 30 (for statistical significance), we take the weighted average over the subsample lying within the age bin, and consider the uncertainty as the median uncertainty over it. The linear fit coefficients for the

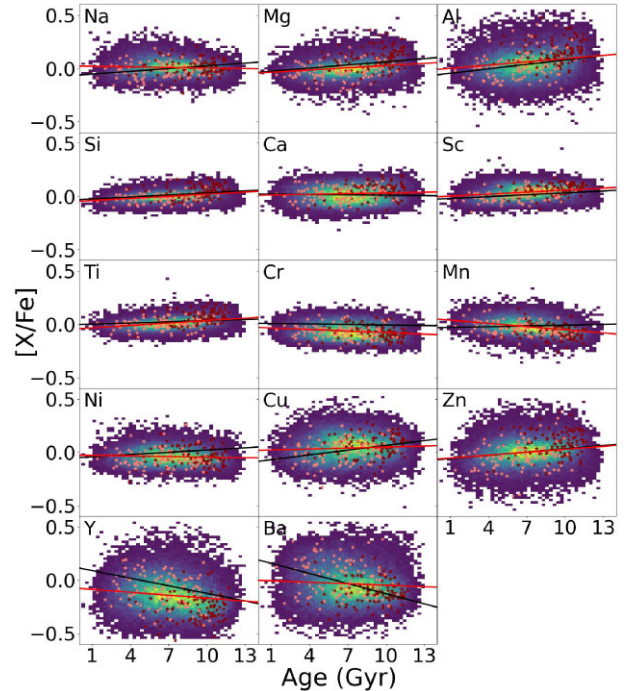


Figure 8. Individual abundances as function of stellar age for the solar twins analysed in this work. Coloured density plots correspond to our labels as determined using *The Cannon*. Age estimates are taken from GALAH DR3 (Buder et al. 2021). Black solid lines correspond to the linear regression fits of abundance-age trends found by Spina et al. (2016) and Bedell et al. (2018). Red solid lines correspond to the linear regression fits found with our data. Stars selected for further analysis in Section 4 are displayed as salmon and dark red stars for low- and high-eccentricity stars, respectively.

Table 4. Linear coefficients (with uncertainties) of abundance-age trends found for the elements used in our study. Abundances of GALAH DR3 solar twins come from our *The Cannon* derived abundances, ages come from GALAH-DR3 (Buder et al. 2021) estimates.

Element	a	σ_a	b	σ_b
Na	-0.001 77	0.000 69	0.022 19	0.005 36
Mg	0.006 99	0.001 04	-0.042 65	0.007 95
Al	0.010 19	0.001 92	-0.009 52	0.013 76
Si	0.006 65	0.000 57	-0.045 71	0.004 29
Ca	0.002 23	0.000 63	0.010 05	0.004 85
Sc	0.006 27	0.000 76	-0.004 39	0.005 63
Ti	0.007 59	0.000 55	-0.041 19	0.004 2
Cr	-0.004 87	0.000 68	-0.029 14	0.005 25
Mn	-0.010 01	0.000 65	0.048 98	0.004 86
Ni	-0.001 82	0.000 67	-0.027 78	0.004 39
Cu	0.003 18	0.001 48	0.016 62	0.010 33
Zn	0.009 19	0.001 47	-0.064 43	0.010 88
Y	-0.009 23	0.001 92	-0.077 16	0.014 9
Ba	-0.004 59	0.001 42	-0.001 57	0.011 22

trends of each element in our catalogue along with their uncertainties are reported in Table 4.

We qualitatively obtain consistent trends for Mg, Al, Si, Ca, Sc, Ti, and Zn. There are minor offsets, such as Cr, but that was already discussed in Buder et al. (2021). Since we train with GALAH DR3 values, it is expected that the offset remains here. We note that the trends found for Na, Mn, Ni, Cu, Y, Ba are in conflict with the ones found by Spina et al. (2016) and Bedell et al. (2018) for high

resolution solar twin stars, where we observe that our data follows slopes of different signs. The differences in these trends can be caused by several factors, such as different levels of precision and accuracy of these abundances between our results and Spina et al. (2016) and Bedell et al. (2018). Other sources of differences might be caused by the fact that our sample is much larger than the ones used by Spina et al. (2016) and Bedell et al. (2018), implying that we might have more stars coming from distant regions in the Galaxy, where the chemical enrichment might occur at different rates (Ratcliffe et al. 2024).

With respect to Fig. 8, we also want to stress that the underlying training set of GALAH DR3 is subject to significant selection effects and systematic parameter inaccuracies. This includes the overestimation (or systematic clumping) of stellar ages of stars around 2 Gyr, due to the missing separation of young star isochrones in our solar twin parameter space. We also expect to sample more intermediate age thin disc stars from the underlying set of stars due to their relative abundance within the GALAH selection function (neglecting the Galactic plane and sampling within magnitude ranges).

4 PHYLOGENETIC TREES WITH GALAH SOLAR TWINS

In this section, we use our new catalogue in the construction of phylogenetic trees. For our purpose, we select from the catalogue two groups of solar twins with different orbit eccentricity. More specifically, we select the 100 stars with lowest eccentricities and 101 stars with highest eccentricities in the sample. These values come as valued-added information in GALAH and are derived from *Gaia* DR3 data (see Section 2 for details). The eccentricities range from 0.000 723 and 0.000 733 for the low-eccentricity group, and from 0.39 to 0.63 for the high eccentricity group. Our goal is to study what the phylogenetic trees built using our measurements tell us about the relatedness of our sample, and how our measurements help in this goal compared to the GALAH DR3 ones. For this experiment, we compare the phylogenies constructed from both data sets. The stars selected for the analysis are listed in Table B1. In that table, we are including the stars ID as labelled in the tips of our trees, in addition to the *Gaia* DR3 IDs for further references, and their ages and eccentricities as downloaded from the valued added catalogues of GALAH DR3. This sample has a median SNR of 32, and the typical uncertainties for the elements in the sample range from 0.094 dex using GALAH DR3 data to 0.026 dex using our *The Cannon* measurements. Their chemical abundance distributions can be seen in Fig. 7 and 8 with yellow and red symbols.

A natural question arises here about our choice of stars. Indeed, there is no particular reason to choose this sample above any other one. Applying the NJ algorithm to the full complement of 40 000 stars is computationally impractical, especially when assessing the robustness of the built trees (see discussions below), due to its time complexity of $O(n^3)$ (Yang 2014), where n is the number of tips. Visualizing a tree with so many tips is also impractical. Therefore, given that we are still at an exploratory phase of using phylogenetics for Galactic evolution, we prefer to make choices of smaller sets of data, and focus our scientific aim to a particular question. In the following sections, we aim to compare the phylogenetic signal between our *The Cannon* and the standard GALAH-DR3 abundances, for which a reduced number stars is sufficient.

4.1 Building trees with the neighbour-joining algorithm

Trees are built following the methods used in Jackson et al. (2021) and Jofré et al. (2017) by using the classical method called neighbour-joining (NJ), a computationally fast divisive cluster algorithm proposed by Saitou & Nei (1987) that iteratively joins nodes (in our case, stars) closely related by a given pairwise distance matrix. This matrix of distances has size $N \times N$ with N being the number of stars considered in the analysis, and the distance is calculated as a *Manhattan Distance* between the chemical abundances of each star in the pair. The tree reconstruction is made in a greedy way by iteratively joining a pair of nodes from the distance matrix that minimizes the *Q criterion*. For a tree with r nodes, the pair (i, j) is joint by minimizing

$$Q_{ij} = (r - 2)d_{ij} - \sum_{k=1}^r (d_{ik} + d_{jk}), \text{ for } i < j \leq r, \quad (2)$$

where d is the distance between the pairs (See section 3.3.3 of Yang 2014, for more details of this terminology). Then a new node is created and is called an internal node, which acts as predecessor of the two joint nodes. The procedure continues by recomputing the distances of the remaining nodes to the new internal node, removing the rows and columns in the distance matrix related to the two joined nodes and adding the ones related to the new one. This reduces the dimension of the distance matrix by 1. The process iterates again until the distance matrix is of size 2×2 , joining the two remaining nodes and returning the built tree.

We build our distance matrix using our selected sample of stars from both eccentricity groups (see Table B1). To compute the distance, we consider a vector of measured abundances for each star, and use the *Manhattan Distance* which is the absolute difference of two vectors. This means that we are using chemical distances for the stars in our sample in a similar way to Jofré et al. (2017). In our case, we perform a further selection of the abundances, namely those whose age-abundance trend is monotonic in Fig. 8. This is known to increase the additivity of the distance matrix, hence making NJ trees that are closer to the true phylogenetic tree (Retzlaff & Stadler 2018). We thus exclude Ca, Cr, and Ni since their age-abundance trends are flat, thus not evolving in time. Since these abundance ratios do not change in our sample, they add noise in our tree reconstruction (Yang 2014; Jackson et al. 2021).

To account for the uncertainties in the data, we build NJ trees from distance matrices computed by empirically sampling a random value out of a normal distribution for each of the abundances, centred at the reported measurement and with a standard deviation of its reported uncertainty. We build 2000 trees by sampling the abundances according to their uncertainties and study their distribution.

From these trees, we select the best tree to be the one which has the highest node support. To do so, we follow the process used in biology that searches the *Maximum Clade Credibility* (MCC) tree out of a sample of trees. Clades correspond to a group of nodes that includes all the descendants of a common predecessor node in the tree. We note that with trees being built empirically, we cannot immediately assume that the trees follow an ancestor–descendant hierarchy; hence, an internal node can not be immediately associated with a predecessor. Using simulated data is needed to learn the prospects and limitations of interpreting clades and evolutionary histories from empirical trees (de Brito Silva et al. 2023). In the MCC, each clade (or node with all the descendant nodes in our trees) in each sampled tree is given a score that reflects the fraction of times that the same pattern appears in all the sampled trees. If the clade occurs for all the trees then the support value is 1 (100 per cent). This indicates high consistency in

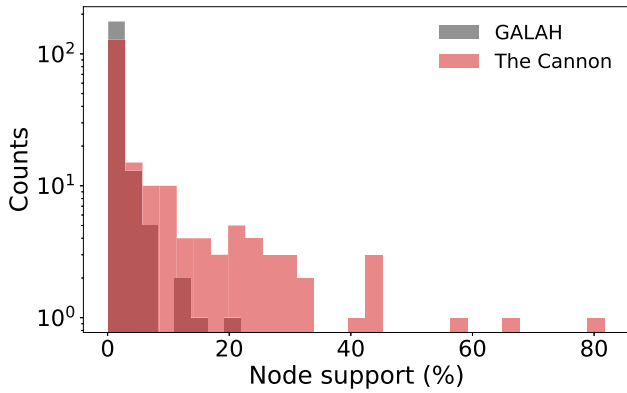


Figure 9. Node support percentage for MCC trees of Fig. 10. In grey, support percentages for the MCC tree built using GALAH DR3 data. In red, support percentages for the MCC tree built using *The Cannon* data.

the data for that topological relationship. The product of these scores is defined as the tree score, so that the MCC tree is the one with the highest tree score. Here, we employ this method to select our best tree and evaluate its robustness, despite not being able to ensure that our nodes and branching pattern can be directly related to clades (see more discussions below).

The distribution of support values for the nodes is shown in Fig. 9. In grey and red, we have the support values for the MCC tree built using GALAH and *The Cannon*, respectively. By comparing the distributions, we observe that *The Cannon* MCC tree is overall better supported. The GALAH highest support value in the GALAH MCC tree is 20 per cent, meaning that every clade seen in the MCC tree only appears in at most 20 per cent of the remaining sampled trees. However, even though the *The Cannon* MCC tree has more support, the overall values do not exceed 50 per cent. This means that even at the high precision in the abundances of our new catalogue, the trees are overall poorly supported.

The root of the tree is the basal split that separates the most distant (in an evolutionary way) object from all the rest. The NJ algorithm produces unrooted trees because in this tree reconstruction method there is no evolutionary model considered and hence no way predict the ancestral state in the relationships of our stars. Hence, even though we are able to apply the MCC method to find the most supported tree, we are not able to attribute a clade in our trees as a group of nodes that includes all the descendants of a predecessor node in an evolutionary context. In our case, it is more appropriate to refer to possible groups in the trees as clans instead of clades.

4.2 GALAH versus Cannon trees

In Fig. 10, we show the MCC trees built from our sample stars. For better visualization of our trees, we choose the tip corresponding to the star labelled with ID 0 as our reference star. This means, the tree is displayed in a way that all branch lengths are visualized with respect to ID 0. That star is the one with highest eccentricity in our sample. Its *Gaia*DR3 ID is 5396076243592498944, and it has an eccentricity of 0.63. This allows us to study the relationship of all stars with respect to that high eccentric one.

We stress that the tree is not rooted. We could have chosen any star as our reference star for visualizing the relationships, e.g. the oldest one, the most metal-poor one, or simply a random one. The tree is displayed relative to that star for better comparison between trees, and the branching order that follows the node with that reference

star does not necessarily mean an ancestor–descendant relationship because we do not have a root nor any prior or model about the evolutionary history of our selected stars. In this case, we chose the most eccentric star of our sample simply because its possibilities to come from a region beyond the Solar neighbourhood is higher than a star with circular orbit, implying its chances to be more evolutionarily distant to the rest are higher. Furthermore, the dynamical properties of the stars depend on *Gaia* data only, and are independent of the spectroscopic parameters (as well as ages and chemical abundances). This allows us to use the same star as reference for both data sets that we consider in building the tree (our *The Cannon* abundances and GALAH–DR3).

The left panel of Fig. 10 shows the MCC tree obtained using GALAH abundances and the right panel shows the MCC tree obtained using *The Cannon* abundances. Both trees have the branches coloured by the metallicity as obtained from the corresponding catalogue. The parenthesis in each tip of the trees corresponds to the eccentricity group, where 0 represents circular orbits and 1 represents more eccentric orbits. Specific information about eccentricities and ages of our stars can be found in Table B1.

By comparing the topologies of these trees, we see some similarities. In both cases, the eccentric stars are located close to the star ID 0, and after a few splits, we are able to observe two main branches. Within these branches, we select clans for further studies, which we label Clan A and B for the GALAH tree, and Clan C and D for the *The Cannon* tree. We will discuss these clans with more detail later on.

We also observe that the length of the tip branches differ between the trees. In fact, the tip branch lengths of the GALAH tree are larger than the ones with *The Cannon*. Some GALAH branches reach 4 dex while Cannon ones reach 3 dex. In fact, the GALAH tree has tip branch lengths which are larger than the inner branches. This is an indication of a *hard* tree (Yang 2014), which are trees prone to errors. In the GALAH tree, the difference between two stars, reflected by the sum of horizontal branch lengths connecting the two tips along the tree, is dominated by the tip branch lengths rather than the internal branch lengths, dominating over the hierarchical structure in the tree. This means that more of the chemical differences among the stars is being explained by their tips than any one of the internal branches. This is also noticed in *The Cannon* tree but to a lesser extent. We have to keep in mind that our catalogue contains very similar stars (all solar twins). The difference between them is intrinsically very small therefore the difference dominated by the tip lengths reflect that the overall distance between stars is more or less similar over the entire sample.

In any case, we observe that both trees chemically cluster together stars from both eccentricity groups, but with the *The Cannon* data the grouping is more resolved. In the GALAH tree, Clan A (highlighted with blue in Fig. 10) contains mostly the stars in circular orbits which are primarily metal-rich, while Clan B (in orange in the figure) contains a mix of stars. When using our catalogue, we see that Clan C (enclosed in red in the right hand panel of Fig. 10) contains mostly low eccentric metal-rich stars. Clan D has (in violet in the figure) stars with high eccentricities and rather metal-poor.

The distributions of the ages of both groups and their associations to the different clans can be seen in Fig. 11. The upper panel shows the distributions of SNR and the lower panel the distribution in age. We see no systematic difference in the SNR of our stars therefore we expect no systematic differences in abundance uncertainties hence branch lengths. The distribution of SNR is well represented to the entire sample (see Fig. 5). Regarding the age distributions, we see that Clan A and Clan C have a range of ages while Clan B and Clan

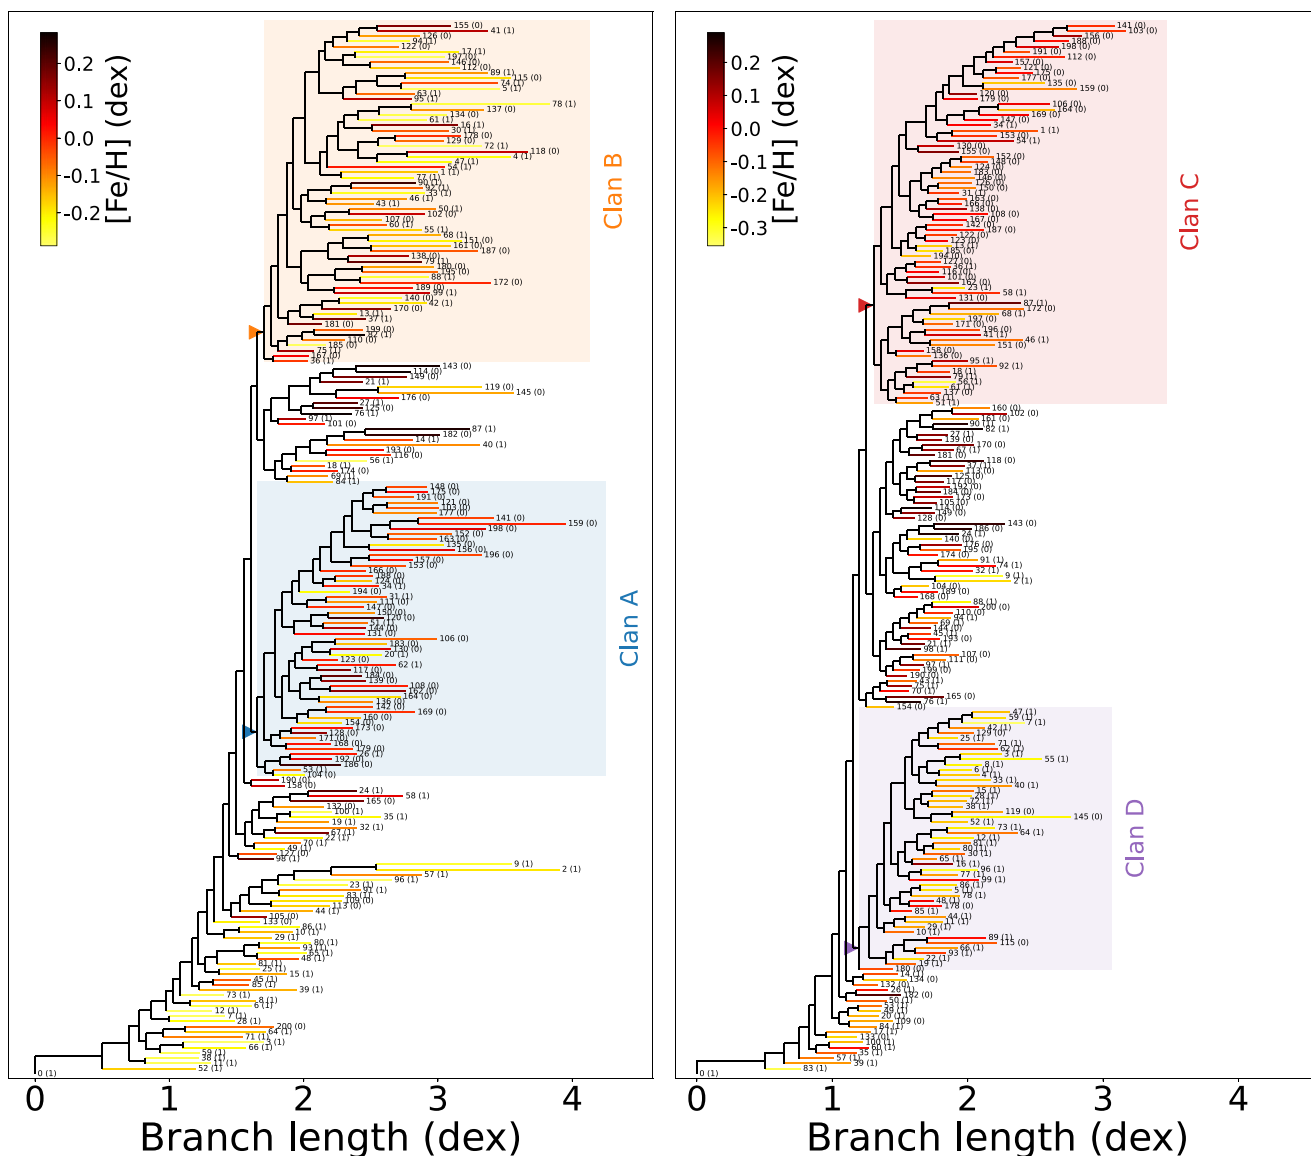


Figure 10. Maximum Clade Credibility (MCC) trees from a sample of 2000 trees for 201 stars of different eccentricity groups with tips coloured by metallicity. On the left, MCC tree obtained with GALAH labels. On the right, MCC tree obtained with *The Cannon* labels. Stars in selected eccentricity groups are enumerated in the tips with values between 0 and 200, 0 being the fixed branch placed at the most eccentric solar twin in the catalogue (eccentricity of 0.63), 1–100 being the following most eccentric stars and 101–200 referring to the less eccentric stars in the catalogue (see Table B1). Selected clans A–B and C–D can be seen from the coloured areas of GALAH and *The Cannon*, respectively.

D have mostly old stars. This is consistent with Clan A and Clan C having mostly stars from the circular orbit group and Clan B and Clan D stars from the eccentric orbit group.

It is worth to comment on the selection of our clans. Between Clan C and Clan D there is a branch of stars that is more similar to Clan C than Clan D but the internal branches are short compared to the length of the tip branches, and the topology is overall more balanced. That branching pattern is indeed similar to a random tree, lacking phylogenetic signal (de Brito Silva et al. 2023). We have thus focused on the clans that are the farthest from each other in the tree.

4.3 Astrophysical interpretation of the selected clans

We now look into the Clan A and B from the GALAH DR3 tree and Clan C and D from the *The Cannon* tree. To do so, we explore the trend of age and $[Y/Mg]$, commonly referred to as *chemical-clock*. Indeed, in the analysis performed on solar twins by Nissen (2015), a tight relationship between age and $[Y/Mg]$ was found. This trend was explained with the argument that yttrium, which is an element produced by AGB stars, increases with increasing Fe, while Mg, which is an element produced by SN II, decreases with increasing Fe. Since Fe increases with time, this difference in dependency with Fe causes a strong dependency of $[Y/Mg]$ with age.

After that study, several works have studied the applicability of this trend considering different kinds of stars, finding that solar-metallicity giants in the solar neighbourhood behave similarly to the

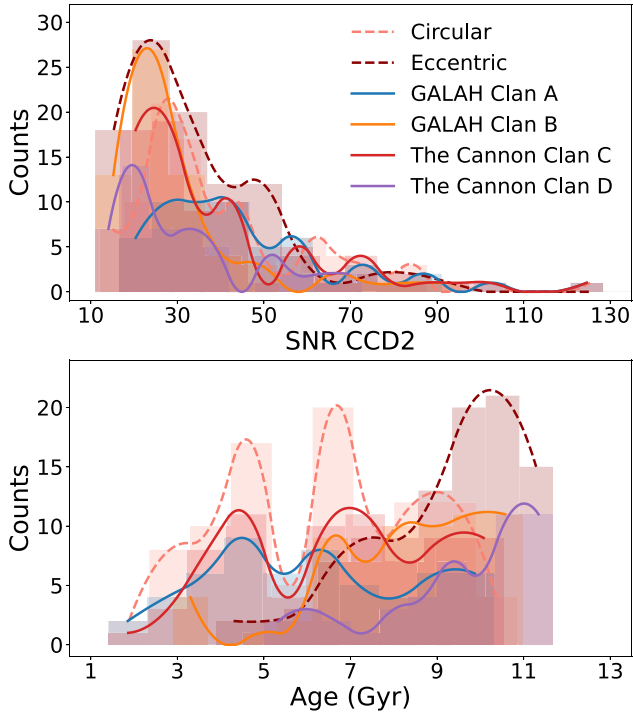


Figure 11. Distributions for SNR CCD2 (top panel) and age (bottom panel), for eccentricity groups in dashed lines and clans selected from trees of Fig. 10 in solid lines.

solar twins (Slumstrup et al. 2017; Casamiquela et al. 2021a) but at lower metallicities, this relationship might weaken (Delgado Mena et al. 2019; Casali et al. 2020; Vitali et al. 2024). It is also suspected that this relation is subject to systematics in the age determination (Berger et al. 2022). Further studies have found this relation might not hold for stars outside the solar neighbourhood (Casamiquela et al. 2021a), which can be explained by the fact that this relation has a strong dependency of the the star formation rate, which is different at different birth radii (Ratcliffe et al. 2024).

Considering that $[Y/Mg] = [Y/Fe] - [Mg/Fe]$, the errors in the abundance ratio used in this work are computed as the quadratic sum of the errors reported for both abundances. Age estimates as well as age errors are taken from Buder et al. (2021). In Fig. 12, we show the age– $[Y/Mg]$ trends for the selected clans in Fig. 10 following the same colours. We compute a linear regression fit of the stars in each group, and plot the fit with a line of the same colour as the corresponding clan. The legend indicates the value of the slope and the uncertainty of the fit. The dashed black line corresponds to the slope of the linear fit found by Nissen et al. (2020), for reference.

We see that the age– $[Y/Mg]$ trends have different slopes for the different clans. However, among all the trends, it is remarkable the agreement of the slope in the trend found for Clan C and the Nissen et al. (2020) fit. Clan C is the group which is composed by a majority of low-eccentric stars using *The Cannon* abundances. We note that the stars analysed by Nissen et al. (2020) have eccentricities normally below 0.1 (see also Nissen 2015; Jofré et al. 2017; Jackson et al. 2021).

Interpreting this finding in terms of the phylogenetic nature of this group is however tricky. As shown recently by Ratcliffe et al. (2024), the age– $[Y/Mg]$ relationship found by Nissen et al. (2020, and references therein) can be interpreted only by considering that birth radii also plays a fundamental role. It is only possible to explain

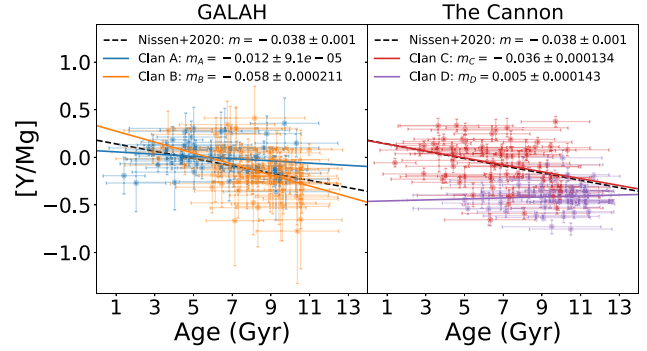


Figure 12. Age– $[Y/Mg]$ relations of the stars in the clans selected from the trees of Fig. 10. Linear regression fits are drawn to quantify the slopes in these relations. Trend found by Nissen et al. (2020) displayed as a dashed black line. $[Y/Mg]$ errors taken from the internal uncertainties reported by GALAH DR3 and our *The Cannon* catalogue. Age errors taken from GALAH DR3 Buder et al. (2021).

that a sample of stars with a restricted metallicity range in the solar neighbourhood can have a range in ages if they come from different Galactic radii. Like this, each star traces a different star formation rate and reaches the same $[Fe/H]$ at different time-scales. The fact that Clan C is composed of stars that are mostly on circular orbits, but have a range in ages, suggests that the oldest stars might have migrated from inner regions. We would expect that old stars with circular orbits that have not migrated should be significantly more metal-poor. These are not here because we have only selected solar-metallicity stars for our study.

Our selection in metallicity is however not too restrictive. In fact, we have a range of 0.6 dex in metallicity in the sample, and Clan C contains stars of all metallicities, indicating that some ISM evolution at the solar radius must be present. Disentangling which stars are product of the inheritance of the ISM at the solar radius and which have migrated or are visiting due to dynamical heating (higher eccentricities) is tricky since all these processes are mixed in the disc (Feltzing, Bowers & Agertz 2020) and evolve as time passes (Aumer, Binney & Schönrich 2016; Bird et al. 2021; Lu et al. 2022).

The slope of Clan D is rather flat, but that could be due to the fact that the stars in Clan D are predominantly eccentric. These stars might have originated from different Galactic radii and are less exposed to have a shared history. Most of them are also old, making the resulting fit biased. It is currently believed that stars originating from different Galactic radii might have different age– $[Y/Mg]$ relations (Casali et al. 2020; Casamiquela et al. 2021a). Furthermore, the trend and its relationship with birth radius evolves with time. For oldest stars, this ratio could have been flat across the Galaxy (Ratcliffe et al. 2024), and this is consistent with our findings.

For the GALAH Clans A and B, we see a smaller difference when considering the uncertainties. Moreover, Clan A has an age– $[Y/Mg]$ trend which is flatter than Clan B, which is the opposite to what we find with the *The Cannon* abundances. It is thus hard to explain that Clan A, which contains predominantly low-eccentric stars, deviates more from the Nissen et al. (2020) relation than Clan B, which contains a mix of stars. This might be a consequence of higher uncertainties in the abundances of GALAH DR3.

A natural question may arise as to whether we are able or not to recover the same trends by simply doing dynamical cuts for the sample stars. To answer this question, we consider Clans A and C, because they are the groups that mainly contain low-eccentric stars. We compare the slopes of their chemical clocks with the ones we

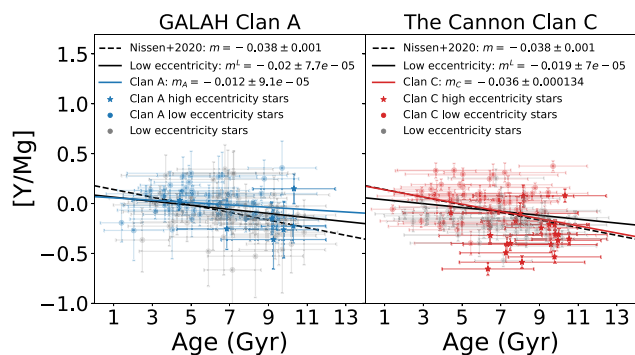


Figure 13. Age–[Y/Mg] relations of the stars in Clans A (GALAH DR3) and C (*The Cannon*) selected from the trees of Fig. 10. Linear regression fits are drawn to quantify the slopes in these relations. A linear fit using all low-eccentricity stars selected for the analysis is displayed as a solid black line. Trend found by Nissen et al. (2020) displayed as a dashed black line. [Y/Mg] errors taken from the internal uncertainties reported by GALAH DR3 and our *The Cannon* catalogue. Age errors taken from GALAH DR3 Buder et al. (2021).

would obtain if we considered all the 100 low-eccentricity stars in the sample. Fig. 13 summarizes this result. On the left panel, we plot again in blue the stars belonging to Clan A and the blue solid line is the resulting linear fit to these stars. Since Clan A does not only contain low-eccentric stars, for better illustration of our findings, we plot with different symbols the stars with circular orbits (in circles) and eccentric orbits (in stars). The resting stars in low-eccentricity orbits are plotted with grey circles, and the linear fit is shown in the black solid line. Finally, for reference, the chemical clock fit found by Nissen et al. (2020) is shown with the dashed line, as before. All slopes are specified in the legend. For Clan C, which considers the abundances obtained by us with *The Cannon*, the same information is shown in the right-hand panel, and the stars are coloured in red.

As we see in Fig. 13, by just considering the low-eccentricity stars, we are not able to find trends that are consistent within the literature, even considering the uncertainties. This might be because we are missing important older stars which now are on less circular orbits. A cut on dynamical properties only is therefore not sufficient since the selected groups are incomplete (Soubiran & Girard 2005; Hawkins et al. 2015). When using the tools provided by phylogenetic analyses, we can chemically identify different groups of stars and find patterns that could be associated to their shared history. We stress that we obtain this result only for our new high-precision abundances, demonstrating also the importance of having very high-precision abundances for a better selection of stellar populations.

4.4 Discussion

In this section, we discuss the prospects and limitations of a phylogenetic analysis on a sample of solar twin chemical data.

4.4.1 On the clustering methods

The NJ algorithm is essentially a clustering algorithm, not particularly different to others available in the literature (see e.g. Ratcliffe et al. 2020, for descriptions and discussions of different clustering algorithms used in chemical data). What makes the NJ algorithm however attractive here is that first we do not have to specify the number of clusters we aim to find, unlike other fast clustering

algorithms such as K-means. This is important when studying the relationships and shared history of a group of stars as a whole, where we are not primarily interested in finding groups but in studying the way in which the entire system is ordered and how this order might tell us something about their evolutionary history.

Second, the NJ algorithm is not designed to only cluster the data, but to visualize the amount of divergence between pairs of objects. This translates into branch lengths that have a meaning of difference. NJ trees therefore are not expected to have the objects aligned at the tips of the tree, in contrast to other dendrograms obtained by clustering algorithms such as DBSCAN or HDBSCAN (Casamiquela et al. 2021b) or other methods (Ratcliffe et al. 2020). The procedure to define the branching pattern only depends on the distance matrix. Other clustering algorithms require to specify parameters of closeness and density in the parameter space, which cannot be constrained in an objective way (Casamiquela et al. 2021b).

In fact, here, we are not primarily interested in finding clusters and quantifying the number of clusters and their properties, but to visualize how the data are structured in their hierarchical order. Because of the heritable information, we consider to build the trees, that order can be used to interpret shared histories, which is the essence of phylogenetics.

4.4.2 On the robustness of solar twin trees

We comment on the poor support of our trees, where in the best case, we still have a large majority of nodes with a support below 50 per cent (see Fig. 9). Considering that the NJ algorithm takes a distance matrix as input and joins the elements that are closest to each other, by construction, it will generate a tree that will reflect the hierarchical order of the distance matrix. But when the range of differences is small, a perturbation of that distribution given the uncertainties will imply a very different tree. The distance method for tree reconstruction becomes uncertain if the distances are too small for the entire sample (Yang 2014).

There is another source of uncertainty here, which is the fact that by considering loose stars in the disc, we can not rule out the possibility that two stars will have the same origin (e.g. be siblings of the same star formation episode of the same molecular cloud). The NJ will place these two stars in two different tips of leaves in the tree (by construction), but they in fact represent one leaf. Without a previous selection of stars belonging to distinct star formation episodes, the NJ algorithm will fail in ranking the leaves, because any hierarchical order obtained for stars from the same star formation episode will be driven by the errors in the abundances and the intrinsic dispersion of such populations.

We also have to keep in mind that there will be a noise because of the ISM inhomogeneities that does not reflect evolution (Kos et al. 2021; Ness et al. 2022). The hope is that noise is less than the change due to evolution (Manea, Hawkins & Maas 2022). In that sense, nodes of poor support could also be used to identify conatal stars. To test these possibilities, simulated data are more suitable, because in that case, we know with certainty the origin of the data (de Brito Silva et al. 2023).

This translates on a node support of less than 50 per cent in the case of *The Cannon*, and below 30 per cent in the case of GALAH DR3 abundances, making the interpretation of the chemical–clock relations of the selected clans uncertain. Recently, de Brito Silva et al. (2023) showed using simulations that uncertainties below 0.08 dex were necessary to have enough phylogenetic signal for

history reconstruction. That work however can not be directly applied to our study, because our metallicity range is much smaller, our real Galaxy did not live in isolation like the simulated galaxy of de Brito Silva et al. (2023), and our stars might come from the same birth cloud whereas de Brito Silva et al. (2023) had stellar particles which may reflect single stellar populations depending on the numerical resolution of the simulation.

We comment that repeating this analysis with different selections of 200 stars did not reproduce the chemical clock trends found in Fig. 13, making us conclude that the selection effects of our stars is biasing our results, and that the phylogenetic signal for solar twins in GALAH data is still too weak for this kind of study, regardless of the high precision of our derived abundances. We need a sample with a more extended chemical space and metallicity range to keep advancing in this topic with observations.

4.4.3 On the chemical evolution of solar twins

Using solar twins for this kind of study might also add further challenges in the interpretation of the results. Having a sample of a restricted range in metallicity makes it impossible to trace back the population that was formed from the pristine gas (e.g. with no metals). The fact that all stars are metal enriched tells us that we are studying a population from a stage in which considerable chemical enrichment already happened. Stars of different ages and different $[Y/Mg]$ but similar $[Fe/H]$ might well be formed at different Galactic environments and arrive to the solar neighbourhood through a dynamical process that can be radial migration or heating.

Given the different time-scales in the pollution of Mg, Y, and Fe into the ISM, it is not straightforward to interpret a clan which has a tight relationship of $[Y/Mg]$ and age but has a limited range in $[Fe/H]$. Is it that the NJ is simply clustering stars of different birth radii which trace different evolutionary histories of the ISM? Perhaps the AGB stars, which live longer than the progenitors of SN II, had enough time to radially migrate and pollute the ISM with Y at a different location than their birth place in contrast to the pollution of Mg by their massive siblings (Johnson et al. 2021). It is possible that the ISM has a shared evolutionary history at a wider range in radii due to migration. The interplay between migration, heating, blurring, and mixing in the ISM of the Milky Way is still poorly understood (Feltzing, Bowers & Agertz 2020). Phylogenetic methods might thus offer an interesting opportunity to learn more about these processes.

We also need to remind ourselves of the selection effect in our sample, where the selection of solar twins systematically underrepresents young stars. As explained in Sharma et al. (2018), the age estimation for both the youngest and oldest stars via the isochrone fitting is biased towards intermediate ages.

As seen from Figs 9, 12, and 13, GALAH data still do not have the precision in the abundances required to perform a robust phylogenetic study. However, GALAH data have been central for this analysis, because *The Cannon* uses the best GALAH results to perform a reanalysis which then allows us to apply phylogenetic techniques in other GALAH stars. Considering that future data releases of GALAH are expected to increase in precision, the prospects of phylogenetic studies in GALAH are still very promising.

5 CONCLUSIONS

In this paper, we have performed a systematic application of the machine learning algorithm *The Cannon* to a set of solar twins observed and analysed in GALAH DR3 (Buder et al. 2021) with

the aim to provide a catalogue of high precision abundances of 38 716 solar twins and use a set of this catalogue for a phylogenetic study of GALAH data. Other scientific applications of high-precision abundances of solar twins include setting constraints on planet engulfment processes (Bedell et al. 2018; Maia et al. 2019) or the level of homogeneity in star formation regions such as open clusters or wide binaries (Liu et al. 2019; Hawkins et al. 2020; Espinoza-Rojas et al. 2021). Therefore, our catalogue can be used to explore these subjects, in addition to our primary intention to perform a phylogenetic analysis.

In the systematic application of *The Cannon* for the generation of this catalogue, we investigated the impact of the labels considering different training sets. We first varied the size of the training set, and then studied the label recovery when removing outliers. This analysis helped us to conclude that a training set with 150 GALAH stars of $SNR > 117$ was sufficient for predicting precise labels of stellar parameters and 14 chemical abundances of solar twins observed with GALAH.

Our results agree within 50 K in temperature, 0.09 dex in $\log g$, 0.03 in $[Fe/H]$, and in 0.05 dex in abundances with GALAH for stars $SNR > 50$. For lower SNR, the results agree less, within 60 K in temperature, 0.1 dex in gravity, 0.07 dex in metallicity and 0.1 dex in the other abundances. This is expected considering that also GALAH data are more uncertain for lower SNR. The internal uncertainties of our model at lower SNR do not significantly increase compared to high-SNR results. The consistency of predicted labels for repeat observations remains comparable to the GALAH ones, provided both GALAH and *The Cannon* used the same pixels to extract the information.

Our new catalogue allows us to test phylogenetic studies on the solar neighbourhood that require high precision abundances. We analysed 200 stars separated in two eccentricity groups, namely a group with circular orbits and another one with orbits of eccentricity around 0.4, and we compared the trees obtained for these groups using GALAH and *The Cannon* abundances. In both cases, we were able to find clans which were distinct in eccentricities, with one clan notably grouping the stars with circular orbits. While the node support in the *The Cannon* tree is higher than for the GALAH tree, the overall support in both trees is not outstanding. This is expected for a sample of stars which are so similar to each other that the hierarchical differences between stars are very small and thus uncertain. It is also possible that many of these solar twins are tracers of the same star formation episode, causing a conflict in the tree shape which is forced by the neighbour joining algorithm. To truly study the support and amount of information carried out in chemical abundances, simulated data should be used instead, even if simulated data are prone to systematic uncertainties (de Brito Silva et al. 2023).

The trees still allowed us to discuss the astrophysical nature of the clans found. To this aim, we compared age- $[Y/Mg]$ relation (e.g. chemical clock) of the clans with the literature (Nissen et al. 2020), obtaining an agreement for one of our clans with our *The Cannon* abundances, but no agreement with the clans found in the GALAH tree. The agreement between the chemical clock of Nissen et al. (2020) and ours was obtained noting that the clan included mostly stars with circular orbits, but some older stars with eccentric orbits as well. Indeed, the age- $[Y/Mg]$ relation found for only stars with circular orbits does not agree with the chemical clock obtained by Nissen et al. (2020), because of the lack of old stars with circular orbits. We discussed that this result was highly dependent on the selection of the stars. The full potential of phylogenetic studies in large catalogues can thus be further explored using

stellar samples with more extended abundance ranges than solar twins.

Our work sets a further basis for the promising future of Galactic phylogenetics, in which we can use large spectroscopic surveys like GALAH with machine learning to improve the chemical abundances which then can be used as input for phylogenetic analyses and so reconstruct the history of our home Galaxy, the Milky Way.

ACKNOWLEDGEMENTS

This work has been funded by Millennium Nucleus ERIS NCN2021_017, Centros ANID Iniciativa Milenio and ANID 2023 Master's Scholarship Program 22230932. We acknowledge FONDECYT Regular Grants 1200703 and 1231057. PJ thanks the Stromlo Distinguished Visitor Program at the ANU. EJJ acknowledges support from FONDECYT Iniciación en Investigación 2020 Project 11200263. We warmly acknowledge Sara Vitali, Danielle de Brito Silva, and Scarlet Elgueta for fruitful conversations throughout the development of this project.

DATA AVAILABILITY

Our catalogue of high precision solar twins from GALAH is available at CDS via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5) or via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/MNRAS>. The spectra used in this work can be directly obtained from the GALAH survey. Spectra of repeated observations can be downloaded upon request.

REFERENCES

- Abdurro'uf et al., 2022, *ApJS*, 259, 35
- Adibekyan V. Z., Sousa S. G., Santos N. C., Delgado Mena E., González Hernández J. I., Israelian G., Mayor M., Khachatryan G., 2012, *A&A*, 545, A32
- Aguilera-Gómez C., Jones M. I., Chanamé J., 2023, *A&A*, 670, A73
- Amarsi A. M. et al., 2020, *A&A*, 642, A62
- Ambrosch M. et al., 2023, *A&A*, 672, A46
- Aumer M., Binney J., Schönrich R., 2016, *MNRAS*, 462, 1697
- Bailer-Jones C. A. L., Rybizki J., Foesneau M., Mantelet G., Andrae R., 2018, *AJ*, 156, 58
- Battistini C., Bensby T., 2015, *A&A*, 577, A9
- Baum D. A., Smith S. D., Donovan S. S. S., 2005, *Sci. Perspect.*, 310
- Bedell M. et al., 2018, *ApJ*, 865, 68
- Bensby T., Feltzing S., Oey M. S., 2014, *A&A*, 562, A71
- Berger T. A., van Saders J. L., Huber D., Gaidos E., Schlieder J. E., Clayton Z. R., 2022, *ApJ*, 936, 100
- Bird J. C., Loebman S. R., Weinberg D. H., Brooks A. M., Quinn T. R., Christensen C. R., 2021, *MNRAS*, 503, 1815
- Buder S. et al., 2018, *MNRAS*, 478, 4513
- Buder S. et al., 2021, *MNRAS*, 506, 150
- Casali G. et al., 2020, *A&A*, 639, A127
- Casamiquela L. et al., 2021a, *A&A*, 652, A25
- Casamiquela L., Castro-Ginard A., Anders F., Soubiran C., 2021b, *A&A*, 654, A151
- Casey A. R. et al., 2017, *ApJ*, 840, 59
- de Brito Silva D. et al., 2023, preprint ([arXiv:2310.12235](https://arxiv.org/abs/2310.12235))
- De Laverny P., Recio-Blanco A., Worley C. C., Plez B., 2012, *A&A*, 544, A126
- De Silva G. M. et al., 2015, *MNRAS*, 449, 2604
- Delgado Mena E. et al., 2019, *A&A*, 624, A78
- Escorza A. et al., 2019, *A&A*, 626, A128
- Espinoza-Rojas F., Chanamé J., Jofré P., Casamiquela L., 2021, *ApJ*, 920, 94
- Felsenstein J., 1988, *Annu. Rev. Genetics*, 22, 521
- Feltzing S., Bowers J. B., Agertz O., 2020, *MNRAS*, 493, 1419
- Freeman K., Bland-Hawthorn J., 2002, *ARA&A*, 40, 487
- Gaia Collaboration et al., 2018, *A&A*, 616, A1
- Gilmore G. et al., 2022, *A&A*, 666, A120
- Guiglion G. et al., 2020, *A&A*, 644, A168
- Gustafsson B., Edvardsson B., Eriksson K., Jørgensen U. G., Nordlund Å., Plez B., 2008, *A&A*, 486, 951
- Hawkins K. et al., 2020, *MNRAS*, 492, 1164
- Hawkins K., Jofré P., Masseron T., Gilmore G., 2015, *MNRAS*, 453, 758
- Heiter U., Jofré P., Gustafsson B., Korn A. J., Soubiran C., Thévenin F., 2015, *A&A*, 582, A49
- Jackson H., Jofré P., Yaxley K., Das P., de Brito Silva D., Foley R., 2021, *MNRAS*, 502, 32
- Jofré P. et al., 2014, *A&A*, 564, A133
- Jofré P., Das P., Bertranpetit J., Foley R., 2017, *MNRAS*, 467, 1140
- Jofré P., Heiter U., Soubiran C., 2019, *ARA&A*, 57, 571
- Jofré P., Heiter U., Tucci Maia M., Soubiran C., Worley C. C., Hawkins K., Blanco-Cuaresma S., Rodrigo C., 2018, *Res. Notes Am. Astron. Soc.*, 2, 152
- Jofré P., Jackson H., Tucci Maia M., 2020, *A&A*, 633, L9
- Johnson J. W. et al., 2021, *MNRAS*, 508, 4484
- Kobayashi C., Karakas A. I., Lugaro M., 2020, *ApJ*, 900, 179
- Kos J. et al., 2017, *MNRAS*, 464, 1259
- Kos J. et al., 2021, *MNRAS*, 506, 4232
- Leung H. W., Bovy J., 2019, *MNRAS*, 483, 3255
- Lind K., Korn A. J., Barklem P. S., Grundahl F., 2008, *A&A*, 490, 777
- Liu F., Asplund M., Yong D., Feltzing S., Dotter A., Meléndez J., Ramírez I., 2019, *A&A*, 627, A117
- Lu Y. L., Ness M. K., Buck T., Zinn J. C., Johnston K. V., 2022, *MNRAS*, 512, 2890
- Maia M. T., Meléndez J., Lorenzo-Oliveira D., Spina L., Jofré P., 2019, *A&A*, 628, A126
- Manea C., Hawkins K., Maas Z. G., 2022, *MNRAS*, 511, 2829
- Manea C., Hawkins K., Ness M. K., Buder S., Martell S. L., Zucker D. B., 2023, preprint ([arXiv:2310.15257](https://arxiv.org/abs/2310.15257))
- Miglio A. et al., 2013, *MNRAS*, 429, 423
- Nandakumar G. et al., 2022, *MNRAS*, 513, 232
- Ness M. K., Wheeler A. J., McKinnon K., Horta D., Casey A. R., Cunningham E. C., Price-Whelan A. M., 2022, *ApJ*, 926, 144
- Ness M., 2018, *PASA*, 35, e003
- Ness M., Hogg D. W., Rix H. W., Ho A. Y. Q., Zasowski G., 2015, *ApJ*, 808, 16
- Nissen P. E., 2015, *A&A*, 579, A52
- Nissen P. E., 2016, *A&A*, 593, A65
- Nissen P. E., Christensen-Dalsgaard J., Mosumgaard J. R., Silva Aguirre V., Spitoni E., Verma K., 2020, *A&A*, 640, A81
- Nissen P. E., Gustafsson B., 2018, *A&AR*, 26, 6
- O'Brian T., Ting Y.-S., Fabbro S., Yi K. M., Venn K., Bialek S., 2021, *ApJ*, 906, 130
- Piskunov N., Valenti J. A., 2017, *A&A*, 597, A16
- Prša A. et al., 2016, *AJ*, 152, 41
- Randich S. et al., 2022, *A&A*, 666, A121
- Ratcliffe B. L., Ness M. K., Johnston K. V., Sen B., 2020, *ApJ*, 900, 165
- Ratcliffe B., Minchev I., Cescutti G., Spitoni E., Jönsson H., Anders F., Queiroz A., Steinmetz M., 2024, *MNRAS*, 528, 3464
- Retzlaff N., Stadler P. F., 2018, *Theory Biosci.*, 137, 133
- Ruiz-Lara T., Gallart C., Bernard E. J., Cassisi S., 2020, *Nat. Astron.*, 4, 965
- Saitou N., Nei M., 1987, *Mol. Biol. Evol.*, 4, 406
- Sharma S. et al., 2018, *MNRAS*, 473, 2004
- Sheinis A. et al., 2015, *J. Astron. Telesc. Instrum. Syst.*, 1, 035002
- Skrutskie M. F. et al., 2006, *AJ*, 131, 1163
- Slumstrup D., Grundahl F., Brogaard K., Thygesen A. O., Nissen P. E., Jessen-Hansen J., Van Eylen V., Pedersen M. G., 2017, *A&A*, 604, L8
- Soubiran C., Girard P., 2005, *A&A*, 438, 139
- Souto D. et al., 2018, *ApJ*, 857, 14
- Spina L., Meléndez J., Karakas A. I., Ramírez I., Monroe T. R., Asplund M., Yong D., 2016, *A&A*, 593, A125
- Ting Y.-S., Conroy C., Rix H.-W., Cargile P., 2019, *ApJ*, 879, 69
- Valenti J. A., Piskunov N., 1996, *A&AS*, 118, 595

Vitali S., Slumstrup D., Jofré P., Casamiquela L., Korhonen H., Blanco-Cuaresma S., Winther M. L., Aguirre Børsen-Koch V., 2024, arXiv e-prints, p. arXiv: 2401.02328

Wheeler A. et al., 2020, *ApJ*, 898, 58

Yang Z., 2014, *Molecular evolution: a statistical approach*. Oxford Univ. Press, Oxford

SUPPORTING INFORMATION

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

suppl_data

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.

APPENDIX A: IMPACT OF TRAINING CHOICES ON LABEL RECOVERY WITH THE CANNON

In this appendix section, we discuss the results of different training choices on the label recovery of *The Cannon*. We perform the three main studies which are discussed in the following sections. The first study is the impact of the predicted labels and their uncertainties for different training set sizes. The second study is the reliability of the covariance uncertainties given the different training set sizes, and the third study is the accuracy and precision of the predicted labels when the training set has label outliers.

A1 Bias as function of training set size

We consider only the high-SNR sample of 5028 stars described in Section 2. From there, we perform a random split where we take 80 per cent of this sample for training set selection, and the remaining 20 per cent as a fixed test set for the different models built. Thus, we have fixed 1006 stars for testing and 4022 stars to build different training sets.

From the 4022 spectra available for training, we study training sizes $N = 10, 20, 30, 50, 100, 150, 200, 250, 300, 350, 400, 500, 1000, 2000, 3000$. We further choose randomly 10 subsamples of size $N_1 = 10$ and train 10 different *The Cannon* models, each one using a different subsample of size N_1 as its training set. This allows us to evaluate the impact of the choice of stars in each set. Then for $N_2 = 20$, we select the first subsample of size $N_1 = 10$ and randomly add the missing $N_2 - N_1$ spectra adding up to a training set of N_2 spectra. This random addition of the first subsample of size N_1 is done 10 times, generating 10 subsamples of size $N_2 = 20$. We again train 10 different *The Cannon* models, each one using a different subsample of size N_2 as its training set. This procedure is followed for all the rest of the training sets and sizes.

With these models, we predict the labels for the 1008 stars in our test set and compare them with the values in GALAH DR3, which is shown in Fig. A1. In the top panel, we show the median difference (bias) between the predicted labels and GALAH DR3 in training as function of the training set size. Each line corresponds to a different label with error bands representing the 16th and 84th percentiles in the distribution of the results obtained by the 10 different *The Cannon* models. Some labels are highlighted with colour as example. The bias between the predicted labels and GALAH DR3 in test as function of training size is shown in the bottom panel. The difference between the GALAH DR3 labels and predicted *The Cannon* labels are small for all training sizes N . $\log g$ and $[Y/Fe]$ show a minor bias increase with

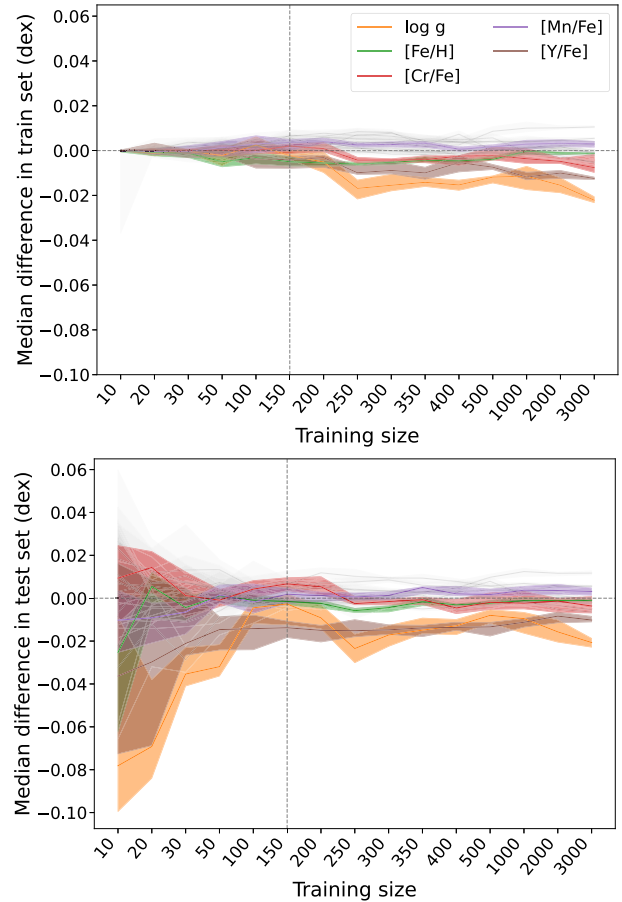


Figure A1. Top panel: Bias in train as function of training size N . Bands represent the 16th and 84th percentile of the values obtained for the 10 different trained models with training sets of size N . Bottom panel: Bias Median in test as function of training size. For better visualisation of the results, effective temperature is not shown here.

N up to 0.01–0.02 dex. For $N = 10$, the differences are essentially zero for all the labels, showing the overfitting effect in *The Cannon*.

We comment the bias in the parameter $\log g$ and its similarity with $[Y/Fe]$, which are highlighted with orange and brown colours, respectively. GALAH DR3 surface gravities are not determined directly from the spectra, but from photometry and the parallax (Buder et al. 2021) because GALAH spectra do not contain sufficient dependency of this parameter. It is hence expected that the *The Cannon* model does a poor job in predicting this parameter.

The impact of the bias in $\log g$ with $[Y/Fe]$ is because our *The Cannon* model considers the ionized Y lines in GALAH spectra. Ionized lines have a dependency on surface gravity. If the surface gravity is poorly determined, it is expected that a method deriving abundance of ionized lines of a given strength will respond by balancing the ill-determination of surface gravity with an ill-determination of that abundance. Most of the other abundances are derived from neutral lines, which are less sensitive to gravity. From the bottom panel of Fig. A1, we observe that for most of the labels, the differences show a decreasing trend with training size N . The trend reaches a plateau at around $N = 150$ which is marked with a vertical dashed line.

To further assess the potential problems of overfitting in the training process, we compute the *mean squared error (MSE)*. This allows us to assess the quality of the predictions made by the different

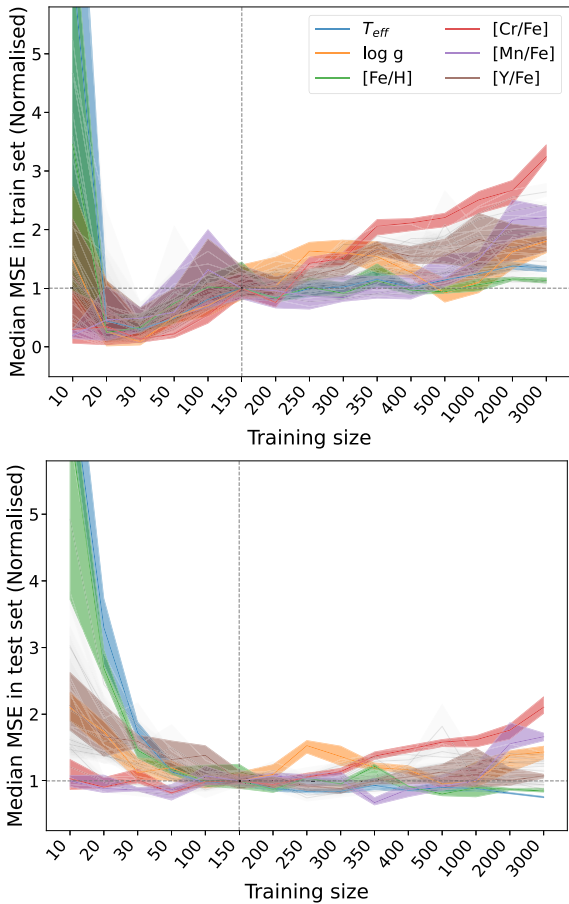


Figure A2. Top panel: Median of MSE in train as function of training size N . Error bands represent the 16th and 84th percentile of the values obtained for the 10 different trained models with training sets of size N . Bottom panel: Median of MSE in test. Top and bottom panel normalized in value at training size $N = 150$.

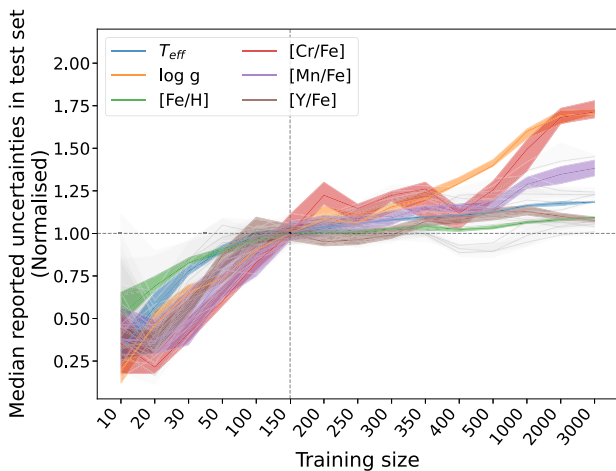


Figure A3. Median of the reported uncertainties in test as function of training size. Normalised at value in training size $N = 150$. Error bands represent the 16th and 84th percentile of the values obtained for the 10 different trained models with training sets of size N .

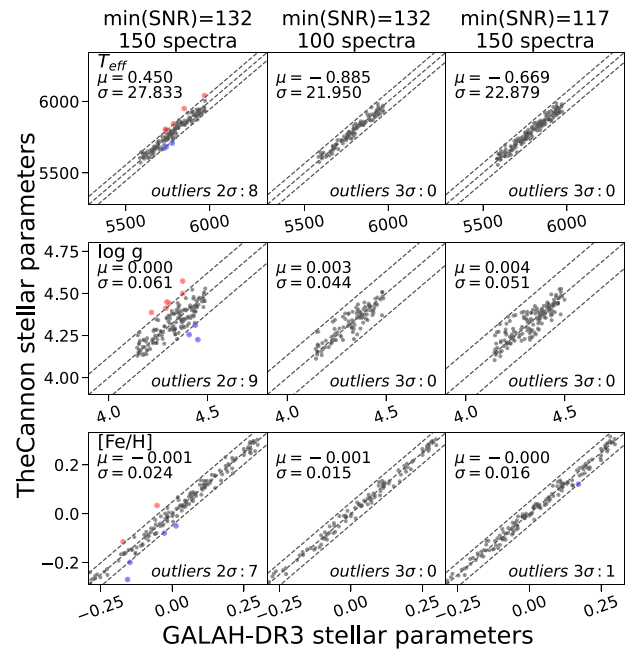


Figure A4. One-to-one comparison of stellar parameters T_{eff} , $\log g$, $[\text{Fe}/\text{H}]$. Comparison of 3 set-ups/models. Model 1: 150 spectra of minimum SNR of 132 (left column), Model 2: 100 spectra of minimum SNR of 132 (middle column), and Model 3: 150 spectra of minimum SNR of 117 (right column). In x -axis, GALAH DR3 labels. In y -axis, *The Cannon* model estimates for labels. Outer dashed lines correspond to 2σ and 3σ boundaries for Model 1 and Models 2, 3, respectively. Overestimates above the upper dashed line are plotted in red, underestimates below the lower dashed line are plotted in blue. Upper left and bottom right of each panel shows the median μ and standard deviation σ of the difference, and the number of outliers found outside each boundary, respectively.

The Cannon models in the test step, as well as the *self-test*, i.e. the test/label prediction step over the training set itself. In other words, we can evaluate how well the model recovers the labels for the stars considered in the training process. This is seen in Fig. A2. The top and bottom panel show the median MSE in training and test, respectively, as a function of the training size N . Each line corresponds to a different label. We normalize the results to the ones obtained at $N = 150$ for a better visualization, and the error bands cover the range of values between the 16th and 84th percentile over the distribution of the results obtained by the 10 different *The Cannon* models.

When $N = 10$, the MSE in train is large with a high discrepancy between the 10 different *The Cannon* models. The MSE significantly drops for $N = 20$, but starts increasing again with the training size. The bottom panel shows that the MSE in test starts high for small training sets, decreasing monotonically until a training size of $N = 150$ and remain small for higher N . Our analysis shows that for a training set smaller than 150 stars, *The Cannon* is affected by overfitting if 14 labels from solar twin GALAH spectra are estimated. For training sets larger than 150 stars, we do not see a significant improvement in the bias and the MSE of the model. Therefore, we conclude that 150 stars is an optimal size for our *The Cannon* model of solar twin stars.

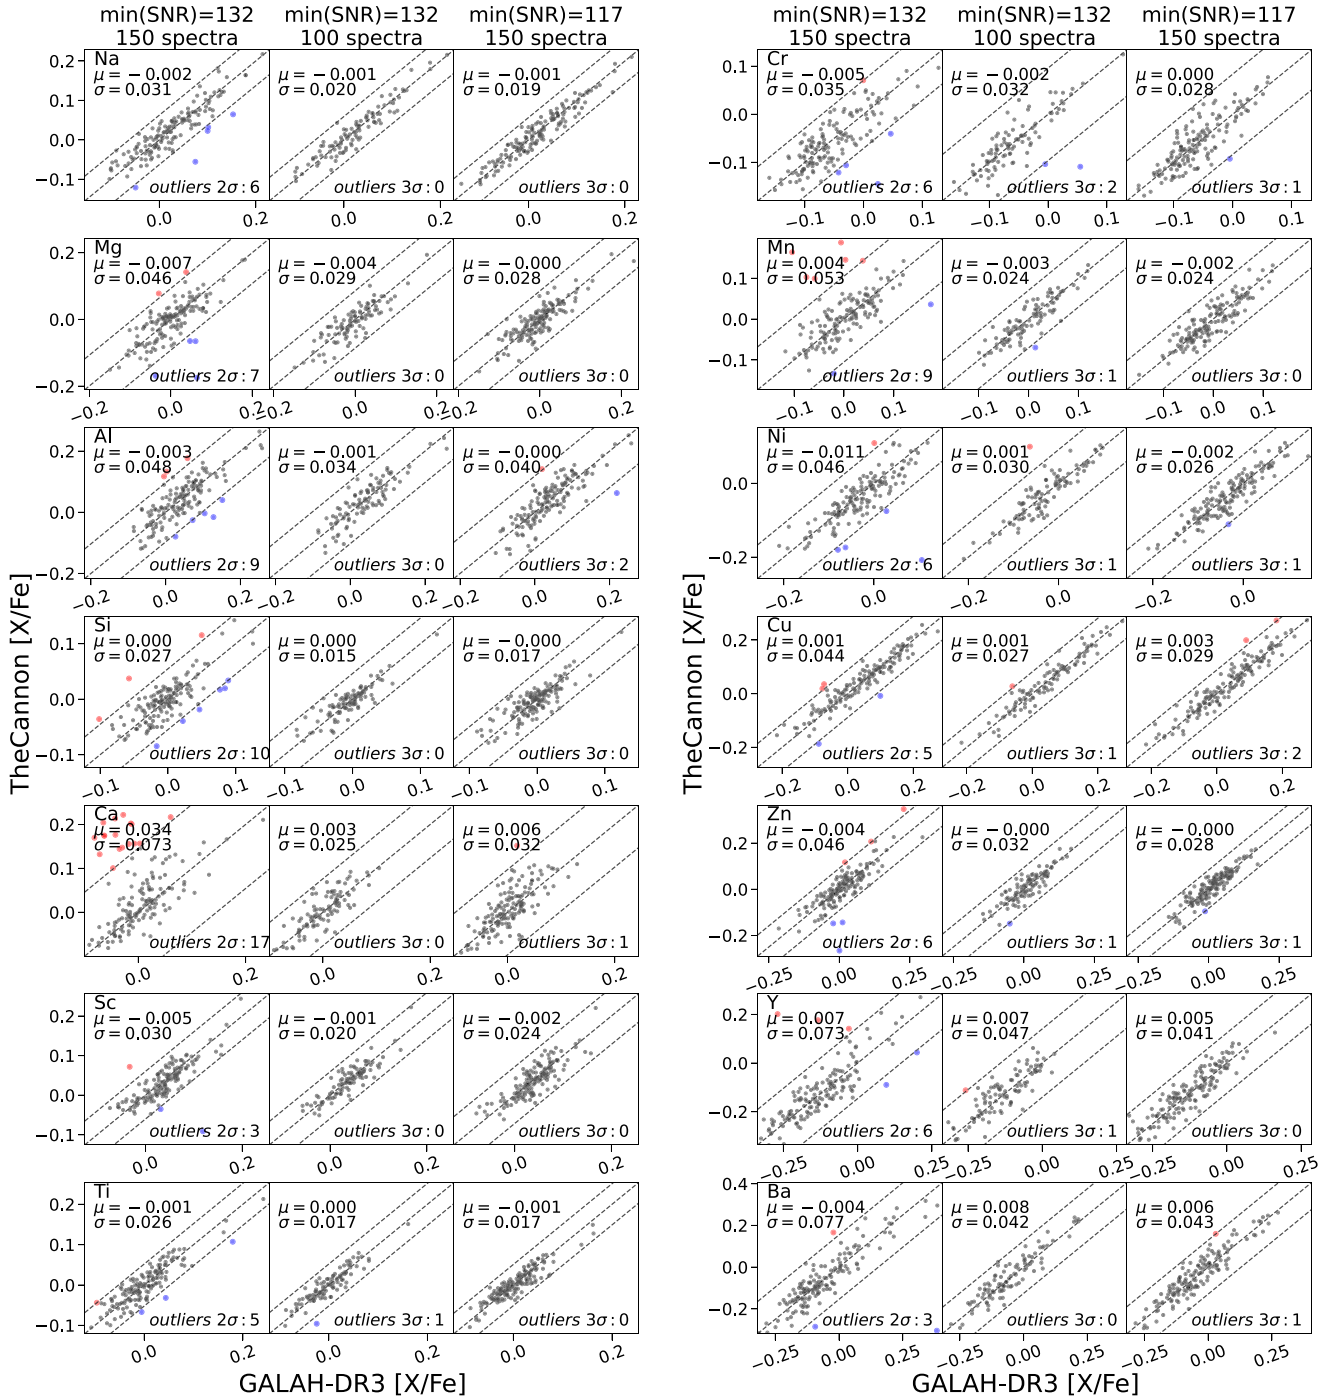


Figure A5. One-to-one comparisons for chemical abundances for the 3 set-ups/models. Model 1: 150 spectra of minimum SNR of 132 (left column), Model 2: 100 spectra of minimum SNR of 132 (middle column), and Model 3: 150 spectra of minimum SNR of 117 (right column). In x -axis, GALAH DR3 labels. In y -axis, *The Cannon* model estimates for labels. Outer dashed lines correspond to 2σ and 3σ boundaries for Model 1 and Models 2, 3, respectively. Overestimates above the upper dashed line are plotted in red, underestimates below the lower dashed line are plotted in blue. Upper left and bottom right of each panel shows the median μ and standard deviation σ of the difference, and the number of outliers found outside each boundary, respectively.

A2 Reliability of precision estimates (covariance uncertainties)

In Fig. A3, we show the median reported internal uncertainties for *The Cannon* as a function of the training size N . The lines and error bands follow the same definition as in Figs A1 and A2. We recall that the internal uncertainties in SME are computed from the diagonal

of the covariance matrix given by SME fitting procedures. With *The Cannon*, we compute the internal uncertainties in the same way.

There is a similar trend for all labels, where the median internal uncertainties in test increase with N . The stellar parameters T_{eff} and $[\text{Fe}/\text{H}]$ increase slowly, reaching a value up to 20 percent at $N = 150$. When focusing on Y, we find that values of $150 <$

$N \leq 300$ reports smaller uncertainties, but steadily increase for larger N . For $\log g$, Cr and Mn, after $N = 150$ the uncertainties considerably increase, reaching values nearly 70 per cent higher in the case of $\log g$ and Cr, and 30 per cent higher for Mn. For the smallest training set sizes, we attribute the low uncertainties to overfitting. In particular for $\log g$, we attribute this increase to the disagreement of spectroscopic and photometric $\log g$ (compare to top panel of Fig. A1). For GALAH DR3, $\log g$ was actually not estimated from the spectra, but photometric information due to the usually lower information content in spectra. This approach can, however, fail for a variety of reasons like binarity or wrong photometric or astrometric information and thus introduce an increasingly systematic trend in the training. While we tried to exclude peculiar spectra from our training set via visual inspection, we suspect that for training sizes above 500, residual peculiar spectra caused the rise of training imperfections and the reported uncertainties. Some further imperfections are also not visible noted, and lead to training labels which are incorrect. Such training imperfections, as evidenced by the rising MSE in Fig. A2 for Cr, can then propagate into the test uncertainties.

Hereafter, we consider $N = 150$ stars as our final training set size. We can conclude that this size is sufficient for making a *The Cannon* model that predicts most labels with high accuracy and sufficient precision. There is not much improvement in both accuracy and precision for higher values of N .

A3 Robustness against outliers

We investigate how robust is our *The Cannon* model against potential outliers in the training set. We choose the optimal size of 150 stars for this analysis. To do so, we first select 150 stars with the highest SNR in our sample (hereafter, Model 1). This yields a training set of stars with SNR above 132 in all CCDs. Then, we train a *The Cannon* model and perform a self-test. The results for Model 1 can be found on the left panels in Fig. A4 for stellar parameters and Fig. A5 for chemical abundances. We set up the 3σ boundaries for all labels and define for each one of them the outliers those who lie outside this boundary. In the figures, these outliers are plotted with red and blue colours that lie outside the dashed lines. In general, the predictions about the parameters agree with GALAH DR3. However, in Fig. A5, we begin to see outliers for some of the chemical abundances, in particular as Ca and Mn. There are four Ca lines, but two of them lie in telluric lines where the correction may not always be perfect (Buder et al. 2018). In GALAH DR3, the spectrum uncertainty is increased for these lines to account for strong blends of the telluric lines. Here, with *The Cannon*, we take the spectrum uncertainty directly from the data base, obtaining overabundances of Ca because the model finds a high absorption feature (Buder et al. (2018)).

We took two different approaches to deal with the outliers. The first one was to remove all the 50 outliers found in the self-test and train a new *The Cannon* model with the remaining 100 stars of minimum SNR 132 (hereafter, Model 2). The second one was to remove the outliers in the first model and set a lower threshold in SNR to add more stars and build up a new training set of size 150 to then train a new *The Cannon* model. In this second case, that this model produced new outliers, so we reiterated the process by removing such outliers and setting a lower minimum SNR threshold to build a new training set of size 150. After 4 iterations, we converged to a *The Cannon* model trained with a set of 150 stars with minimum SNR of 117 (hereafter, Model 3).

The results of these two approaches can be seen in the middle and right panels of Figs A4 and A5, for Model 2 and Model 3,

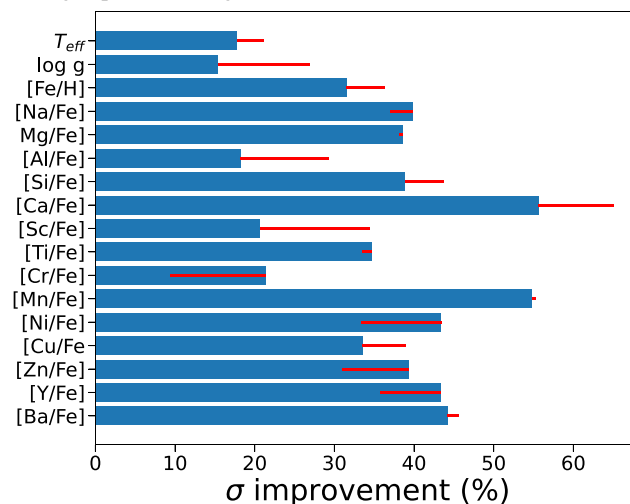


Figure A6. σ improvements with respect to Model 1 for all labels (in percentage). Improvement given by model 3 in blue bins and improvement/decay in percentage made by model 2 in red lines.

respectively. There are no outliers for the stellar parameters in both approaches. The same holds for the chemical abundances in Fig. A5 where most labels have either no outliers or very few.

A4 Choosing final training set

From Figs A4 and A5, we can see that Models 2 and 3 are an improvement with respect to Model 1 in terms of agreement in with respect to GALAH DR3. Fig. A6 shows the overall percentage improvement of the dispersion σ in each label, given by the two latter models with respect to Model 1. The blue bins represent the percentage of improvement for Model 3, and the red lines represent the increment/decay in percentage of the improvement in σ by the Model 2. For stellar parameters, we observe a higher improvement given by Model 2, with a difference of 3.3 per cent, and 4.9 per cent for the T_{eff} and metallicity, respectively. For $\log g$, we have a considerable difference of 11.5 per cent in favour of Model 2. For chemical abundances, we observe negligible differences for Mg, Ti, Mn, Ba where the differences are up to 1.4 per cent. For Al, Si, Ca, Sc, Cu, we find higher differences up to 13.7 per cent in favour of Model 2, and for Na, Cr, Ni, Zn, Y, we obtain differences up to 11.9 per cent in favour of Model 3. The mean improvement over all labels is 36.2 per cent and 34.8 per cent for models 2 and 3, respectively.

Taking into consideration that the overall difference of 2.6 per cent in mean improvement for all the labels is very small, we choose Model 3 for training. This model has an optimal size of 150 as well as a better coverage in the parameter space of stellar parameters and 14 chemical abundances for solar twins.

APPENDIX B: STARS SELECTED FOR PHYLOGENETIC ANALYSIS

Table B1 shows the ages and eccentricities of 50 stars used in the tree. The information about the rest can be found online.

Table B1. Example of 50 stars selected for phylogenetic analysis (see Fig. 10). Tip ID corresponds to the integer displayed in the tips of the trees. For reference, we show the *Gaia* DR3 IDs, and the ages and eccentricities (ϵ). The information of the rest of the stars can be found online as supplementary material.

Tip ID	<i>Gaia</i> DR3 ID	Age (Gyr)	ϵ
0	5396076243592498944	8.74	0.6303
1	3814762674970623104	7.50	0.6109
2	6690566020667409024	8.49	0.5699
3	56608356156072448	10.90	0.5324
4	2532978746890377472	10.26	0.5305
5	6557301912844603264	10.53	0.5257
6	4710178345898217728	10.76	0.5183
7	3747102371529922560	10.11	0.5142
8	2582225773914697728	5.00	0.5024
9	6247059932386652288	3.91	0.4902
10	6483285057806757376	9.39	0.4796
11	4475030814892134912	11.21	0.4793
12	5416434491653412480	9.05	0.4750
13	3201058165300460928	9.51	0.4748
14	3910048792175067648	10.14	0.4660
15	5827969736762248192	5.17	0.4607
16	6665136481380692992	5.97	0.4605
17	6811053318039677184	10.53	0.4584
18	6094499086259119360	9.81	0.4574
19	2558440382468221184	10.33	0.4530
20	5565471677892685184	10.29	0.4527
21	3206464940714422144	7.26	0.4523
22	3462833192176962688	8.99	0.4523
23	3753282382791867904	9.78	0.4504
24	6579941746317849344	5.57	0.4498
25	3203847042185006080	11.39	0.4497
26	2616281875274482560	9.26	0.4491
27	685943836661554944	7.10	0.4469
28	5566658772490401920	8.59	0.4453
29	6687193131310348032	11.20	0.4448
30	5780117135278458112	10.51	0.4446
31	3857760795161492992	10.32	0.4433
32	3405317601488435328	10.27	0.4410
33	5484461722739849984	9.42	0.4406
34	6583876554835641088	4.42	0.4405
35	6140741109344938368	10.20	0.4404
36	3393159923462202624	7.76	0.4404
37	3694036023363825152	8.24	0.4401
38	3618875671935767936	10.92	0.4401
39	6033793262610025984	7.65	0.4314
40	2600019239306534528	9.16	0.4310
41	2715831692413101056	8.16	0.4301
42	6147111267399070464	10.34	0.4301
43	6379802489437427840	9.64	0.4297
44	6456328640466199424	10.45	0.4292
45	6144956155890048256	9.97	0.4289
46	4823564043700716672	8.10	0.4282
47	5722776126416465792	10.66	0.4274
48	5582111338970433664	11.23	0.4265
49	5451886835341219712	11.44	0.4260
50	6656031455656664192	8.57	0.4257

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