

The Milky Way’s enrichment across five nucleosynthetic channels from 4 million LAMOST spectra

ADAM WHEELER,¹ MELISSA NESS,^{1,2} SVEN BUDER,^{3,4,5} AND JEFFREY D. SIMPSON⁶

¹*Department of Astronomy, Columbia University, Pupin Physics Laboratories, New York, NY 10027, USA*

²*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*

³*Max Planck Institute for Astronomy (MPIA), Königstuhl 17, 69117 Heidelberg, Germany*

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

⁵*ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Canberra, ACT 2611, Australia*

⁶*School of Physics, UNSW, Sydney, NSW 2052, Australia*

ABSTRACT

Large stellar surveys are revealing the chemodynamical structure of the Galaxy across a vast spatial extent. However, the many millions of low-resolution spectra observed to date have only been partially exploited. We employ The Cannon, a data-driven approach to calculating abundances, to derive individual element abundances from low-resolution ($R = 1800$) LAMOST spectra, using the GALAH survey as our reference. We deliver abundance measurements for five (dwarfs) or six (giants) of O, Eu, mean α , Sc, mean s -process, Mg, Al, Mn, and Ba, which represent five different nucleosynthetic channels, for 3.9 million stars, to a precision of 0.05 - 0.23 dex. Using wide binary pairs, we demonstrate that our measurements provide additional chemodynamical discriminating power, beyond metallicity alone. We show the coverage of our catalogue with radial, azimuthal and dynamical abundance maps, and examine the neutron capture element chemodynamics across the disk and halo, which indicate different origins for the *in-situ* compared to the accreted halo population. LAMOST has near-complete Gaia coverage, which provides an unprecedented perspective on chemistry across the Milky Way.

Keywords: methods, spectra, abundances

1. INTRODUCTION

Large stellar surveys such as *Gaia* (Collaboration et al. 2016), APOGEE (Majewski et al. 2017; Holtzman et al. 2018; García Pérez et al. 2016), GALAH (De Silva et al. 2015), Gaia-ESO (Gilmore et al. 2012), RAVE (Steinmetz et al. 2006), LAMOST (Newberg 2012; Zhao 2012) and SEGUE (Yanny et al. 2009) are providing the data to empirically characterize the Milky Way disk and infer the primary drivers of its formation and evolution (e.g. Hayden et al. 2015; Nidever et al. 2014; Helmi et al. 2018; Frankel et al. 2018; Blancato et al. 2019, others).

Detailed abundances are one of the primary measurements made from stellar spectra. Their determination is a primary motivation for medium- and high-resolution spectroscopic surveys for several reasons: they provide effective chemical fingerprints of stars (Freeman & Bland-Hawthorn 2002), link directly to the environment in which they were born (e.g. Krumholz et al. 2019), de-

scribe the chemical diversity of the disk (e.g. Weinberg et al. 2019) and the chemical pathways of enrichment (e.g. Rybizki et al. 2017). Combined with stellar kinematics, abundance measurements are core to the pursuit of Galactic archaeology.

Conventionally, detailed abundances have been derived from medium- and high- resolution stellar spectra (i.e. APOGEE $R = 22,500$, GALAH $R = 28,000$, and RAVE $R = 8000$), where element abundances are determined from sets of selected un-blended absorption features by fitting to stellar models. Until recently, the measurements for low-resolution spectra, like LAMOST and SEGUE, were typically limited to stellar parameters and α -enhancements (T_{eff} , $\log(g)$, $[\text{Fe}/\text{H}]$, $[\alpha/\text{Fe}]$) (e.g. Lee et al. 2011). More recently, data-driven approaches and full spectral modeling have opened up the opportunity to obtain individual abundances from low-resolution spectra, where lines are blended. Ting et al. (2018) have shown that oxygen abundances can be inferred from spectra in wavelength regions containing no atomic oxygen lines through the features of species in the CNO atomic-molecular network. Indirectly inferred

abundances also have a long empirical history. The Ca II triplet, for example, is an often used metallicity index (Armandroff & Zinn 1988; see Vásquez et al. 2015 for a recent calibration). To date, with the exception of Xiang et al. (2019), the efforts to extract individual abundances from LAMOST have largely focused on a few elements, namely an integrated α -element abundance and the elements C and N, which are particularly important, as these elements indicate age (e.g. Li et al. 2016; Xiang et al. 2017; Ho et al. 2017b,a; Zhang et al. 2019).

In this work, we use data-driven approach to label low-resolution LAMOST spectra with several abundances. LAMOST is one of the largest stellar surveys to date, with over 5×10^6 publicly available spectra, at $R = 1800$. The survey has extensive coverage of the Milky Way’s disk, halo and, in particular, the outer disk, the detailed chemodynamics of which are largely unexplored. Specifically, we employ The Cannon (Ness et al. 2015), a model characterized in large part by its simplicity, to derive individual abundances from LAMOST. Other data-driven methods include The Payne (Ting et al. 2018a), which like The Cannon works by explicitly modelling spectra as a function of labels (stellar parameters and abundances), and Leung & Bovy (2018), which use neural networks to learn labels directly from spectra without explicit inference. Xiang et al. (2019) recently released a catalog of 16 abundances (C, N, O, Na, Mg, Al, Si, Ca, Ti, Cr, Mn, Fe, Co, Ni, Cu, and Ba) for LAMOST DR 5 using a neural-net-based model calibrated by both labelled spectra (using overlap with both LAMOST and APOGEE) and physical modelling. This work has many common aspects with our own, but is different in detail.

Our approach requires reference objects, stars with high-quality spectra and precise labels (stellar parameters and abundances), that are representative of the survey objects. They are used to calibrate a model that produces synthetic spectra from stellar labels. This model is then used to estimate labels for the full set of survey stars, in our case, the LAMOST catalog. Both the APOGEE and GALAH surveys have stars in common with LAMOST which can serve as possible reference objects. APOGEE provides higher precision abundance measurements compared to GALAH. This enables, for example, the clear disambiguation of the the low- and high- α sequences, as seen in the radial maps of Hayden et al. (2015) and Nidever et al. (2014). However, the dimensionality of the abundance space measured by APOGEE is low (Ness et al. 2018; Price-Jones & Bovy 2018; Ness et al. 2019) (although note weak lines of neutron capture elements have been identified in this regime

(Cunha et al. 2017; Hasselquist et al. 2016)). GALAH, on the other hand, provides abundance measurements across a more extensive set of nucleosynthetic channels, including the neutron-capture (r and s) processes. The neutron-capture element enhancements have been previously explored only through boutique analyses of small samples of stars observed at high resolution (e.g. Bensby et al. 2014; Spina et al. 2018) and at the solar neighbourhood, to which GALAH is largely confined (e.g. Buder et al. 2019; Schönrich & Weinberg 2019). GALAH also provides abundances for main-sequence stars, allowing us to extend our modelling to that regime.

We want to explore the promise of the largest number of element abundance families as possible, so we took the roughly 10,000 stars in common between GALAH and LAMOST to build a model using the LAMOST spectra and GALAH stellar parameters and abundances. While the GALAH labels are less precise than those from APOGEE and thus yield less precise LAMOST labels, the LAMOST catalog is large enough that combining stars will enable very precise mean measurements of abundances on a population basis (e.g. Ness et al. 2019; Blancato et al. 2019). Using GALAH as a source for our input labels allows us to propagate measurements of r -process and s -process abundances to the outer disk and halo.

In deriving a set of individual abundances for LAMOST, this work complements the LAMOST catalogue, which provides stellar parameters and bulk metallicity (a term used interchangeably with $[\text{Fe}/\text{H}]$ in this work) only. We deliver detailed abundances for elements from five nucleosynthetic families: *light* elements, which are dispersed by asymptotic giant branch (AGB) stars and core-collapse supernovae (CCSN), and whose atmospheric abundances can change due to dredge-up; α -elements, which are dispersed primarily by CCSN; *iron-peak* elements, which are dispersed by both CCSN and type Ia supernovae (SNIa); *odd-Z* elements, which are dispersed by both CCSN and SNIa and expected to display similar trends to the α elements; s -process elements, which are though to be produced and dispersed in AGB stars; and r -process elements, which are produced in extremely neutron-rich environments. It is not clear at the present whether neutron-star mergers are the primary site of the r -process, or if other sites make appreciable contributions (e.g. Arnould et al. 2007; Côté et al. 2018; Siegel et al. 2019). For each star, we deliver six (for dwarfs) or seven (for giants) abundances of O (light), Eu (r -process), mean α , Sc (iron-peak), mean s -process, Mg (α), Al (odd Z), Mn (iron-peak), and Ba (s -process). Having derived these abundances, we demonstrate the scientific value of multi-element abundances

of large numbers of stars. We do this using pairs of stars across the disk and halo, examining the abundance similarity of wide binaries, that have been identified by their kinematics alone. We also map the chemodynamical abundance structure of the disk and halo, making links to signatures of evolution such as radial migration and Galaxy assembly.

In §2 we describe the GALAH and LAMOST data and the quality cuts we applied. §3 provides a brief overview of The Cannon. In §4 we discuss model checks and evaluate the error of our label estimates. §5 discusses our public catalog and key scientific results, and §6 discusses their implications.

2. DATA

Our data comprises the R=1800 DR 4 v2 LAMOST spectra, the R=28,000 DR 2.1 GALAH spectra and stellar parameter and abundance labels (Buder et al. 2018), as well as the *Gaia* proper motion and parallax measurements for our stars. From GALAH we use T_{eff} , $\log(g)$, v_{mic} , and $[\text{Fe}/\text{H}]$, along with abundances with respect to Fe, of O, Si, Ca, Ti, Eu, Sc, Y, Mg, Al, Mn, and Ba. Figure 1 shows the Galactic footprints of GALAH and LAMOST. A portion of each survey’s spectrum for a typical training set star is shown in Figure 2.

2.1. Quality cuts and data cleaning

One of the formal assumptions of The Cannon is that the training labels are known exactly, so constructing a high-fidelity training set is crucial. To build our training set, we first determined the set of stars in common between GALAH and LAMOST. We performed a $1''$ sky match between GALAH DR 2.1 and LAMOST DR 4 v2 to identify these reference object candidates, of which there were roughly ten thousand. We then removed all stars from the potential training set with signal to noise ratio (S/N) less than 30 in either the LAMOST z band (`snrz`) or the GALAH blue channel (`snr_c1`). We also removed any star for which `chi2_cannon` (a column in the GALAH catalog, not a product of our analysis) was greater than 4, which indicates that the best fit spectral model is a poor fit, and any star for which `flag_cannon` was nonzero, which can indicate a variety of problems with abundance determination. These cuts removed roughly half of the stars from consideration. We found that cutting on the reported GALAH label errors did not improve our performance against the validation set.

To further exclude low-quality measurements from our training set, we therefore generated and evaluated the fit of the best-fit Cannon model spectrum for each reference stellar spectrum, for every element in the GALAH

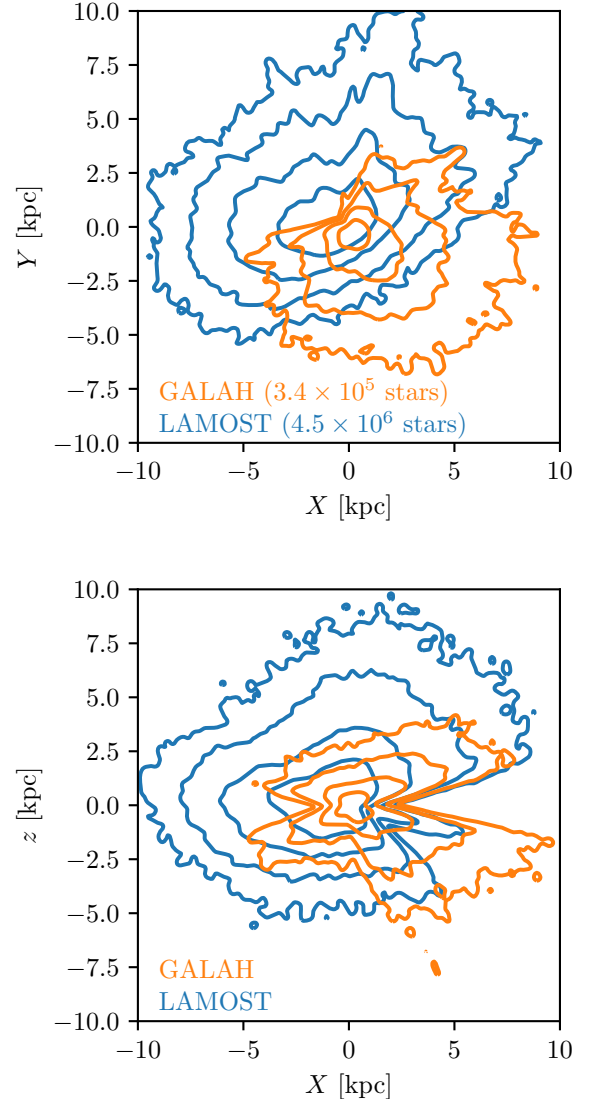


Figure 1. Face on (top) and edge on (bottom) contours in surface density ($5 \times 10^{4,3,2,1} \text{ kpc}^{-2}$) shown in heliocentric Galactic coordinates for GALAH and LAMOST, which probes much farther into the outer disk and halo.

catalog. The GALAH pipeline uses separate Cannon models for each elemental abundance in order to restrict each model to the wavelengths of unblended lines. Each model has different best-fit parameters, which we were not able to retrieve. They are, however, within the errors of the mean reported stellar parameters for each star (Buder et al. 2018). For the stellar parameter labels, we used the values in the GALAH DR2.1 catalog, along with A_K values calculated with the Rayleigh-Jeans color excess (RJCE) method (Majewski et al. 2011) applied to ALLWISE (Wright et al. 2010; Mainzer et al. 2011) and 2MASS (Skrutskie et al. 2006) broadband photometry,




Figure 2. A comparison of the GALAH and LAMOST spectra for the same star, 2MASS 00010184+0407201, *Gaia* DR2 2740354684364096000. The top panels show part of the star’s GALAH spectrum (S/N = 65) and Cannon model, while the bottom two show the same for LAMOST (S/N = 179). On the left, see the large H_{β} line and surrounding features, on the right, see the fit around a known Mn feature (highlighted in red on the left). Both the measured spectra and Cannon models are shown with their $1\text{-}\sigma$ error.

as was done for the GALAH models. We calculated χ^2 between the best-fit GALAH model and the observed GALAH spectrum for every star in our training set in the region of the strongest lines of each element (the `chi2_cannon` flag pertains to the global fit). Appendix C lists the wavelength regions used, which are the same windows used in the GALAH pipeline. The distribution of χ^2 values for some elements peaked lower than expected from nominal measurement error alone by a factor of 2-3, meaning that a cut on some multiple of χ^2/dof was not theoretically justified. We removed all stars with χ^2 values above the 85th percentile, for any of its abundances. This led to a significant improvement in our cross-validation results, as discussed in our methods, on the order of 15%-40% percent). Using the 75th percentile, as a more conservative cut, gave us no improvement in cross-validation tests. These cuts leave 1722 stars in the training set. We do not exclude stars flagged in the GALAH based on `flag_x_fe` because we performed our own per-abundance χ^2 cut and because removing stars where the GALAH model may be extrapolating reduces the size of our training set too drastically. We emphasize however that The Cannon is likely to extrapolate well (in a well understood way using the simple polynomial model we employ) for many abundances.

2.2. Reference Objects

After the quality cuts described in Section 2.1, we were left with a training set that spans the Kiel diagram (Figure 3). We modeled giants and dwarfs separately, with




Figure 3. Our 1722 training objects in the LAMOST Kiel diagram space colored by (GALAH) $[\text{Fe}/\text{H}]$. Since there are more metal-poor stars in the giant training set, our giant model is unbiased down to lower metallicity.

the division between models given in terms of LAMOST $\log g$ and T_{eff} by

$$\log g = \begin{cases} 4.18 & T_{\text{eff}} < 5200\text{K} \\ (-6 \times 10^{-4})T_{\text{eff}}/\text{K} + 7.3 & T_{\text{eff}} \geq 5200\text{K}. \end{cases} \quad (1)$$

The split gives us 532 giants and 1190 dwarfs as our reference objects. We decided which elements to infer for each model by balancing our ability to recover each abundance in cross-validation (§4) with the objective of having roughly a single element per nucleosynthetic channel. For both models, we include T_{eff} , $\log g$, v_{mic} , $[\text{Fe}/\text{H}]$, $[\text{O}/\text{Fe}]$, and $[\text{Eu}/\text{Fe}]$ as labels. For the dwarfs, we also used iron-relative abundances in: error-weighted mean α (from Mg, Si, Ca, and Ti), Sc, and mean s-process (from Ba and Y). For the giants, we also used Mg (α), Al (odd- Z), Mn (iron-peak), and Ba (s -process).

We only used mean abundances in the same nucleosynthetic family if they appeared strongly correlated in the training set. We tried using dereddened *Gaia* G band magnitude instead of $\log g$, which would allow us to apply a prior at test time, but we found that this did not improve our results in practice. The model had trouble predicting extinction, partially because our training sets do not include any high-extinction stars. Including extinction as a label did not improve our ability to predict any of the abundances, so we opted not to.

For subsequent analysis, we cross-matched with *Gaia* by taking the source within 1 arcsecond of the LAMOST star with the lowest G -band magnitude. Throughout this paper, we use the parallactic distance estimates from Bailer-Jones et al. (2018).

3. THE MODEL

The following is a brief description of The Cannon (see Ness et al. 2015, for a more extended discussion.). For this work we build a JULIA-based implementation of The Cannon, which is fully documented and available at [this URL](#)¹ and via the JULIA package manager. The source code uses the same nomenclature as the description here, and allows for optional masking of labels (self-consistent training with the model constrained so that each label is only “on” at specified wavelengths).

For each star, n , we take the flux value in the spectral pixel with wavelength λ to be $F_{n\lambda}$ and its (Gaussian, independent) measurement uncertainty to be $\sigma_{n\lambda}$. To prepare the spectra for The Cannon we first redshift-corrected the spectra using the z value provided in the LAMOST data table and interpolated each star to a common wavelength grid. We then continuum-normalized the spectra by dividing out the continuum, approximated by smoothing the spectra with a Gaussian kernel with a 50 Å standard deviation, truncated at 150 Å from the center, in the same manner as Ho et al. (2017b,a). This normalised flux is then near unity in

the absence of emission or absorption features. For each reference star, we also define ℓ_n be the vector containing its physical parameters and abundances (its *labels*). These are the quantities we ultimately wish to infer for the rest of the LAMOST spectra, at test time.

In our case, our labels ℓ for our n reference stars are:

$$\ell_n = \langle T_{\text{eff}}, \log(g), v_{\text{mic}}, [\text{Fe}/\text{H}], [X_1/\text{Fe}], \dots, [X_N/\text{Fe}] \rangle \quad (2)$$

where X_1, \dots, X_N are the elements whose abundances we wish to determine. It is good practice for both numerical stability and model flexibility to express all labels in units such that they are distributed around zero and have similar magnitudes. We do this by subtracting from each label its (training set) mean and dividing it by its (training set) dispersion. This transformation is then undone after the inference has taken place.

In numerous published uses of The Cannon (including this one), the flux in each pixel is described by a 2nd degree polynomial of the elements of the label vector whose coefficients, θ , are determined by a training set of spectra for which, ideally, both accurate and precise labels are available.

For a given spectral pixel and star, we then have our spectral flux, F , for our n reference objects at each wavelength, λ defined as:

$$\begin{aligned} F_{n\lambda} &= \theta_\lambda^0 && \text{(constant term)} \\ &+ \theta_\lambda^{T_{\text{eff}}} T_{\text{eff}} + \dots + \theta_\lambda^{X_N} [X_N/\text{Fe}] && \text{(linear terms)} \\ &+ \theta_\lambda^{T_{\text{eff}}^2} T_{\text{eff}}^2 + \dots + \theta_\lambda^{X_N^2} ([X_N/\text{Fe}])^2 && \text{(squared terms)} \\ &+ \theta_\lambda^{T_{\text{eff}} \log(g)} T_{\text{eff}} \log(g) + \dots \\ &+ \theta_\lambda^{X_N X_{N-1}} [X_N/\text{Fe}] [X_{N-1}/\text{Fe}] && \text{(cross-terms)} \\ &+ \text{error.} \end{aligned}$$

To specify an error model, we can write the above as a likelihood function

$$F_{n\lambda} | \ell_n, \theta_\lambda, s_\lambda \sim \mathcal{N}(\boldsymbol{\eta}(\ell_n) \cdot \boldsymbol{\theta}_\lambda, \sigma_{n\lambda}^2 + s_\lambda^2) \quad (3)$$

where \mathcal{N} is the normal distribution, s_λ is model uncertainty (either inherent stochasticity or physics that hasn’t been captured by the model) at wavelength λ , $\boldsymbol{\theta}_\lambda$ is the vector of coefficients describing how the flux at λ varies with label value, and $\boldsymbol{\eta}$, the quadratic expansion (called the vectorizing function in Casey et al. (2016)), maps from labels to every 0th, 1st, and 2nd order combination of components of the label vector,

$$\boldsymbol{\eta}(\ell_n) = \langle 1, T_{\text{eff}}, \dots, [X_N/\text{Fe}], T_{\text{eff}}^2, T_{\text{eff}} \log(g), \dots \rangle. \quad (4)$$

¹ github.com/ajwheeler/TheCannon.jl




Figure 4. The Cannon likelihood as a probabilistic graphical model. During training, the latent variables in the “stars” panel (ℓ_n) are fixed using stars for which labels are known, in our case from GALAH. When inferring stellar labels, the latent variables in the “pixels” panel (θ_λ and s_λ) are fixed to their point estimates from training and the maximum likelihood estimates for all ℓ_n are calculated.

If ℓ_n is a vector of length N , $\boldsymbol{\eta}(\ell_n)$ is a vector of length $(N^2 + 3N)/2$. A more flexible model could be constructed by replacing $\boldsymbol{\eta}$ to an expansion with higher order terms, or to other combinations of labels. However, quadratic models have been shown to be sufficient in practice, and the combinatoric increase in model parameters that would be necessary for higher a higher order polynomial is undesirable.

Figure 4 shows the likelihood function as a probabilistic graphical model, which depicts the relationships between observed and latent quantities. Ideally, the full joint distribution over training data and output labels would be sampled from directly (with e.g. Markov-chain Monte-Carlo), but such an approach is not computationally feasible. Instead the problem is divided into a training step, in which a point estimate of each θ_λ and s_λ is estimated from the training set, and an inference step, in which the labels of each star are estimated.

During the training step, each θ_λ and s_λ is jointly fixed to its maximum-likelihood estimate (MLE), given the labeled spectra in the training set. For fixed s_λ , the model is linear with fixed Gaussian error in θ_λ , so its MLE, $\hat{\theta}_{|s_\lambda}$, can be calculated analytically. Finding \hat{s}_λ is then a matter of numerically maximizing the log-likelihood,

$$\log \mathcal{L}(s_\lambda) = - \sum_n \frac{1}{2} \left(\frac{(\boldsymbol{\eta}(\ell_n) \cdot \hat{\boldsymbol{\theta}}_\lambda|_{s_\lambda} - F_{n\lambda})^2}{\sigma_{n\lambda}^2 + s_\lambda^2} + \ln(\sigma_{n\lambda}^2 + s_\lambda^2) \right) + \text{const}, \quad (5)$$

in one dimension. During the inference step, the MLE of ℓ_n is calculated with s_λ and θ_λ fixed to their point estimates. There is no trick to get us out of multivariate optimization here, since $\boldsymbol{\eta}$ is nonlinear.

4. MODEL EVALUATION

We use 12-fold cross-validation (CV) in order to verify that the model is able to recover stellar labels. We partition the reference objects into twelve random subsets, then predict the labels of each subset using the other eleven as training data. This gives us a prediction for each reference star that has not leveraged its GALAH labels. Figures 5 and 6 show CV performance for the giants and dwarfs respectively, along with the scatter, bias, and correlation coefficient for each label. Our CV-assessed abundance precision ranges from 0.05 to 0.23 dex for dwarfs and 0.07 to 0.22 dex for giants. Figure 7 shows the precision (twice the scatter in Figures 5 and 6) of each of our abundances relative to the range over which the model does not extrapolate. This quantity is often what is relevant when comparing the labels of different stars, rather than characterizing a single star. Examination of the labels inferred for spectra from repeat observation of the same star show differences consistent with CV-precision.

We also use CV to identify the thresholds beyond which our model fails to extrapolate correctly due to lack of training data. We say that the model is extrapolating when it under- or over-predicts the label being considered 90% of the time in CV. The specific calculation is as follows: For a given label, l (e.g. $l = T_{\text{eff}}$), we approximate $p(l_{\text{true}}, l_{\text{inferred}})$ with a kernel-density estimate (KDE) with bandwidth chosen by Silverman’s rule (Silverman 1986), then use this approximate distribution to find the values of l_{inferred} at which $p(l_{\text{true}}|l_{\text{inferred}})$ excludes l_{inferred} at the 90% level. These boundaries are shown as horizontal lines in Figures 5, and 6. Stars that fall beyond these boundaries are flagged in our catalog.

4.1. Model interpretability

The Cannon is simple enough that its parameters are open to direct interpretation. This sets it apart from more complex modeling approaches such as, e.g. neural networks.

It is clear, for example, that our model learns T_{eff} in large part from the Balmer series, as this is where $\theta_{T_{\text{eff}}}$ becomes large. By examining model scatter, s_λ as a function of wavelength, we can tell that our model is less precise in the regions of CN bands, at the beginning of the spectral region of LAMOST. In some wavelength regions, s_λ drops to 0, likely because continuum-normalization introduces small correlations between nearby pixels that are not accounted for by the model. It is also apparent that the model is leveraging the whole spectrum to predict abundances, rather than strong lines only. We performed tests by isolating only regions where individual abundance features are present in the spectra, and forced the model’s coeffi-

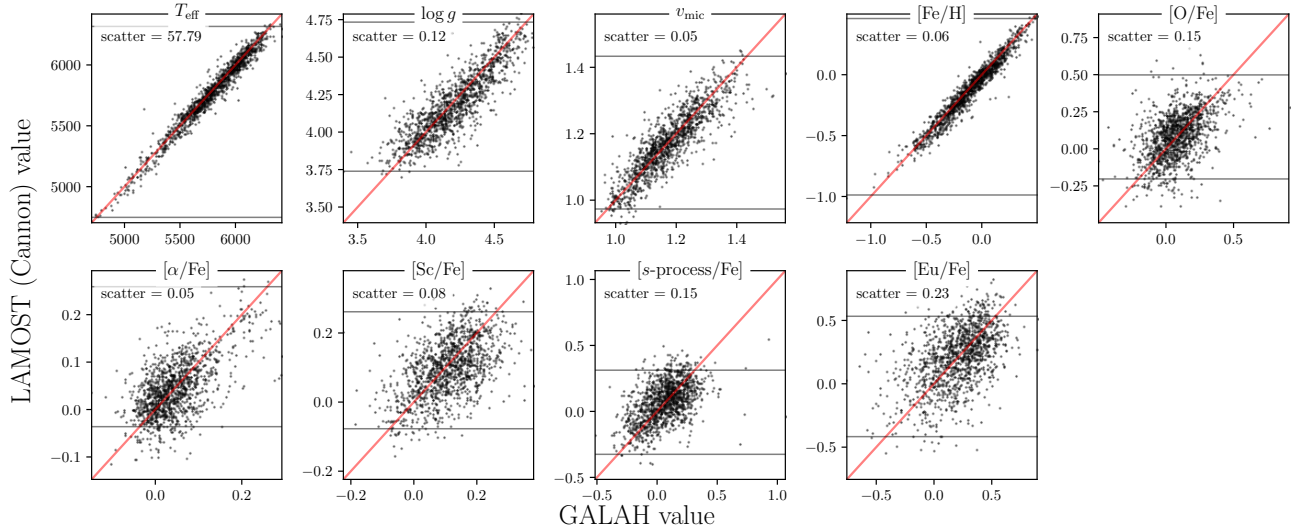


Figure 5. Cross-validation recovery of training set labels for the dwarf model. Horizontal lines show boundaries beyond which the model fails to extrapolate.

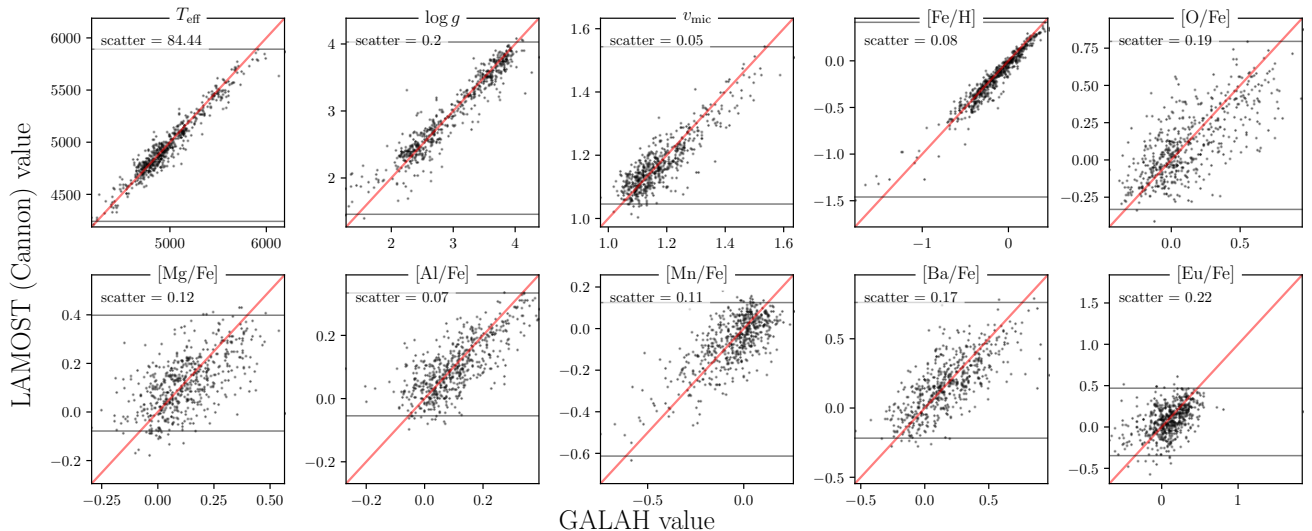


Figure 6. Analogous to Figure 5. Cross-validation recovery of training set labels for the giant model. Horizontal lines show boundaries beyond which the model fails to extrapolate. While the scatter in $[\text{Fe}/\text{H}]$ is higher at lower values, the model appears to be nearly unbiased down to $[\text{Fe}/\text{H}] \approx -1.5$.

cients to be zero at training time outside of these regions. This is the approach via which GALAH determined their abundance labels, using so called abundance windows. For the LAMOST spectra, this approach of using windows fails to recover abundance ratios in cross-validation. Section 4 discusses some implications of this fact, and our scatter and linear coefficients are plotted as a function of wavelength in Appendix A.

5. RESULTS

5.1. Catalogue

We produce a catalog of stellar parameters and individual abundances for 4,541,883 observations of 3,744,284 stars across the Kiel diagram (Figure 8), which is available at zenodo.com/xyz. We combine observations of the same star by reporting (z -band) S/N -weighted averages of their labels. Along with our inferred stellar parameters and abundances, we provide the LAMOST and *Gaia* identifiers for each star, as well as its Galactic position, radial velocity and estimated actions. We also provide windowed and whole-spectrum χ^2 values and flags to tell when the model is extrapolating. For each star, we calculated approximate actions




Figure 7. The precision of each of our abundances relative to the range over which the model is approximately unbiased. We generally infer abundance ratios with precision at the 30-50% level. Smaller values indicate higher relative precision of that abundance and presumably higher discriminating power between stars.

with `galpy` (Bovy 2015) using the Stäckel fudge (Binney 2012; Bovy & Rix 2013), applied using *Gaia* distances (Bailer-Jones et al. 2018) and proper motions, and LAMOST radial velocities. We assumed that the Sun sits at $X = 8$ kpc, $z = 0.25$ kpc (Jurić et al. 2008) and is moving with $v_X = 11.1$ km/s, $v_Y = -232.24$ km/s, $v_z = 7.25$ km/s (Schoenrich & Binney 2009). Table 1 provides the full catalog schema.

Here we highlight several caveats to the use of this data:

- We allow our model to take advantage of the full information content of the spectrum. It therefore not only learns from the most fundamental features of each label, but from correlated features (as we can identify using our model coefficients, which is an advantage of a simple interpretable model). Examination of the model’s coefficients reveals that the whole spectrum is leveraged in order to predict each abundance. Our CV tests show that our model works: it performs well with no hyperparameter tuning. Our individual abundances trends, including Eu, are in accordance with boutique high resolution studies of this element in the Galactic disk (Overbeek et al. 2016; da Silva et al. 2016). Our analysis of wide binaries in the solar neighborhood (El-Badry et al. 2019) is indicative of the additional discriminating power beyond an overall metallicity these abundances provide (see Section 5.2). However, the abundances are not being measured directly. The fidelity of our predicted labels relies on our reference objects (confined to the solar neighborhood) being representative of the survey data. For

this reason, it’s reasonable to expect our abundances to be more accurate for disk stars than halo stars. In order to identify many case where the model fails to generalize from the training set, we provide χ^2 values calculated across the whole spectrum and individually in narrow windows centered on strong lines, for each element. If the best-fit spectrum is a poor fit around known features of a given element, it is likely highly enriched or depleted in that element. In fact, this approach is a good way to find such stars with anomalous abundance patterns. Indeed, Casey et al. (2019), Kemp et al. (2018), and Norfolk et al. (2019) have used the departure of a basic stellar parameter model generated with The Cannon, from the spectra, to find LAMOST stars that are enhanced in Li, K, and Ba and Sr, respectively.

- A caveat which is general to data-driven methods is that the model will not necessarily extrapolate correctly outside the parameter space spanned by the training set. We provide flags to indicate when individual abundances are in the regime where they may be incorrectly extrapolated, as well as a flag indicating when T_{eff} , $\log(g)$, v_{mic} , or $[\text{Fe}/\text{H}]$ may be incorrectly extrapolated (Table 1). We determine when our model is extrapolating as described previously, in Section 4.
- While error estimates for each abundance ratio are desirable, producing accurate ones would be prohibitively costly with our current inference infrastructure. We advise the user to use our CV-assessed error and caution them to be aware that treating our abundance measurements as homoscedastic is a necessary compromise.
- Examination of open clusters in our catalog reveals that our inferred abundance ratios for dwarf stars are subject to strong systematics as a function of T_{eff} (see Figure 17, which shows $[\alpha/\text{Fe}]$ vs T_{eff} for all dwarfs in Praesepe for which we have abundances). The dwarf stars are not used in the majority our analysis in this paper, as we do not correct for these systematics. There are astrophysical reasons for abundance trends with T_{eff} and $\log(g)$, such as atomic diffusion (Gao et al. 2018; Souto et al. 2019), although such strong trends are not expected in all abundances, as we find in our catalog. Similar systematics appear to be present in the official LAMOST $[\text{Fe}/\text{H}]$ values for dwarfs, which suggests that these trends are not introduced by our label transfer, but are inherited from our *ab-initio* stellar models via our training set. Our analysis of wide binaries (see Section 5.2) suggests that this is a precision bottleneck for our model.




Figure 8. A portion of the best-fit model spectra (black) and real data (colored), both with $1\text{-}\sigma$ uncertainties, for 6 randomly-chosen stars in our catalog. Though simple, the model is flexible enough to fit the data across the Kiel diagram.

There are no obvious systematics in the red giant stars in our catalogue, save for $[\text{Ba}/\text{Fe}]$, discussed below. However, LAMOST does not contain enough red giants in known open clusters or wide binaries to determine the presence and magnitude of any systematics conclusively. The systematic trends we see in dwarf abundances could be “calibrated out” using the nearly 3000 stars in LAMOST DR4 open clusters (with two or more targets), (Cantat-Gaudin et al. 2018), and 142 known wide binaries (El-Badry et al. 2019). Correcting for these systematics in the dwarf population is beyond the scope of our analysis. Despite this systematic effect, our abundances for dwarf stars are still useful for conducting analyses in restricted temperature ranges. See, for example, our examination of abundances of wide binaries in Section 5.2. When examining the abundance trends across the disk, we exclude the dwarf stars, and focus on the $\approx 1 \times 10^6$ red giant stars in our catalogue. These giants span a vast spatial extent, and alone demonstrate the scientific potential of the distribution of stellar abundance data across the Galaxy.

Unless otherwise stated, in the sections below we employ stars in our catalog for which for which chi2 is less than 7000. Other cuts were not found to have an effect on the results presented below.

5.2. Detailed abundances of wide binaries

El-Badry et al. (2019) (hereafter EB19) used *Gaia* to identify wide binaries in the solar neighborhood and examined their properties as a function of $[\text{Fe}/\text{H}]$. We subsequently are able to examine the detailed abundances of wide binaries from EB19, for which both stars are LAMOST dwarfs. **MKN: this raises the question**




Figure 9. **top:** Error-weighted chemical distances for EB19 wide binaries and for random field pairs selected with the same cuts and chosen to have the same $\Delta[\text{Fe}/\text{H}]$ distribution as the EB19 sample. The wide binaries are more chemically similar than implied by their similarity in bulk metallicity alone. **bottom:** Differences in $[\text{O}/\text{Fe}]$ for wide binaries and field pairs.

- why not giants?, maybe say, for which both stars in the binary are in LAMOST, and state these are all dwarf stars Because of the strong systematic trends with T_{eff} that are present in our dwarf

Table 1. Catalog schema. Here, x and X stand for each of the chemical symbols for the elements whose abundances being measured. Our Galactic coordinate system is right-handed. We also make available the table of per-observation labels.

column name	type	unit	description
<code>source_id</code>	integer		<i>Gaia</i> DR2 source id
<code>designation</code>	string		LAMOST unique star identifier
<code>giantmodel</code>	boolean		true if labels were estimated with giant model
<code>teff</code>	float	K	T_{eff}
<code>logg</code>	float		$\log(g)$
<code>vmic</code>	float		v_{mic}
<code>keil_extrap</code>	boolean		true if T_{eff} or $\log(g)$ (the axes of the Keil diagram) are in regime where model fails to extrapolate for any observation
<code>chi2</code>	float		whole-spectrum χ^2
<code>fe_h</code>	float		[Fe/H]
<code>fe_h_extrap</code>	boolean		true if [Fe/H] value is in regime where model fails to extrapolate for any observation
<code>x_fe</code>	float		[X/Fe]
<code>chi2_x_fe</code>	float		χ^2 calculated in windows around strong lines of X
<code>x_fe_extrap</code>	boolean		true if [X/Fe] value is in regime where model fails to extrapolate for any observation
<code>snrz</code>	float		LAMOST z -band S/N
<code>ra</code>	float	deg	right-ascension
<code>dec</code>	float	deg	declination
<code>R</code>	float	kpc	R , in Galactic cylindrical coordinates
<code>phi</code>	float	rad	ϕ , in Galactic cylindrical coordinates
<code>z</code>	float	kpc	z , in Galactic cylindrical coordinates
<code>vR</code>	float	km/s	R -velocity
<code>vT</code>	float	km/s	ϕ -velocity
<code>vz</code>	float	km/s	z -velocity
<code>JR</code>	float	kpc km/s	radial action
<code>Jphi</code>	float	kpc km/s	angular momentum
<code>Jz</code>	float	kpc km/s	vertical action

abundances, we constrain our analysis to wide binaries with $\Delta T_{\text{eff}} < 250$ K.

To confirm their additional chemodynamical discriminating power of our inferred abundances, we examined the abundance similarity of wide binaries compared to a reference sample non-binary pairs. We constructed a set of random pairs of field stars, where each pair has the same metallicity as the binary pair. The reference stars also conform to the quality cuts made in EB19 with $T_{\text{eff}} < 250$ [K]. We used rejection sampling to ensure that they had as closely as possible the same $\Delta[\text{Fe}/\text{H}]$ distribution as the EB19 sample. By comparing the abundance distribution of the random field pairs with the wide binaries, we can characterize the amount of information contained in our detailed abundances *above and beyond* that contained by the bulk

metallicity, [Fe/H]. To capture the difference in chemistry between stars we use precision-scaled euclidean distance,

$$\sqrt{\sum_i \left(\frac{\Delta[X_i/\text{Fe}]}{\sigma_i} \right)^2}, \quad (6)$$

where the X_i 's are the elements measured, and the σ_i 's are their CV-assessed uncertainty. Figure 9 shows the distribution of these chemical distances for both the wide binaries and the field pairs with the same $\Delta[\text{Fe}/\text{H}]$ distribution. The difference between these distributions shows that detailed chemical abundances provide additional information about star's birth sites. Each abundance included pushes the chemical difference distribution of the binaries and random pairs further apart. The wide binaries peak at a smaller chemical distance of 2.2

than the reference pairs, which peak at 0.8. Thus, the closer abundances of wide binaries compared to field stars, at the same overall metallicity, $[\text{Fe}/\text{H}]$, demonstrates the power of combining both dynamics and detailed abundances to identify and characterise stars with the same birth site. We did not find that binaries with a larger separation are more chemically different, in contrast with the results of Ramirez et al. (2019).

If a cut in ΔT_{eff} is not made, the systematic error in each abundance becomes similar in magnitude to the dispersion of chemistry in the solar neighborhood (0.1 – 0.5 dex, depending on abundance). Without this ΔT_{eff} cut, and with this subsequent high systematic error, random field pairs and wide binaries appear to have very similar chemical difference distributions. Even more stringent requirements for ΔT_{eff} results in even more distinct chemical distance distributions for the wide binaries, but at the expense of the number of qualifying wide binaries. In fact, the chemical differences we see in the wide binaries are much smaller than the error we get in cross-validation. This suggests that systematic T_{eff} -dependant effects dominate our CV-assessed errors (see Figure 17). If, in future work, if we were able to reduce or eliminate this effect, perhaps by conditioning a model on chemically homogeneous open clusters, we could produce much higher-fidelity detailed abundances. Currently, scientific exploits of our ≈ 3 million dwarf stars should employ narrow ranges of T_{eff} .

Our tests on the measured chemical differences between wide binary stars indicate not only that the detailed chemistry provide evidence of a common birth site. They also show that systematic effects are a large fraction of our error budget—a promising sign that we can do better with low-resolution spectra in the future.

We also similarly investigated the chemical differences for a sample of co-moving pairs in (Kamdar et al. 2019) compared to a reference set of field stars at the same $[\text{Fe}/\text{H}]$ and found that they were also chemically more similar than an equivalent set of field pairs, although less so than the wide binaries. Simpson et al. 2019 used GALAH abundances to determine whether 15 co-moving pairs found in *Gaia* were co-natal; the same approach could be used here.

5.3. Mapping chemistry in the Milky Way

We have one of the largest homogeneous stellar sample of abundances. This sample is ideal for mapping the abundance distribution of the Milky Way across a large spatial extent. First, we map the disk across (R, z) , to characterize the spatial abundance trends in that plane. Similar maps can be created with a different set of abundances using APOGEE but that data set is most con-

centrated to the disk and the inner Galaxy, while the LAMOST giants more extensively span the halo and outer disk. APOGEE data clearly reveals the flaring in intermediate-age populations in the (R, z) plane (e.g. Ness et al. 2016). This is presumably a consequence of radial migration (e.g. Roškar et al. 2008), where by stars increase in scale height as they move outward in the disk (Minchev et al. 2012). Due to the correlations between abundances and ages (Bedell et al. 2018; Feuillet et al. 2018, 2019; Ness et al. 2019), we might expect to also see such flaring in mean-abundance maps, although this is potentially confounded by the metallicity dependence of the age-abundance relationships (Ness et al. 2019). Detailed analyses of the chemodynamical distribution across (R, z) that seek to make any quantitative claims require a careful consideration of the LAMOST selection function. Characterisation of flaring profile of the disk also require stellar ages, as noted by Minchev et al. (2014, 2018). Here, we aim to show the potential of this data for more in-depth analysis, that accounts for the selection function.

Figure 10 shows the (R, z) plane coloured by mean label value for nearly 800,000 giant stars, for abundance ratios of Fe, O, Eu, Mg, Al, Mn, and Ba as well as T_{eff} and $\log(g)$. These maps span $-4 \text{ kpc} < z < 4 \text{ kpc}$ and $7.5 \text{ kpc} < R < 15 \text{ kpc}$. The disk is clearly distinct from the halo—around the solar neighbourhood: at $R = 8 \text{ kpc}$, for example, the halo transition appears as a smooth mean abundance change centered on $|z| \approx 2 \text{ kpc}$. Flaring is seen in the individual elements, particularly for O, Mg, Eu and Ba. All of these elements show different flaring, of varying strength and profile. All abundances increase or decrease monotonically with $|z|$ at fixed R , except for $[\text{Al}/\text{Fe}]$, which increases with $|z|$ until $|z| \approx 2 \text{ kpc}$, beyond which it decreases with $|z|$.

The apparent barium-depleted “cone” centered on the sun is caused by systematic trends in $[\text{Ba}/\text{Fe}]$ as a function of T_{eff} and $\log(g)$, in combination with LAMOST’s selection function. If we plot only the red clump stars as identified by Ting et al. (2018b) (roughly 2×10^5 stars), which exhibit a narrow range of stellar parameters and which have very precise photometric distances, this feature disappears. The shape and morphology of flaring in the elements is preserved when examining the red clump stars only. Finally, we note that both $[\text{Ba}/\text{Fe}]$ and $[\text{O}/\text{Fe}]$ appear to be asymmetrically distributed about the Galaxy’s midplane. This asymmetry in the mean abundance value around the midplane persists in maps of the (R, z) plane made with only red clump stars, suggesting that they are not related to T_{eff} -dependant systematics. As seen in Figure 10, this feature doesn’t correlate clearly with T_{eff} or $\log(g)$; nor does it trace ex-




Figure 10. The (R, z) plane colored by mean label value for nine labels of our 800,000 giant stars, summed over all ϕ . The flaring of the disk can be seen in $[\text{Eu}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$, and $[\text{Mn}/\text{Fe}]$.

inction as traced by dust maps, or mean signal to noise of the stars. We do not rule out the possibility that the midplane asymmetry seen in these elements is caused by selection effects, particularly in light of the fact that these asymmetrical features are stretched along lines of sight.

Figure 11 shows mean-abundance maps in the (X, Y) plane for kinematic thin disk stars ($J_z < 30 \text{ kpc km s}^{-1}$), as well as mean T_{eff} and A_V maps, for comparison. We highlight the apparent azimuthal structure in $[\text{Mn}/\text{Fe}]$ and $[\text{O}/\text{Fe}]$, which isn't easily explained

by spurious correlation with T_{eff} , A_V , S/N , z , or J_z . These are two abundances for which inferred values have lower relative uncertainty (Figure 7). Again, we note that a complete treatment of this data would involve explicit modelling of the LAMOST selection function. Note that the sensitivity of $[\text{Ba}/\text{Fe}]$ to T_{eff} is clearly visible. Like the “cone” in the (R, z) plane, correlation of $[\text{Ba}/\text{Fe}]$ with heliocentric distance disappears in maps including only red clump stars.

Azimuthal trends in abundance are known to exist in Galactic gas, and are often attributed to spiral struc-




Figure 11. Mean abundances of LAMOST giants in the thin disk ($J_z < 30 \text{ kpc km s}^{-1}$), mapped in the (X, Y) plane. T_{eff} and A_V are also mapped for comparison.

ture (e.g. Wenger et al. 2019). Variations in the height of the midplane in combination with LAMOST’s selection function could give rise to azimuthal abundance gradients, but it is not clear why this would be manifest in some abundances and not others unless due to abundance-age correlations, which does not appear to be the case here). Additionally, there is no correlation between the strength of vertical (Figure 10) and azimuthal (Figure 11) gradients, which would be expected if the azimuthal trends were due to midplane variations.

Figure 12 shows mean-abundance maps of LAMOST giants for the thin disk ($J_z < 30 \text{ kpc km s}^{-1}$) in the (R, v_ϕ) plane. By using LAMOST radial velocities, we are able to probe further out into the disk than with the Gaia DR2 RVS sample. In this plane a particularly prominent feature are the ‘ridges’ first reported by Kawata et al. (2018). A number of interpretations have been given for the origin of these ridges, including perturbations introduced by spirals, the bar, an external perturber or a combination of these (e.g. Antoja et al. 2018; Fragkoudi et al. 2019b,a Laporte et al. 2019).

Of particular interest is the longest ridge (outlined by a dashed line). Wheeler et al. 2019 (*in prep*) will discuss its dynamical origin. This ridge is most visible in the $[O/Fe]$ and $[Mn/Fe]$, a feature shown also in cosmological simulations of barred galaxies (Fragkoudi et al. 2019a). These are two elements that display the clearest azimuthal abundance gradients, and which our inferred abundances have the lowest relative uncertainty (Figure 7).

5.4. The disk-halo transition seen in chemistry

The distributions of the neutron capture elements, $[Ba/Fe]$ and $[Eu/Fe]$, vs $[Fe/H]$, colored by v_ϕ , are plotted in spatial bins in Figures 13 and 14. Because of the T_{eff} -dependant systematics in $[Ba/Fe]$, we have only plotted stars with $4800 \text{ K} < T_{\text{eff}} < 5000 \text{ K}$ in Figure 13. Using other temperature ranges doesn’t qualitatively change the plot, but not restricting T_{eff} yields higher dispersion in $[Ba/Fe]$. v_ϕ allows us to clearly distinguish between the disk and halo populations at the scale heights (and vertical actions, J_z) where both are present (primarily the center row in Figures 13 and 14), illuminat-




Figure 12. Mean-abundance maps of thin disk ($J_z < 30 \text{ kpc km s}^{-1}$) LAMOST giants in the (R, v_ϕ) plane. The dashed box outlines a ridge of higher stellar number density, thought to correspond to stars trapped by the outer-Linblad resonance of the bar.

ing the chemical differences between them. Disk stars are prograde across R , and concentrated to low z , with most disk stars having $v_\phi \gtrsim 100 \text{ km/s}$. There are fewer metal-poor disk stars as R increases, with a narrower more metal rich distribution at larger R , seen across the smallest z range. Halo stars hand are seen at larger z and are characterised by their v_ϕ distributed approximately around 0. At $11 \text{ kpc} < R < 13 \text{ kpc}$, most halo stars have $v_\phi \lesssim 100 \text{ km/s}$ with a distribution of $v_\phi = 0 \pm 50 \text{ km/s}$ **MKN: please check, I invented this number - you can check by just getting the mean and std of the top right panel. Let's also check-in with one another about changes I made here.** The halo stars also appear to have increasingly negative velocities in the inner Galaxy. At our intermediate height from the plane, $2 \text{ kpc} < z < 4 \text{ kpc}$, the metal-poor stars, $[\text{Fe}/\text{H}] < -1.0$, are predominantly retrograde at our smallest R range, $3 \text{ kpc} < R < 5 \text{ kpc}$. Cutoffs

of $J_R \approx 100 \text{ km s}^{-1} \text{ kpc}$ and $J_\phi \approx 1500 \text{ km s}^{-1} \text{ kpc}$ also clearly distinguish between disk and halo stars at large R , but are less distinct at small R .

The distribution of (kinematic) halo stars in the $([\text{Ba}/\text{Fe}], [\text{Fe}/\text{H}])$ plane has a transition at $[\text{Fe}/\text{H}] \approx -1$ (most clearly seen in the middle row of Figure 13, $2 \text{ kpc} < |z| < 4 \text{ kpc}$). This metallicity corresponds to the transition between the disk and halo, as well as the approximate boundary between the accreted and (at least one component of the) *in-situ* halo (Bonaca et al. 2017, called “the splash” in Belokurov et al. 2019 and the “heated thick disk” by Di Matteo et al. 2018). At least for barium, the abundance planes at $4 \text{ kpc} < |z| < 6 \text{ kpc}$, suggest that an overlap in the chemical plane of different sequences, perhaps associated with the accreted and *in-situ* halo. Both $[\text{Eu}/\text{Fe}]$ and $[\text{Ba}/\text{Fe}]$ have larger dispersion at high $[\text{Fe}/\text{H}]$, but the sequence of europium abundances varies less across z .




Figure 13. The $([\text{Ba}/\text{Fe}]-[\text{Fe}/\text{H}])$ plane, colored by J_R , and plotted in spatial bins in the Galaxy, with (up to) 700 randomly selected stars plotted in each bin. Because of the T_{eff} -dependant systematics in our inferred $[\text{Ba}/\text{Fe}]$ values, we have only plotted stars with $4800 \text{ K} < T_{\text{eff}} < 5000 \text{ K}$. With increasing $|z|$, halo stars (as distinguished by their lower ϕ -velocities) become more dominant.




Figure 14. Same as figure 13, with $[\text{Eu}/\text{Fe}]$ in place of $[\text{Ba}/\text{Fe}]$ and with stars chosen without restriction on T_{eff} .

6. DISCUSSION & CONCLUSIONS

We have trained a data-driven model (The Cannon) to measure detailed abundances from low-resolution LAMOST spectra, delivering up to 7 abundances for 3.9×10^6 stars to a precision of $0.05 - 0.23$ dex. 2.9×10^6 of these are dwarf stars, for which we infer labels with $0.05 - 0.23$ dex precision. 8.8×10^5 are red giants for which we infer abundances with $0.07 - 0.22$ dex precision. Our best-fit model spectra are easily reproducible using our catalog,

implementation of The Cannon, and model coefficients², which are available online. We used the red giants to examine the spatial distribution of abundances in the disk and halo and the dwarf stars to investigate the chemical similarity of wide binaries.

Our analysis of the chemical similarity of dwarf stars in wide binaries compared to field stars showed these stars are from a common birth site and enabled us to quantify the additional resolving or discriminating

² **AJW: fix** See github.com/ajwheeler/LAMOST_abundances

power in the vector of our derived abundances beyond an overall metallicity.

Using the red giants, we first mapped the profile of the disk in the (R, z) plane in the elements O, Eu, Mg, Al, Mn, and Ba. These maps show the flaring of the disk and the distinction in abundances between the halo population at high latitudes and the disk. Second, we examined face-on projections across the disk in this set of abundances. These projections hint at some non-axisymmetric patterns in the abundances. Indeed, the *Gaia* mission has revealed numerous dynamical deviations from axisymmetry in the disk and perturbations in the solar neighbourhood (Antoja et al. 2018; Sellwood et al. 2019; Trick et al. 2019). Third, we construct mean-abundance maps in the (R, v_ϕ) plane and discuss potential signature of the bar’s interaction with the disk. Finally, we investigated the abundance planes of [Ba/Fe]-[Fe/H] and [Eu/Fe]-[Fe/H] across (R, z) , making similar maps to Hayden et al. (2015), but with the neutron capture abundances. These maps showed the disk and halo trends across [Fe/H] at all (R, z) . These different trends might be used to separate any in-situ halo from heated disk stars, from an accreted halo. As the set of abundances we deliver give higher discriminating power to identify chemically similar stars compared to [Fe/H] alone, we expect the multiple families of abundances will be useful for studies of the plethora of chemodynamical substructure in the Milky Way halo (see e.g. Belokurov et al. 2019; Helmi et al. 2018; Di Matteo et al. 2018; Myeong et al. 2019; Antoja et al. 2018).

We derive abundances with diverse nucleosynthetic channels and are demonstrably uncovering some of the breadth of chemical information in the Milky Way. However, caveats are discussed in Section 5.1 and we further detail some of these here.

In contrast to Ho et al. (2017b,a), we find that our cross validation results do not vary strongly with S/N. This indicates that our precision is limited by that of reference labels themselves, and, if improved, we would obtain higher precision results for our test objects that scale as expected with SNR.

Because our model is in some cases not inferring abundances from the corresponding lines themselves, it may not be robust to stars with properties or enrichment histories not represented in the training set. This means that, while stars that are highly enriched or depleted in an element may not have their abundances accurately inferred, the best fit model should have large residuals in regions of its known lines. Casey et al. (2019) used this effect to identify stars that are unusually rich in lithium, but this approach could be extended to all el-

ements with strong lines in the LAMOST wavelength range.

Training a model without α or Mg appears yields a systematic offset in inferred neutron-capture abundances, likely because the model will exploit correlations between these nucleosynthetic families if they are not “controlled for”.

The large number of low-resolution spectra available now (RAVE and SEGUE), in addition to LAMOST) and coming in the near future (e.g. WEAVE, Dalton et al. 2012; MOONs, Cirasuolo et al. 2014; 4MOST, de Jong et al. 2019; DESI, DESI Collaboration et al. 2016) and Sloan V (Kol ???), makes honing our ability to learn from these data a fruitful endeavor. We also discussed future improvements to our methodology, especially the possibility of using open clusters to reduce the effect of systematic trends with stellar parameters in inferred abundances. Other promising methodological directions include using more robust inference for model parameters and labels, perhaps allowing more rigorous error estimation, and allowing missing labels in the training set, which would enable us to use training data from multiple surveys.

Software: Optim.jl (Mogensen & Riseth 2018), Matplotlib (Hunter 2007), galpy (Bovy 2015)

ACKNOWLEDGMENTS

The authors thank Kathryn Johnston, James Aplegate, David Helfand, Tomer Yavetz, Klemen Cotar, Francesca Fragkoudi, and Matthew Abruzzo for useful discussion and notes.

AJW is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1644869.

MKN is in part supported by a Sloan Research Fellowship. We thank and acknowledge the Kavli Institute for Theoretical Physics at the University of California, Santa Barbara, and in particular the Gaia19 workshop. This work was also in part inspired during the Aspen Center for Physics ‘Dynamics of the Milky Way System in the Era of Gaia’ workshop, which is supported by National Science Foundation grant PHY-1607611.

SB acknowledges funds from the Alexander von Humboldt Foundation in the framework of the Sofja Kovalevskaja Award endowed by the Federal Ministry of Education and Research. This research has been supported by the Australian Research Council (grants DP150100250 and DP160103747). Parts of this research were supported by the Australian Research Council (ARC) Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013.

JDS acknowledges the support of the Australian Research Council through Discovery Project grant DP180101791

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

Guoshoujing Telescope (the Large Sky Area Multi-Object Fiber Spectroscopic Telescope LAMOST) is a National Major Scientific Project built by the Chinese Academy of Sciences. Funding for the project has been provided by the National Development and Reform Commission. LAMOST is operated and managed by the National Astronomical Observatories, Chinese Academy of Sciences.

We acknowledge computing resources from Columbia University’s Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant 1G20RR030893-01, and associated funds from the New York State Empire State Development,

Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010.

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, and NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology. WISE and NEOWISE are funded by the National Aeronautics and Space Administration.

REFERENCES

????

- Antoja, T., Helmi, A., Romero-Gomez, M., et al. 2018, *Nature*, 561, 360–362, arXiv: 1804.10196
- Armandroff, T. E., & Zinn, R. 1988, *The Astronomical Journal*, 96, 92
- Arnould, M., Goriely, S., & Takahashi, K. 2007, *Physics Reports*, 450, 97–213, arXiv: 0705.4512
- Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Mantelet, G., & Andrae, R. 2018, *AJ*, 156, 58
- Bedell, M., Bean, J. L., Meléndez, J., et al. 2018, *ApJ*, 865, 68
- Belokurov, V., Sanders, J. L., Fattahi, A., et al. 2019, arXiv:1909.04679 [astro-ph], arXiv: 1909.04679
- Bensby, T., Feltzing, S., & Oey, M. S. 2014, *Astronomy & Astrophysics*, 562, A71, arXiv: 1309.2631
- Binney, J. 2012, *MNRAS*, 426, 1324
- Blancato, K., Ness, M., Johnston, K. V., Rybizki, J., & Bedell, M. 2019, arXiv:1906.05297 [astro-ph], arXiv: 1906.05297
- Bonaca, A., Conroy, C., Wetzell, A., Hopkins, P. F., & Keres, D. 2017, *The Astrophysical Journal*, 845, 101, arXiv: 1704.05463
- Bovy, J. 2015, *ApJS*, 216, 29
- Bovy, J., & Rix, H.-W. 2013, *ApJ*, 779, 115
- Buder, S., Asplund, M., Duong, L., et al. 2018, *MNRAS*, 478, 4513
- Buder, S., Lind, K., Ness, M. K., et al. 2019, *Astronomy & Astrophysics*, 624, A19, arXiv: 1804.05869
- Cantat-Gaudin, T., Jordi, C., Vallenari, A., et al. 2018, *Astronomy & Astrophysics*, 618, A93, arXiv: 1805.08726
- Casey, A. R., Hogg, D. W., Ness, M., et al. 2016, arXiv e-prints, arXiv:1603.03040
- Casey, A. R., Ho, A. Y. Q., Ness, M., et al. 2019, arXiv e-prints, arXiv:1902.04102
- Cirasuolo, M., Afonso, J., Carollo, M., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, Proc. SPIE, 91470N
- Collaboration, G., Prusti, T., de Bruijne, J. H. J., et al. 2016, *Astronomy and Astrophysics*, 595, A1
- Cunha, K., Smith, V. V., Hasselquist, S., et al. 2017, *ApJ*, 844, 145
- Côté, B., Fryer, C. L., Belczynski, K., et al. 2018, *The Astrophysical Journal*, 855, 99, arXiv: 1710.05875
- da Silva, R., Lemasle, B., Bono, G., et al. 2016, *Astronomy & Astrophysics*, 586, A125, arXiv: 1510.06314
- Dalton, G., Trager, S. C., Abrams, D. C., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8446, Proc. SPIE, 84460P
- de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *The Messenger*, 175, 3
- De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, 449, 2604
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv e-prints, arXiv:1611.00036
- Di Matteo, P., Haywood, M., Lehnert, M. D., et al. 2018, arXiv:1812.08232 [astro-ph], arXiv: 1812.08232

- El-Badry, K., Rix, H.-W., Tian, H., Duchêne, G., & Moe, M. 2019, arXiv:1906.10128 [astro-ph], arXiv: 1906.10128
- Feuillet, D. K., Frankel, N., Lind, K., et al. 2019, arXiv:1908.02772 [astro-ph], arXiv: 1908.02772
- Feuillet, D. K., Bovy, J., Holtzman, J., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 477, 2326–2348, arXiv: 1803.06352
- Fragkoudi, F., Grand, R. J. J., Pakmor, R., et al. 2019a, arXiv:1911.06826 [astro-ph], arXiv: 1911.06826
- Fragkoudi, F., Katz, D., Trick, W., et al. 2019b, *Monthly Notices of the Royal Astronomical Society*, stz1875, arXiv: 1901.07568
- Frankel, N., Rix, H.-W., Ting, Y.-S., Ness, M. K., & Hogg, D. W. 2018, *The Astrophysical Journal*, 865, 96, arXiv: 1805.09198
- Freeman, K., & Bland-Hawthorn, J. 2002, *Annual Review of Astronomy and Astrophysics*, 40, 487–537
- Gao, X., Lind, K., Amarsi, A. M., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 2666–2684, arXiv: 1804.06394
- García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, *The Astronomical Journal*, 151, 144
- Gilmore, G., Randich, S., Asplund, M., et al. 2012, *The Messenger*, 147, 25–31
- Hasselquist, S., Shetrone, M., Cunha, K., et al. 2016, *ApJ*, 833, 81
- Hayden, M. R., Bovy, J., Holtzman, J. A., et al. 2015, *The Astrophysical Journal*, 808, 132
- Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, *Nature*, 563, 85
- Ho, A. Y. Q., Rix, H.-W., Ness, M. K., et al. 2017a, *The Astrophysical Journal*, 841, 40
- Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017b, *The Astrophysical Journal*, 836, 5, arXiv: 1602.00303
- Holtzman, J. A., Hasselquist, S., Shetrone, M., et al. 2018, *The Astronomical Journal*, 156, 125, arXiv: 1807.09773
- Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90
- Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008, *The Astrophysical Journal*, 673, 864
- Kamdar, H., Conroy, C., Ting, Y.-S., et al. 2019, arXiv e-prints, arXiv:1904.02159
- Kawata, D., Baba, J., Ciucă, I., et al. 2018, *Monthly Notices of the Royal Astronomical Society: Letters*, 479, L108–L112, arXiv: 1804.10175
- Kemp, A. J., Casey, A. R., Miles, M. T., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 1384–1392, arXiv: 1807.05693
- Krumholz, M. R., McKee, C. F., & Bland-Hawthorn, J. 2019, *ARA&A*, 57, 227
- Lee, Y. S., Beers, T. C., Prieto, C. A., et al. 2011, *The Astronomical Journal*, 141, 90, arXiv: 1010.2934
- Leung, H. W., & Bovy, J. 2018, *Monthly Notices of the Royal Astronomical Society*, doi:10.1093/mnras/sty3217, arXiv: 1808.04428
- Li, J., Han, C., Xiang, M.-S., et al. 2016, *Research in Astronomy and Astrophysics*, 16, 110
- Mainzer, A., Bauer, J., Grav, T., et al. 2011, *ApJ*, 731, 53
- Majewski, S. R., Zasowski, G., & Nidever, D. L. 2011, *ApJ*, 739, 25
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *The Astronomical Journal*, 154, 94
- Minchev, I., Chiappini, C., & Martig, M. 2014, *Astronomy & Astrophysics*, 572, A92, arXiv: 1401.5796
- Minchev, I., Famaey, B., Quillen, A. C., et al. 2012, *Astronomy & Astrophysics*, 548, A127, arXiv: 1205.6475
- Minchev, I., Anders, F., Recio-Blanco, A., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 481, 1645, arXiv: 1804.06856
- Mogensen, P. K., & Riseth, A. N. 2018, *Journal of Open Source Software*, 3, 615
- Myeong, G. C., Vasiliev, E., Iorio, G., Evans, N. W., & Belokurov, V. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 1235–1247, arXiv: 1904.03185
- Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015, *ApJ*, 808, 16
- Ness, M., Hogg, D. W., Rix, H.-W., et al. 2016, *The Astrophysical Journal*, 823, 114, arXiv: 1511.08204
- Ness, M., Rix, H.-W., Hogg, D. W., et al. 2018, *The Astrophysical Journal*, 853, 198
- Ness, M. K., Johnston, K. V., Blancato, K., et al. 2019, arXiv e-prints, arXiv:1907.10606
- Newberg, H. J. e. a. 2012, *Research in Astronomy and Astrophysics*, 12, 735–754, arXiv: 1206.3578
- Nidever, D. L., Bovy, J., Bird, J. C., et al. 2014, *The Astrophysical Journal*, 796, 38, arXiv: 1409.3566
- Norfolk, B. J., Casey, A. R., Karakas, A. I., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 2219–2227, arXiv: 1911.00177
- Overbeek, J. C., Friel, E. D., & Jacobson, H. R. 2016, *The Astrophysical Journal*, 824, 75, arXiv: 1604.05735
- Price-Jones, N., & Bovy, J. 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 1410
- Ramirez, I., Khanal, S., Lichon, S. J., et al. 2019, arXiv:1909.07460 [astro-ph], arXiv: 1909.07460
- Roškar, R., Debattista, V. P., Quinn, T. R., Stinson, G. S., & Wadsley, J. 2008, *The Astrophysical Journal*, 684, L79, arXiv: 0808.0206
- Rybizki, J., Just, A., & Rix, H.-W. 2017, *Astronomy & Astrophysics*, 605, A59, arXiv: 1702.08729

- Schoenrich, R., & Binney, J. 2009, *Monthly Notices of the Royal Astronomical Society*, 399, 1145–1156, arXiv: 0907.1899
- Schönrich, R., & Weinberg, D. H. 2019, *Monthly Notices of the Royal Astronomical Society*, 487, 580–594, arXiv: 1901.09938
- Sellwood, J. A., Trick, W., Carlberg, R., Coronado, J., & Rix, H.-W. 2019, *Monthly Notices of the Royal Astronomical Society*, 484, 3154–3167, arXiv: 1810.03325
- Siegel, D. M., Barnes, J., & Metzger, B. D. 2019, *Nature*, 569, 241
- Silverman, B. W. 1986, *Density Estimation for Statistics and Data Analysis* (Routledge)
- Simpson, J. D., Martell, S. L., Da Costa, G., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 482, 5302–5315, arXiv: 1804.05894
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
- Souto, D., Prieto, C. A., Cunha, K., et al. 2019, *The Astrophysical Journal*, 874, 97, arXiv: 1902.10199
- Spina, L., Mel'endez, J., Karakas, A. I., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 474, 2580
- Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *The Astronomical Journal*, 132, 1645–1668, arXiv: astro-ph/0606211
- Ting, Y.-S., Conroy, C., Rix, H.-W., & Asplund, M. 2018, *ApJ*, 860, 159
- Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2018a, arXiv:1804.01530 [astro-ph], arXiv: 1804.01530
- Ting, Y.-S., Hawkins, K., & Rix, H.-W. 2018b, *The Astrophysical Journal*, 858, L7, arXiv: 1803.06650
- Trick, W. H., Fragkoudi, F., Hunt, J. A. S., Mackereth, J. T., & White, S. D. M. 2019, arXiv:1906.04786 [astro-ph], arXiv: 1906.04786
- Vásquez, S., Zoccali, M., Hill, V., et al. 2015, *Astronomy & Astrophysics*, 580, A121
- Weinberg, D. H., Holtzman, J. A., Hasselquist, S., et al. 2019, *The Astrophysical Journal*, 874, 102, arXiv: 1810.12325
- Wenger, T. V., Balsa, D. S., Anderson, L. D., & Bania, T. M. 2019, arXiv:1910.14605 [astro-ph], arXiv: 1910.14605
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, 140, 1868
- Xiang, M., Liu, X., Shi, J., et al. 2017, *Monthly Notices of the Royal Astronomical Society*, 464, 3657–3678, arXiv: 1610.00083
- Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, arXiv:1908.09727 [astro-ph], arXiv: 1908.09727
- Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, *The Astronomical Journal*, 137, 4377
- Zhang, X., Zhao, G., Yang, C. Q., Wang, Q. X., & Zuo, W. B. 2019, *PASP*, 131, 094202
- Zhao, Y. e. a. 2012, arXiv:1206.3569 [astro-ph], arXiv: 1206.3569

APPENDIX

A. LINEAR COEFFICIENTS

Figures 15 and 16 show the linear coefficients of our models as a function of wavelength.

B. ABUNDANCES OF PRAESEPE

Figure 17 shows systematic trends in $[\alpha/\text{Fe}]$ as a function of T_{eff} by plotting inferred values for stars in Praesepe, which, as an open cluster, is chemically homogeneous to a level well below our precision. Other abundances exhibit similar behavior. The dispersion of $[\alpha/\text{Fe}]$ values exhibited in Praesepe is 0.07 dex—contrast this with the CV-assessed precision for dwarfs: 0.06 dex. This indicates that systematic error as a function of stellar parameters is a major contributor to our error budget.

C. GALAH WINDOWS

Table 2 lists the GALAH line windows for each measured element. We used these windows to calculate per-element χ^2 values for GALAH spectra to eliminate spurious measurements from our training set.




Figure 15. Linear coefficients for the dwarf model.




Figure 16. Linear coefficients for the giant model




Figure 17. Inferred $[\alpha/\text{Fe}]$ vs effective temperature for LAMOST stars in Praesepe. The systematic trend with T_{eff} is spurious, since all stars in Praesepe have the same abundances to below the precision achievable with low-resolution spectra.

Table 2. GALAH line windows for each element

element	windows [\AA]
Al	(6695.78, 6696.17), (6698.41, 6698.92), (7834.95, 7835.47), (7835.84, 7836.43)
Ba	(5853.53, 5853.86), (6496.68, 6497.19)
Ca	(5857.22, 5857.60), (5867.28, 5867.72), (6493.48, 6493.99), (6499.37, 6499.94), (6508.52, 6509.03)
Co	(6632.23, 6632.81), (7712.41, 7713.07), (7837.76, 7838.50)
Cr	(4775.03, 4775.21), (4789.20, 4789.47), (4800.83, 4801.20), (4847.98, 4848.31), (5702.12, 5702.50), (5719.50, 5719.99), (5787.64, 5788.14), (5844.40, 5844.79), (6629.80, 6630.25)
Cu	(5781.92, 5782.42)
Eu	(5818.61, 5818.99), (6644.97, 6645.29)
K	(7698.57, 7699.31)
La	(4716.29, 4716.61), (4748.62, 4748.85), (4803.92, 4804.24), (5805.52, 5805.96)
Li	(6707.37, 6708.26)
Mg	(4729.90, 4730.22), (5710.86, 5711.30)
Mn	(4739.01, 4739.29), (4761.37, 4761.64), (4765.69, 4766.06), (4783.17, 4783.58)
Na	(4751.71, 4751.94), (5682.54, 5682.92), (5687.93, 5688.37)
Ni	(5748.21, 5748.59), (5846.82, 5847.21), (6482.60, 6483.05), (6532.58, 6533.10), (6586.02, 6586.47), (6643.37, 6643.94), (7713.89, 7714.48), (7788.48, 7789.29)
O	(7771.53, 7772.27), (7773.75, 7774.57), (7775.08, 7775.75)
Sc	(4743.56, 4743.98), (4752.99, 4753.41), (5657.68, 5658.12), (5666.92, 5667.30), (5671.59, 5672.09), (5684.02, 5684.30), (5686.72, 5687.21), (5717.02, 5717.52), (5723.90, 5724.28), (6604.39, 6604.97)
Si	(5665.21, 5665.82), (5690.18, 5690.68), (5700.91, 5701.29), (5792.70, 5793.31)
Ti	(4719.32, 4719.60), (4757.96, 4758.28), (4759.07, 4759.48), (4764.40, 4764.82), (4778.06, 4778.43), (4781.56, 4781.93), (4797.84, 4798.12), (4798.35, 4798.63), (4801.80, 4802.21), (4820.11, 4820.66), (4849.04, 4849.41), (4865.28, 4865.83), (4873.88, 4874.20), (5689.25, 5689.80), (5716.25, 5716.80), (5720.27, 5720.65), (5739.24, 5739.68), (5866.02, 5866.79), (6598.89, 6599.53), (6716.52, 6716.90), (7852.19, 7853.01)
V	(4746.51, 4746.78), (4784.32, 4784.60), (4796.74, 4796.97), (4831.52, 4831.75), (4875.26, 4875.72), (5657.07, 5657.68), (5668.13, 5668.62), (5670.60, 5671.10), (5702.78, 5703.88), (5725.27, 5725.82), (5726.87, 5727.36), (5727.42, 5727.97), (5730.99, 5731.49), (5736.82, 5737.26), (5743.20, 5743.70), (6531.18, 6531.62)
Y	(4854.79, 4855.02), (4883.54, 4883.82), (5662.74, 5663.18), (5728.74, 5729.01)
Zn	(4721.99, 4722.27), (4810.36, 4810.63)