Chemical Doppelgangers in GALAH DR3: The Distinguishing Power of Neutron-capture Elements among Milky Way Disk Stars

Catherine Manea[1] , Keith Hawkins[1] , Melissa K. Ness[2,3] , Sven Buder[4,5] , Sarah L. Martell[5,6] , and Daniel B. Zucker[7,8]

[1] Department of Astronomy, The University of Texas at Austin, 2515 Speedway Boulevard, Austin, TX 78712, USA
[2] Department of Astronomy, Columbia University, Pupin Physics Laboratories, New York, NY 10027, USA
[3] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA
[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), Australia
[6] School of Physics, UNSW, Sydney, NSW 2052, Australia
[7] School of Mathematical and Physical Sciences, Macquarie University, Sydney, NSW 2109, Australia
[8] Astrophysics and Space Technologies Research Centre, Macquarie University, Sydney, NSW 2109, Australia

Abstract

The observed chemical diversity of Milky Way stars places important constraints on Galactic chemical evolution and the mixing processes that operate within the interstellar medium. Recent works have found that the chemical diversity of disk stars is low. For example, the Apache Point Observatory Galactic Evolution Experiment (APOGEE) "chemical doppelganger rate," or the rate at which random pairs of field stars appear as chemically similar as stars born together, is high, and the chemical distributions of APOGEE stars in some Galactic populations are well-described by two-dimensional models. However, limited attention has been paid to the heavy elements ($Z > 30$) in this context. In this work, we probe the potential for neutron-capture elements to enhance the chemical diversity of stars by determining their effect on the chemical doppelganger rate. We measure the doppelganger rate in GALactic Archaeology with HERMES DR3, with abundances rederived using *The Cannon*, and find that considering the neutron-capture elements decreases the doppelganger rate from ∼2.2% to 0.4%, nearly a factor of 6, for stars with $-0.1 < [\mathrm{Fe/H}] < 0.1$. While chemical similarity correlates with similarity in age and dynamics, including neutron-capture elements does not appear to select stars that are *more* similar in these characteristics. Our results highlight that the neutron-capture elements contain information that is distinct from that of the lighter elements and thus add at least one dimension to Milky Way abundance space. This work illustrates the importance of considering the neutron-capture elements when chemically characterizing stars and motivates ongoing work to improve their atomic data and measurements in spectroscopic surveys.

Unified Astronomy Thesaurus concepts: Galactic archaeology (2178); Chemical abundances (224); Galaxy chemical evolution (580)

1. Introduction

The recent decade has brought forth an exponential increase in available stellar spectroscopic data, enabling population-level analyzes of the chemical compositions of Milky Way stars at unprecedented scale. Massive spectroscopic surveys such as Apache Point Observatory Galactic Evolution Experiment (APOGEE; Abdurro'uf et al. 2022), Large Sky Area Multi-Object Fiber Spectroscopic Telescope (Cui et al. 2012), Gaia-European Southern Observatory (Gilmore et al. 2022), Hectochelle in the Halo at High Resolution (Conroy et al. 2019), the RAdial Velocity Experiment (Steinmetz et al. 2020), and GALactic Archaeology with HERMES (GALAH; Buder et al. 2021) have provided the Galactic science community with *millions* of stellar spectra. In the wake of this abundance of stellar spectroscopic data, recent work has begun to investigate *how* much information is actually contained in these data sets. Does each element carry unique information, or are many of these abundances correlated? More specifically, do stars at a fixed metallicity tend to display similar chemical profiles in elements across the periodic table, or is there more diversity in

their possible compositions? The literature refers to this notion as the *dimensionality* of chemical abundance space.

The dimensionality of Milky Way abundance space carries both physical and practical implications that affect a wide range of subfields of astronomy. When viewed through the physical lens, the dimensionality of Milky Way chemical abundances traces the stability of nucleosynthetic yields across Galactic time and space and the scale and efficiency at which newly synthesized elements are dispersed and mixed into the interstellar medium. Furthermore, because stellar compositions dictate the architectures and compositions of their planetary systems (e.g., Fischer & Valenti 2005; Nielsen et al. 2023), the dimensionality of abundance space dictates the expected diversity of planetary systems around Milky Way stars. When viewed through the practical lens, the question of chemical dimensionality seeks to assess whether one truly needs to measure tens of elements to fully understand a star's composition, or if measuring a few elements and inferring the rest produces the same quantity and quality of information, thereby significantly enhancing efficiency. Additionally, the dimensionality of Milky Way abundance space contributes to the validity of strong chemical tagging, the method of reconstructing dispersed stellar birth siblings using chemistry alone (e.g., Freeman & Bland-Hawthorn 2002). Strong chemical tagging poses two requirements: (1) that stars born together are chemically homogeneous (e.g., Bovy 2016;

Hawkins et al. 2020; Nelson et al. 2021) and (2) that groups of stars born together are chemically unique (e.g., Lambert & Reddy 2016; Price-Jones et al. 2020; Cheng et al. 2021). If chemical dimensionality is sufficiently high, then the strong chemical tagging requirement that each birth cluster possesses a unique chemical profile could be satisfied.

Ting et al. (2012) was among the first works to directly investigate the dimensionality of Milky Way chemical abundance space. They performed principal component analysis on combined data from several different high-resolution spectroscopic studies to find that stellar abundances tend to possess between six and nine principal components associated with various nucleosynthetic sites, a result they later validate in APOGEE (Ting & Weinberg 2022). Other works, however, find that chemical abundances may have as few as two dimensions. For example, Ness et al. (2019) found that for the low-α disk, with a star's [Fe/H] and age, one can predict its remaining APOGEE abundances to within 0.02 dex. Furthermore, a star's abundance of elements produced in supernovae can be predicted to $\approx$0.015 dex using Fe, Mg, and age (Ness et al. 2022). This is on the order of the intrinsic scatter of these elements within open clusters, where stars are known to be born together (e.g., Bovy 2016). This implies that these three dimensions link to birth radii, but individual elements produced in supernovae cannot be used to distinguish individual birth groups. Sharma et al. (2022) similarly found that the abundances of most GALAH-reported elements can be inferred to within 0.03 dex using just [Fe/H] and age, though they note that certain elements such as Y and Ba are exceptions to this and cannot be predicted well. Weinberg et al. (2022) and Griffith et al. (2022, 2024) also addressed this question. They create a two-dimensional (also called a two-process) model that describes the chemical evolution of the Milky Way according to global enrichment due to time as well as the relative ratio of Type Ia to Type II supernovae. They then subtract this model from the APOGEE and GALAH data and study the residual abundance patterns in the data. In APOGEE, the two-dimensional model is suitable enough to predict a star's APOGEE abundances to within 0.03 dex for all well-measured elements (Weinberg et al. 2022), with the addition of a third dimension modeling asymptotic giant branch (AGB) star nucleosynthesis marginally improving the representation of the data (Griffith et al. 2024). In GALAH, the two-dimensional model produces abundance residuals less than 0.07 dex for most well-measured elements, and the addition of a third dimension, again associated with AGB star nucleosynthesis, further decreases abundance residuals (Griffith et al. 2022).

The majority of the analysis that has been done in the context of APOGEE reports a low-dimensionality of chemical abundance space. However, these analyses have been limited to nucleosynthetic channels traced by the light/odd-z, alpha, and iron-peak elements. While APOGEE sometimes measures Ce and Nd, products of the neutron-capture process, these elements have not been sufficiently investigated in the context of chemical dimensionality. Therefore, the dimensionality of the Milky Way's chemical abundances is still an open question. In this work, we seek to address the specific role that the neutron-capture elements play in the chemical dimensionality of Milky Way abundance space. The GALAH survey offers a more rigorous test of the underlying dimensionality of the Milky Way disk, as it is explicitly designed to capture five channels of nucleosynthetic enrichment (light/odd-Z, alpha,

iron-peak, slow neutron-capture process, and rapid neutron-capture process elements) for the purpose of testing the validity of strong chemical tagging (De Silva et al. 2015). Most critically, due to its optical coverage and high resolution ($R \sim 28{,}000$), GALAH measures the neutron-capture elements. These include elements formed in both the rapid (r-) and slow (s-) neutron-capture processes, two nucleosynthetic families that are not presently well-measured in APOGEE but may add to the dimensionality of Milky Way abundance space (e.g., Lambert & Reddy 2016; Griffith et al. 2022; Sharma et al. 2022).

We study the potentially unique information contained in neutron-capture elements through the lens of "chemical doppelgangers," stars that are highly chemically similar but otherwise dynamically unrelated. Ness et al. (2018) was the first to measure the so-called "chemical doppelganger rate," defined as the rate at which randomly drawn pairs of field stars are measured to be as chemically similar as stars born together. Stars that are born together, such as those in open clusters, have been found to be relatively chemically similar, with intrinsic dispersions in APOGEE-measured elements ranging from 0.005–0.070 dex (e.g., Bovy et al. 2016; Poovelil et al. 2020), though chemical differences associated with atomic diffusion have been observed across stars of different evolutionary states (e.g., Liu et al. 2016, 2019; Gao et al. 2018; Souto et al. 2018, 2021; Beeson et al. 2024). Open cluster stars are typically considered to represent an upper limit for stellar chemical homogeneity and random field stars the lower limit. Ness et al. (2018) found that between 0.3% and 1% of randomly drawn field pairs in APOGEE DR13 (Holtzman et al. 2018) appear to be as chemically similar as stars residing in open clusters. While a fraction of these field pairs are likely true conatal pairs from dispersed open clusters, it is unlikely that all are conatal. To place this result in perspective, if all massive star clusters that have formed in the disk had unique chemical abundance profiles, the expected recovery rate of true birth siblings would be closer to 10^{-4} or 10^{-5} given the star formation history of the Milky Way (e.g., Bland-Hawthorn et al. 2010). Thus, most doppelgangers are likely not true birth siblings, and stars can share remarkably similar chemical profiles despite being born of different star-forming complexes due to a relatively homogeneous chemical evolution of the thin disk. These results are validated by de Mijolla et al. (2021) in APOGEE DR16 (Ahumada et al. 2020). Together, these works suggest that the diversity of Milky Way disk star abundances is qualitatively low.

In this work, we perform the first measurement of the doppelganger rate in GALAH. We center our investigation on whether the neutron-capture elements affect the measured doppelganger rate. If neutron-capture elements add at least one unique dimension to chemical abundance space at the available abundance precision, then we expect the doppelganger rate to decrease with the addition of this nucleosynthetic family. If neutron-capture elements do not add a unique dimension, then we expect the doppelganger rate to remain relatively unchanged with the addition of these elements. Through this test, we implicitly probe whether neutron-capture elements enhance the diversity of Milky Way stars or simply trace the lighter (light/odd-Z, α, and iron-peak) elements primarily produced in supernovae. We use the following questions to guide our analysis:

   i What is the doppelganger rate in GALAH?
   ii How does the inclusion of neutron-capture elements specifically affect the doppelganger rate?
   iii Do there exist pairs of stars that are "doppelganger" in the light, α, and iron-peak elements but show differences in the neutron-capture elements? If so, are there any physical characteristics that differentiate these pairs from pairs that are "doppelganger" in all elements?

In Section 2, we describe the GALAH data set and our choice of open cluster stars that serve as a reference in our doppelganger rate measurements. In Section 3, we use *The Cannon* to re-derive abundance ratios in 16 elements (Fe, O, Al, Mg, Ca, Si, Cr, Cu, Zn, Mn, Zr, Y, Ba, Ce, Nd, and Eu) for the purpose of enhancing precision and ensuring well-constrained uncertainties. We then measure the doppelganger rate using these re-derived abundances. In Section 4, we investigate the impact of the neutron-capture elements on the measured doppelganger rate and compare the physical characteristics of stars that are partial doppelgangers (doppelganger in the light, α, and iron-peak elements) with those that are complete doppelgangers (doppelganger in all measured elements). We discuss our results in the context of previous observational and simulation work in Section 5 and conclude in Section 6.

2. Data

GALAH Data Release 3 (DR3 Buder et al. 2021) serves as the basis of our investigation. GALAH is an optical (4710 Å $< \lambda <$ 7890 Å spread across four noncontiguous cameras), magnitude limited ($V < 14$) spectroscopic survey with high resolving power ($R \sim 28,000$) that targets Milky Way stars at $|b| > 10°$ (De Silva et al. 2015). GALAH DR3 reports stellar parameters (e.g., T_{eff}, log g, spectral broadening, [Fe/H], etc.) and abundances for nearly 600,000 stars derived using Spectroscopy Made Easy (SME; Valenti & Piskunov 1996; Piskunov & Valenti 2017) and 1D MARCS model atmospheres (Gustafsson et al. 1975; Bell et al. 1976; Gustafsson et al. 2008). Nonlocal thermodynamic equilibrium is assumed during spectral line synthesis of 13 elements (H, Li, C, O, Na, Mg, Al, Si, K, Ca, Mn, Fe, and Ba) whereas LTE is assumed for the rest. For each star, GALAH DR3 reports its surface abundances in up to 30 elements spanning six major nucleosynthetic families: the light (Li, C), α (Mg, Si, Ca, O), odd-Z (Na, Al, K), iron-peak (Sc, V, Cr, Mn, Fe, Co, Ni, Cu, Zn, Ti), and slow (s-) and rapid (r-) process neutron-capture (Rb, Sr, Y, Zr, Mo, Ru, Ba, La, Ce, Nd, Sm, and Eu) elements. In addition to the main catalog, DR3 provides the community with several other data products, such as a one-dimensional, radial-velocity-corrected, continuum-normalized, combined spectrum for nearly every star sampled by the survey as well as a catalog of dynamical parameters and age estimates for nearly all GALAH targets. Dynamical parameters are determined using Python package `galpy` (Bovy 2015; see GALAH survey webpage[9] for details on their assumed Galactic potentials and properties). Age estimates for GALAH stars are determined using Bayesian Stellar Parameter Estimation code from Sharma et al. (2018), which uses PARSEC release v1.2S + COLIBRI stellar isochrones (Marigo et al. 2017) and Bayesian estimation to infer intrinsic stellar parameters from observables T_{eff}, log g,

[Fe/H], [α/Fe], parallax, and the Two Micron All Sky Survey J- and H-band magnitudes.

For this investigation, we aim to measure the doppelganger rate using abundances that we re-derive from the GALAH DR3 spectra using *The Cannon* (Ho et al. 2017; see Section 3.1.1 for our motivation for this choice). We elect to perform our investigation on red giant and red clump stars in the GALAH survey, as they are on average brighter and thus probe a larger Galactic volume relative to dwarfs. This enables us to (a) better compare to Ness et al. (2018), which also used giants, and (b) understand the doppelganger rate on a broader spatial scale, expanding on work such as that of Bedell et al. (2018), which investigated the chemical diversity of stars within the local (100 pc) solar neighborhood. Though there are several advantages to using stars in this evolutionary state, age determinations for red giant branch and red clump stars are poorly constrained relative to those of main-sequence stars near the turnoff. This primarily impacts our later analysis (Section 4.3) where we study the age similarity of doppelgangers. To build our stellar sample, we apply a series of selections that we motivate in the following paragraph:

   i `flag_sp = 0`
   ii `flag_fe_h = 0`
   iii `snr_c2_iraf > 20`
   iv `ruwe < 1.2`
   v `1.5 < log g < 3.5`
   vi `0.0033*teff - 13.6 < log g < 0.0036*teff −13.9`
   vii `−1.20 < fe_h < 0.20`
   viii `−0.25 < Cr_fe < 0.15`
   ix `Cu_fe > −0.30`
   x `Zr_fe < 0.60`
   xi `Y_fe < 0.60`
   xii `Ba_fe < 0.80`
   xiii `−0.50 < Ce_fe < 0.40`
   xiv `Nd_fe < 0.60`
   xv `Eu_fe < 0.60`.

Though we re-derive abundances using *The Cannon*, the first two requirements ensure reliable GALAH-reported stellar parameters and [Fe/H] abundances, which consequently eliminates clearly problematic spectral data. The third requirement ensures that all spectra have a signal-to-noise ratio (S/N) above 20 in the ~5700 Å spectral region, and the fourth requirement filters for potential spectroscopic binaries (e.g., Belokurov et al. 2020) that were missed by the `flag_sp` flag. Requirements (v) and (vi) ensure we select red giant and red clump stars. The remaining requirements excise certain stars that have either extreme abundances (e.g., s-process enhanced stars, which are likely post-mass transfer systems and no longer reflect their natal composition) and/or GALAH-reported abundances that are not sampled by our high-quality training set and thus less reliably inferred by our model (see Section 3.1.4).

2.1. Open Cluster Catalog

As mentioned in Section 1, we define doppelgangers to be pairs of field stars that appear as chemically similar to one another as stars residing in open clusters. As such, a reliable set of reference open cluster stars is critical for our investigation. We use the open cluster catalog of Spina et al. (2021) to build this reference set. Spina et al. (2021) builds off of the widely used open cluster catalog of Cantat-Gaudin et al. (2018),

---

specifically improving cluster membership determinations for GALAH-sampled open clusters. They make use of a Support Vector Machine classifier (see their footnote 2 for a detailed description) to re-assess cluster memberships using Gaia astrometry (Brown et al. 2021) and validate their results with careful inspection of the resulting cluster isochrones and radial-velocity distributions. For each cluster star, they report a membership probability. We only consider open cluster stars that have a probability of membership that exceeds 50% ($P_{\mathrm{mem}} > 0.5$). In practice, the vast majority of open cluster stars in our selection have membership probabilities between 90% and 100%, but this selection allows for a few additional stars with membership probabilities of 75%.

3. Method

The primary goal of this work is to measure how often random pairs of field stars sampled by GALAH appear as chemically similar as GALAH stars in open clusters, referred to as the doppelganger rate. Throughout this work, we closely follow the method of Ness et al. (2018). They compute the doppelganger rate using 20 elemental abundance ratios ([Fe/H], [C/Fe], [N/Fe], [O/Fe], [Na/Fe], [Mg/Fe], [Al/Fe], [Si/Fe], [S/Fe], [K/Fe], [Ca/Fe], [Ti/Fe], [V/Fe], [Mn/Fe], [Ni/Fe], [P/Fe], [Cr/Fe], [Co/Fe], [Cu/Fe], and [Rb/Fe]) homogeneously derived from APOGEE DR13 spectra using The Cannon (Ness et al. 2015). To compute the doppelganger rate, they draw random pairs of field stars unassociated with known clusters and compare their chemical similarity to that of random stellar pairs drawn from within open clusters, which they refer to as intracluster pairs. To quantify the abundance similarity of stars in a pair, they compute a χ^2 value for each pair, defined as:

$$\chi^2_{nn'} = \sum_{i=1}^{I} \frac{[x_{ni} - x_{n'i}]^2}{\sigma^2_{ni} + \sigma^2_{n'i}}, \qquad (1)$$

where the two stars in the pair are indexed as n, n', and x, σ are their derived abundance and abundance uncertainty, respectively, in element i. This leads to a global chemical similarity metric for each pair that considers all sampled elements. Doppelgangers are defined as stellar pairs with χ^2 values less than the median of intracluster pairs.

3.1. The Cannon

As in Ness et al. (2018), we measure the doppelganger rate using abundances and abundance uncertainties re-derived using The Cannon (Ho et al. 2017). The Cannon is a data-driven method for determining parameters and abundances from stellar spectra. The Cannon does not explicitly use atomic physics to determine these parameters: instead, it fits a suitably flexible model to the relationship between each spectral pixel's intensity and each input label (e.g., T_{eff}, log g, [Fe/H], etc.) using a high-fidelity training set. We use a second-order polynomial to model the relationship between spectral pixels and the following labels: T_{eff}, log g, broadening velocity, [Fe/H], [O/Fe], [Mg/Fe], [Al/Fe], [Si/Fe], [Ca/Fe], [Cr/Fe], [Cu/Fe], [Zn/Fe], [Mn/Fe], [Zr/Fe], [Y/Fe], [Ba/Fe], [Ce/Fe], [Nd/Fe], and [Eu/Fe]. We subsequently infer these labels at test time in our implementation. This model is similarly used in several other works that require high-precision abundances for hundreds of thousands of stars (e.g., Buder et al. 2018; Wheeler et al. 2020; Walsen et al. 2024). The Cannon is a

generative model, and constructs, for each label inference, a probability distribution function for the observed flux—that is, a theoretical spectrum for each star for which the labels are inferred. This enables the goodness-of-fit to be evaluated for the model spectrum versus the data, for each label and for each star. We direct the reader to Ness et al. (2015) for a thorough description of the methodology of The Cannon.

3.1.1. Why Re-derive GALAH DR3 Abundances Using The Cannon?

At the core of our investigation is the determination of the chemical similarity of pairs of stars. In an ideal setting, we could simply take the absolute difference in the elemental abundance ratios of each star in a pair to determine their degree of chemical similarity. However, this is impossible because all abundance measurements have an associated uncertainty, and thus, we must factor them into our determinations of chemical similarity. In this work, The Cannon is critical for (1) providing improved precision of the measured abundances and (2) providing accurate uncertainty estimates for stars that we can validate are well fit by our model in spectral space.

Abundance precision is a key extrinsic factor that influences the measured doppelganger rate. With large uncertainties in [X/Fe], it is more difficult to distinguish true doppelgangers from stars that show chemical differences unresolved at the current precision level. As such, maximizing precision is key to accurately constraining the doppelganger rate. As explained in Buder et al. (2021), The Cannon is capable of outperforming the GALAH DR3's SME-derived abundance precision. For GALAH DR3, elemental abundances were derived using on-the-fly spectral synthesis, and in order to limit computation time, the syntheses were only performed for a selection of unblended spectral regions associated with each element of interest. This means that some elemental abundances, such as [Mg/Fe], were derived using just one or two spectral lines, leading to larger uncertainties (e.g., Jofré et al. 2019). The Cannon in part achieves increased abundance precision because it leverages the entire spectrum to retrieve abundance information. This means it can use all spectral features—strong lines, weak lines, blended lines, lines with uncertain line data, and even continuum effects—associated with each element to infer an abundance. It is known that changes in the chemical composition of a star's atmosphere will affect the atmospheric opacity profile of the star, particularly when the abundance of electron-donating atoms is altered (Ting et al. 2018). Changes in opacity will consequently affect several different parts of the spectrum, including the continuum and line strengths of other elements. The Cannon is able to capture these trends, though careful consideration of nonphysical correlations between abundances and spectral features must also be taken (see Section 3.1.4).

Well-constrained abundance uncertainties are also key in this investigation. Underestimated uncertainties will artificially decrease the doppelganger rate while overestimated uncertainties will artificially increase it. Abundance uncertainties from classical abundance determination methods are influenced by a collection of sources, including but not limited to uncertainties in the data reduction process, input atomic physics, and choice of continuum placement when fitting a spectrum (e.g., Jofré et al. 2017). Using The Cannon, we can determine a systematic cross-validation uncertainty for each label, that represents the fidelity with which we recover the training labels. This measurement uncertainty incorporates the uncertainty on the

training labels that are inherited from GALAH. However, we can take advantage of repeat visits of individual stars to quantify our internal precision, which represents the overall systematic precision with which we can determine each stellar label, and this is what we ultimately adopt for our abundance uncertainties (see Section 3.1.5). It is important to highlight that in our investigation, abundance accuracy, which we inherit from the training set of stars and cannot control, is not necessarily important. Our only requirement is that spectra that are identical possess identical labels, so any global offsets in abundance do not impact our result, only relative offsets between different stars due to differing chemistry. Systematic trends in the data such as unphysical correlations between [Fe/H] versus $T_{\rm eff}$ can, however, impact our analysis. As such, when chemically comparing stars, we restrict comparison to stars with similar $T_{\rm eff}$ and log g values (see Section 3.3).

3.1.2. Re-derived Stellar Parameters and Abundances

Instead of re-deriving the full array of elements reported by GALAH, we select a subsample of the elements. We select elements that enable us to ask our scientific question of the data but limit our selection to the set of abundances for which we can build a high-fidelity training set. Our selection samples elements from each major nucleosynthetic family, including a light/odd-Z element (Al), α elements (O, Mg, Ca, Si), iron-peak elements (Fe, Cr, Cu, Zn, Mn), first (s-process) peak elements (Y, Zr), second s-process peak elements (Ba, Ce, Nd), and an r-process element (Eu). In addition to these elements, we re-derive $T_{\rm eff}$, log g, and `v_broad` (broadening velocity due to rotation, macroturbulence, etc.) for each spectrum.

3.1.3. Additional Modifications to Spectra Prior to Input into The Cannon

Upon downloading[10] all GALAH DR3 spectra that satisfy the conditions enumerated in Section 2, we interpolate each star's flux and flux error array over the shared wavelength grid recommended by the GALAH team.[11] We perform several tests to assess how manipulating the spectra affects our resulting label precision and find that truncating the spectra to only include the first three CCDs ($\lambda \leqslant 6730$) increases the performance of The Cannon. This increase in performance from neglecting the last CCD is likely due to the strong spikes in the redmost spectral segment that do not originate in the stellar photosphere and make fitting a model to that spectral region difficult. These spectral spikes in the redmost CCD are likely due to imperfect telluric subtraction, a common challenge in the spectral reduction of near-infrared spectra (e.g., Griffith et al. 2022). We are able to neglect this final CCD without a loss of precision because most of the spectral lines associated with our sampled elements lie in the first three CCDs. The only exception to this is O: by removing the fourth CCD, we lose access to the only available O lines in GALAH. Previous works using The Cannon have successfully recovered O abundances without using O lines (Ting et al. 2018), so we proceed with inferring O, but we *do not* consider it in our subsequent measurement of the doppelganger rate. Additionally, we remove three spectral segments containing strong diffuse interstellar bands (e.g., Vogrinčič et al. 2023) near
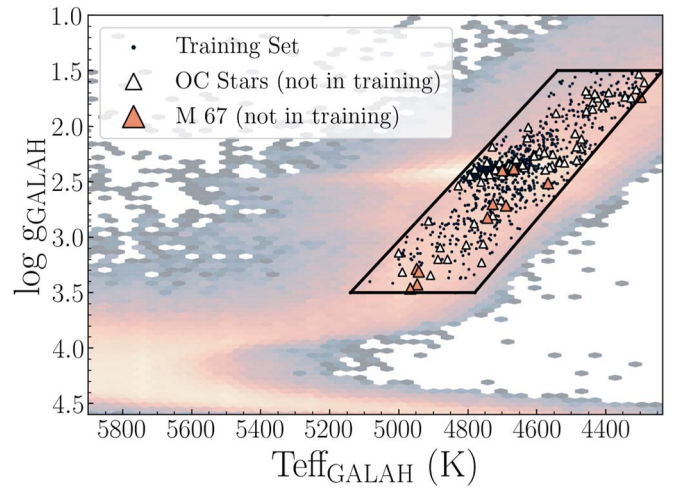

Figure 1. Kiel diagram (GALAH-reported log g vs. $T_{\rm eff}$) for the full GALAH sample (background distribution). Our parameter space of investigation, which was designed to contain red giant and red clump stars, is encapsulated within the black polygon. Black dots mark stars in our training set. We include open cluster stars as triangles to illustrate their parameter space coverage, with filled orange triangles highlighting members of chemically homogeneous open cluster M67. All background stars that fall within the black polygon and satisfy our quality cuts (Section 2) are considered in our analysis.

$\lambda 5798$, $\lambda 5871$, and $\lambda 6614$. These features are caused by interstellar dust absorbing the star's light and are not intrinsic to the stellar photosphere. Thus, to ensure that The Cannon does not use these features for inference, we remove them.

3.1.4. Training Set

When building our high-fidelity training set, we experiment with various selections to identify one that allows for high output abundance precision while also sampling sufficient open cluster stars. Our final choice of training set is limited to stars that satisfy the following requirements, in addition to those mentioned in Section 2:

  i   `flag_X_fe = 0`
  ii   `chi2_sp < 10`
  iii   `snr_c2_iraf > 100`.

The first requirement makes use of an element-specific flag that controls for unreliable GALAH-measured abundances in specific elements. The second requirement reports the χ^2 fit of the best-fitting SME model to the GALAH data. Previous works, such as Nandakumar et al. (2022), have required that `chi2_sp < 4` when building a GALAH-based training set. We found that increasing our requirement to `chi2_sp < 10` enables us to achieve the same precision while allowing for better sampling of high-metallicity ([Fe/H] > 0) red giant stars. The final requirement ensures that we are training The Cannon on high-S/N data to maximize its ability to learn real correlations between spectral pixels and labels as opposed to learning from noise.

In Figure 1, we present a Kiel (log g versus $T_{\rm eff}$) diagram of our final training set (black dots) consisting of 956 stars that sample the surface area of our chosen parameter space, which we define by the polygon in Figure 1. We note that our chosen polygon is not parallel to the red giant branch. This means that we are unable to sample cooler, metal-rich giants beyond [Fe/H] = 0.2. This is out of necessity: there are not enough high-S/N stars with unflagged abundances to extend our training
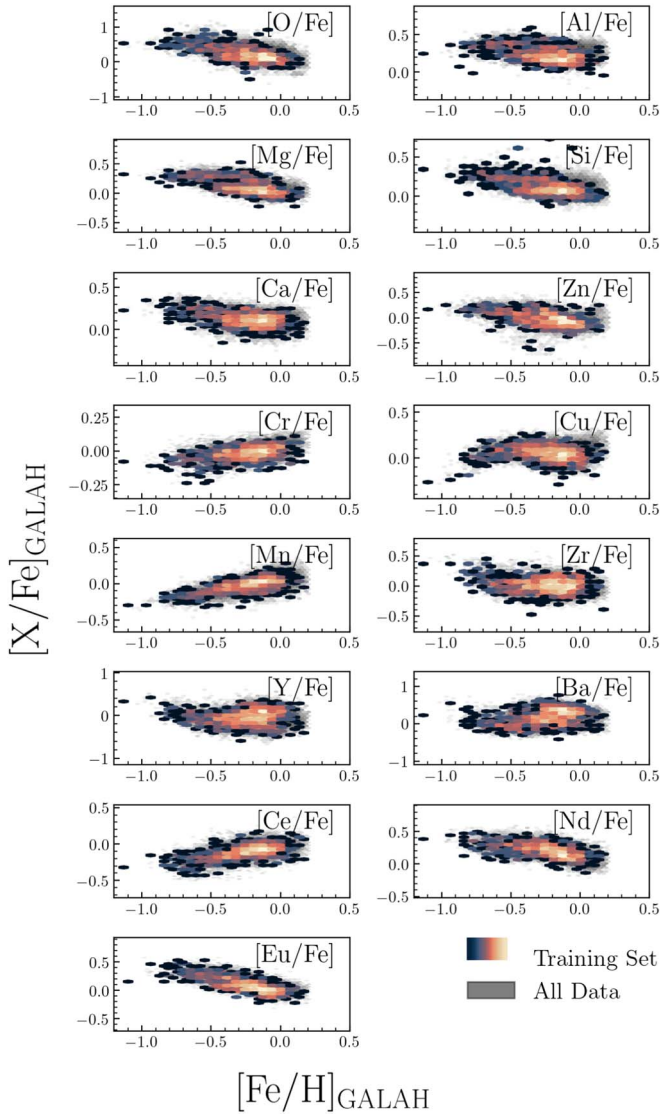
Figure 2. Density distributions in [X/Fe] vs. [Fe/H] for our training sample (color) atop our full sample (background gray). Note that the color maps are colored by logarithmic stellar density. We ensure that the training set spans the parameter space of our full sample to ensure reliable output *Cannon* abundances.

sample to this region of the Kiel diagram. However, this is not a problem, as we conduct our entire analysis in the chosen polygon, and the polygon is well-sampled by the training set. To illustrate the parameter space sampled by our open cluster stars, we include them in the Figure as open triangles, with M67, a chemically homogeneous open cluster (e.g., Bovy 2016; Ness et al. 2018; Poovelil et al. 2020), as orange filled triangles. We note that these open cluster stars are not in our training set. The background distribution represents the full GALAH data set, and all background stars that fall within the polygon and satisfy our quality cuts are re-analyzed using *The Cannon*. In Figure 2, we plot [X/Fe] versus [Fe/H] distributions for our training set stars atop the equivalent for our full sample, highlighting that the surface area of the [X/Fe] versus [Fe/H] distribution of our full sample is fully covered by our training set. This is important for ensuring that the model need not extrapolate when inferring stellar abundances.

We assess the ability for our model to recover the GALAH labels of the training set by performing a series of ten leave-10%-out cross-validation tests. This involves training our model on 90% of the training data and assessing its ability to recover the GALAH-reported labels of the remaining 10% of the training data. In Figure 3, we plot the *Cannon*-recovered label as a function of input GALAH label for all stars in our training set, marking the one-to-one line for reference. The model is successful in recovering the input training data labels to high precision, with recovered labels agreeing with GALAH-reported labels within 0.04–0.08 dex for most elements. Exceptions to this are O, Zr, and Y, which we recover to within 0.11–0.14 dex. We note that this cross validation is an assessment of the fidelity with which we can determine the reference labels, but as it includes the GALAH label uncertainties, it is not a measurement of the internal precision of *The Cannon* on these data (see Section 3.1.5). As a measure of robustness, we assess the fit of the output *Cannon* model spectra to the input GALAH spectrum, both globally and around strong lines of measured elements, via a χ^2 goodness-of-fit metric that considers GALAH flux uncertainties. To determine goodness-of-fit to specific lines, we adopt the GALAH SME line masks presented in Buder et al. (2018). We flag and subsequently ignore all stars with global χ^2 and line-specific χ^2 values that exceed two times the degrees of freedom (e.g., spectral pixels). In Figure 4, we show an example fit to a typical open cluster star to illustrate the quality of *The Cannon*'s model spectra fits.

When using data-driven algorithms such as *The Cannon* to measure abundances from the full spectral range without the use of censoring, a procedure where *The Cannon* is only allowed to learn abundance information from specified strong lines of each element, it is important to acknowledge that the model can infer abundances using correlations between spectral features not directly associated with the element. As mentioned in Section 3.1.1, in some cases, there are physical reasons for the existence of these correlations, as changes in a star's atmospheric composition in an element can influence the spectral behavior of other elemental lines or continuum regions (e.g., Ting et al. 2018). However, in other cases, this can lead to abundance inferences of certain elements that are instead primarily driven by a nonphysical correlation that is introduced by the training set data. In the context of our high-resolution spectra, we expect the primary abundance information for each element to be learned from strong, known lines of the element. We conduct three tests to confirm this. We first inspect the first-order *Cannon* model coefficients, which describe the direct quadratic relationship between each spectral pixel and each label (see Appendix A). In Figures 15 and 16, we plot the first-order *Cannon* coefficients for each label as a function of wavelength. We mark strong, known lines of each element with a red dashed line and confirm that *The Cannon* is drawing its primary abundance information from those line regions. Next, in Figure 5, we repeat this exercise by plotting the median spectrum of our sample and highlighting in orange spectral regions that correspond to the strongest 1% of first-order coefficients for a selection of five elements. We mark known strong lines of each element with a thick orange line. These two tests make evident that the primary abundance information retrieved by *The Cannon* is coming from strong lines of each element, though it is also clear that *The Cannon* is leveraging
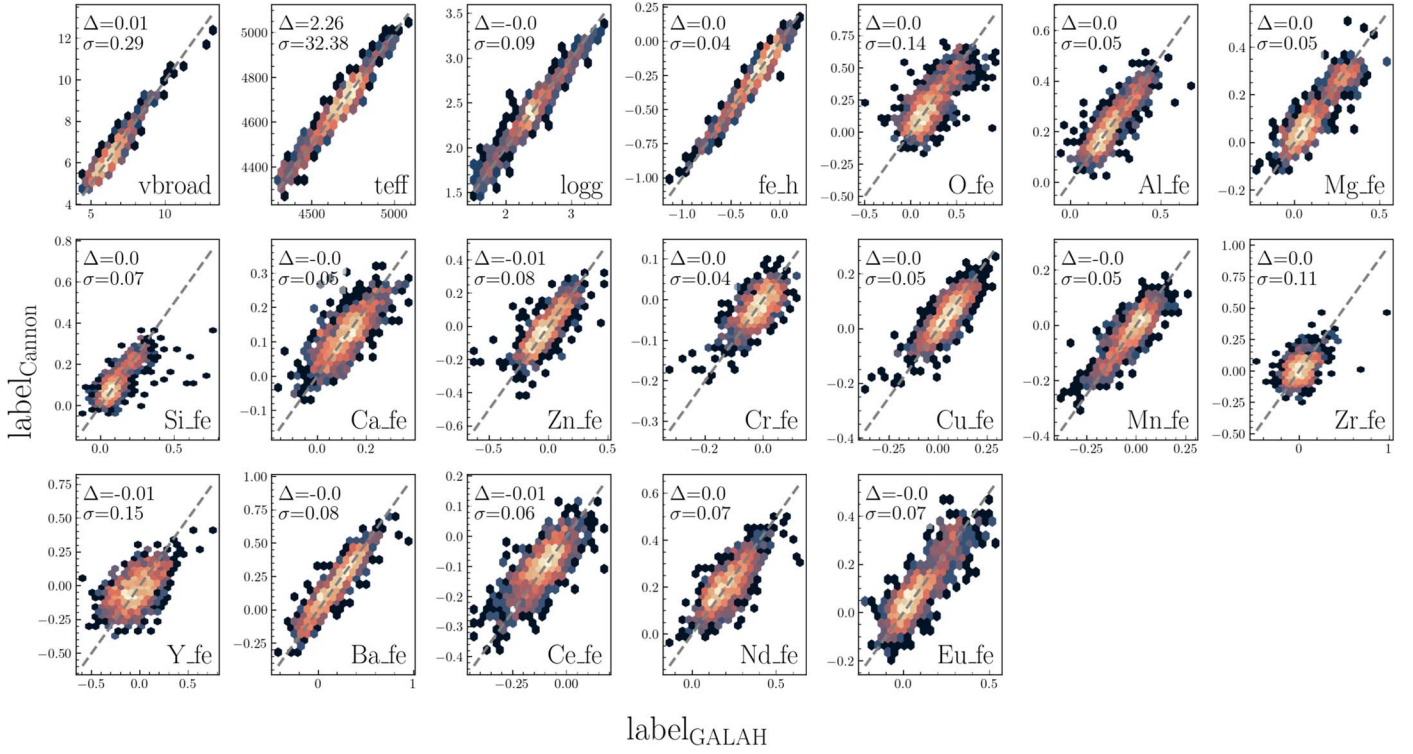
Figure 3. Combined results of our ten leave-10%-out cross-validation tests, where we plot output *Cannon* label as a function of input GALAH DR3 label for all 19 labels-of-interest in our 956 star training set. The color map corresponds to point density, and the dashed line represents the one-to-one line. The offset from (Δ) and scatter (σ) around the one-to-one line is printed in the corner of each panel. Our model is able to retrieve the GALAH-reported abundances of our training set within 0.07 dex for the majority of elements.

Figure 4. A comparison of our *Cannon* model fit (red) to the spectrum of a star in M67 (black). The smaller top panels highlight segments of the spectra corresponding to strong lines of each element. Bottom panel shows a larger cutout of the spectrum. It is evident that *The Cannon* is capable of fitting the GALAH data well, both globally and around lines of interest.

the full spectrum to extract abundance information. This is by design and allows *The Cannon* to achieve its enhanced precision. Finally, we perform a censoring test to assess the impact of censoring on our *Cannon*-derived abundances and final results. We find that applying censoring does not qualitatively affect our final results, and we include a detailed description of this test and its results in Appendix B.

3.1.5. Abundance Uncertainties from The Cannon

To determine our final abundance uncertainties, we must take into account two sources of uncertainty. The first is that reported directly by *The Cannon*, which reflects the dispersion in the final likelihood function for each label. The second is the external, systematic uncertainty that is best parameterized as a
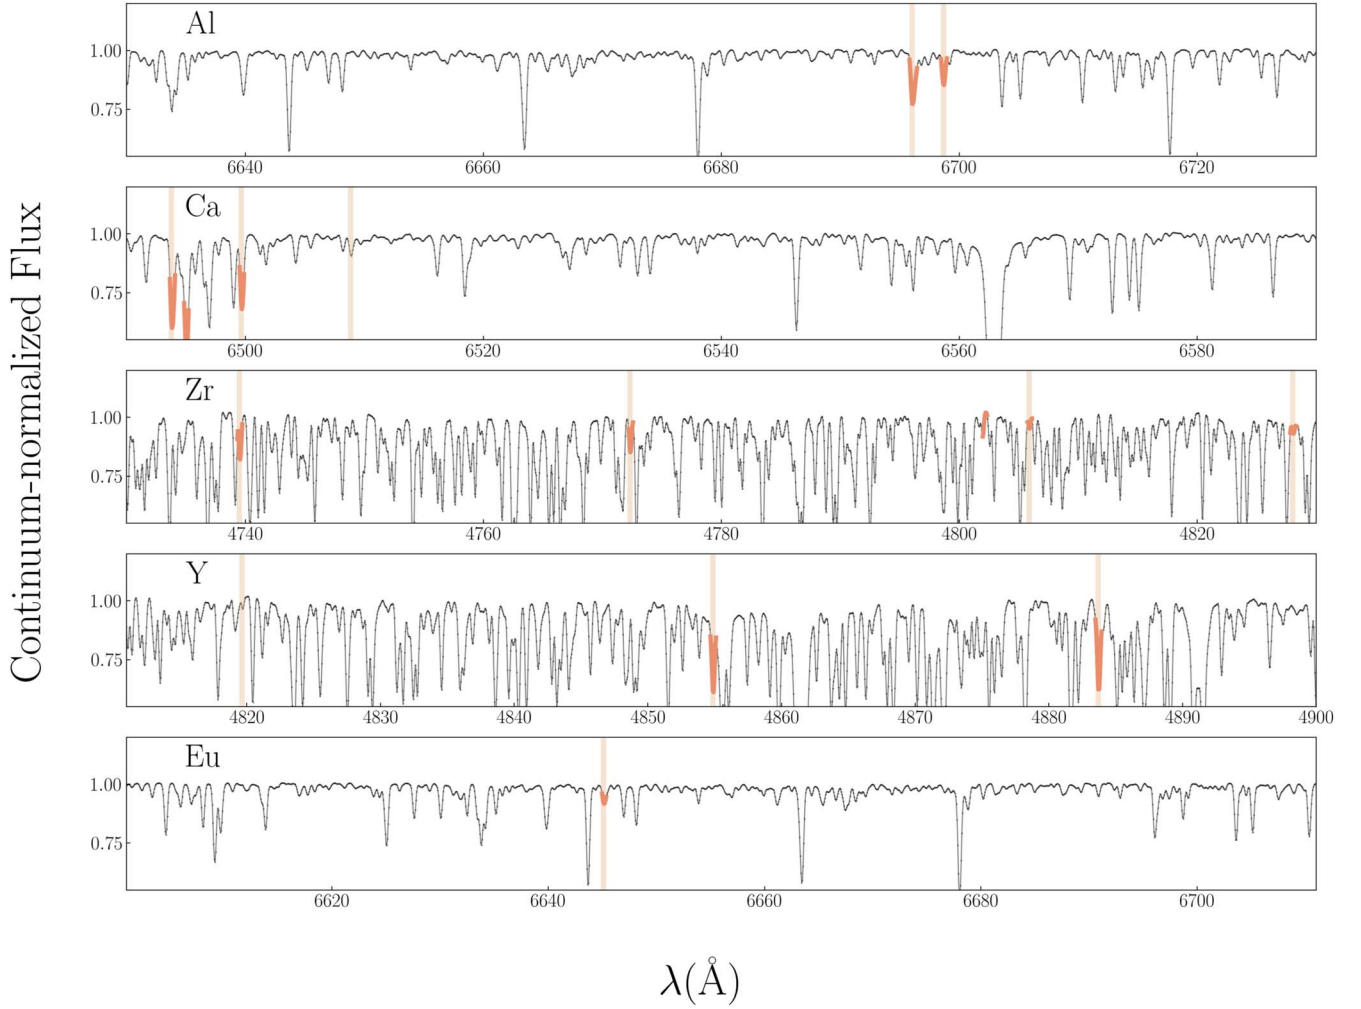
Figure 5. Spectral windows, showing the continuum-normalized flux as a function of wavelength, for the median stellar spectrum in our sample (black) containing regions of absorption lines of Al, Ca, Zr, Y, and Eu, from top to bottom. We highlight in orange spectral regions that correspond with the 1% largest first-order *Cannon* model coefficients in the labeled element. The black spectrum is the base *Cannon* spectrum, which represents the median spectrum of the full data set. Light orange vertical lines correspond to known lines of the element, and thin black vertical lines correspond to known lines of Fe. We note that, even in the most-difficult-to-measure elements such as Zr, the *Cannon* model successfully draws its primary abundance information from the relevant element's lines. See Appendix B Figure 15 for the full array of first-order coefficients as a function of wavelength for each label.

function of spectral S/N (e.g., Ness et al. 2015; Wheeler et al. 2020; Nandakumar et al. 2022). To determine our model's systematic precision, we make use of repeat-visit spectra, spectra taken of the same object and later coadded before being measured for the final main GALAH catalog. Repeat-visit spectra present the opportunity to test our model's stability as a function of S/N. In an ideal case, when running spectra of the same source but with different S/Ns through *The Cannon*, our model should always return the same labels regardless of the S/N of the spectrum. Thus, any change in the model's inferred labels between spectra of the same source but at different S/Ns can quantify our S/N-dependent label uncertainty. For this test, we download the GALAH DR3 `all_spec` catalog, which reports stellar parameters for each individual observed spectrum. We then identify stars with more than one observation by filtering for repeated values in the `dr3_source_id` column. We then download eight nights of data that have significant numbers of targets with repeat-visit spectra (150427, 150428, 150429, 150430, 170912, 170911, 170910, and 170909) and only consider the 387 targets that span the parameter space of our larger data set. We then produce several

instances of each source's spectrum at various S/Ns, starting first with single spectra populating the lowest-S/N bins, followed by coadded versions of the spectra. For example, if a source has three total observations, we are able to produce three low-, three medium-, and one high-S/N version of its spectra for the purpose of this experiment. We then run the spectra through *The Cannon* and, taking the labels reported for the highest-S/N spectrum as "truth," measure the dispersion in the difference in the inferred labels between the low- and high-S/N spectra as a function of S/N. As reported in Nandakumar et al. (2022), our precision increases with S/N exponentially and plateaus beyond an S/N of 40. We fit exponential functions to describe the relationship between S/N and label recovery precision and adopt the S/N-dependent dispersions as the external precision of our inferred abundances.

We compute our final label uncertainties by taking the quadratic sum of the internal model uncertainties reported by *The Cannon* and the external uncertainties from our S/N experiment. In Figure 6, we compare our repeat-visit abundance dispersion (black triangles) with those of GALAH repeat-visit results (gray circles) and GALAH-reported
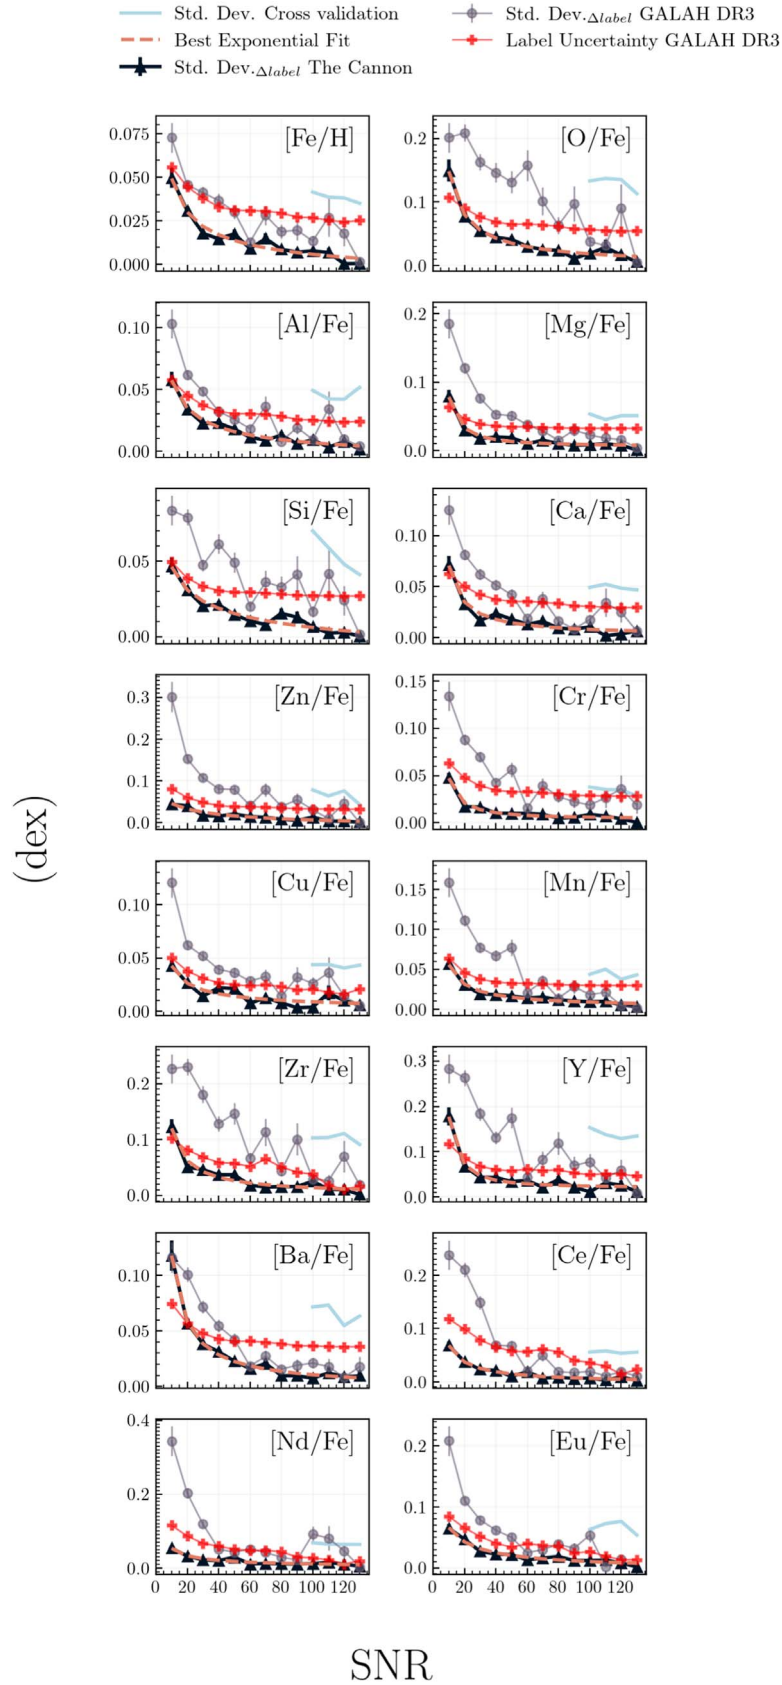
Figure 6. Results of our repeat-visit spectrum investigation into the model's precision for each of the 16 elemental abundance ratios that we re-derive in this work. Solid black triangles present the standard deviation in the difference between the reported *Cannon* label between low-S/N spectra and the highest-S/N spectrum as a function of S/N for 387 objects with repeat observations. An exponential fit to this relationship, which is our final adopted precision, is shown as a dashed orange line. Gray circles represent the equivalent of the black triangles except for the GALAH-reported abundances. Red crosses display the mean uncertainty as a function of S/N. We mark the standard deviation from our cross-validation test (Figure 3) in light blue. This figure illustrates the enhanced precision achieved by *The Cannon*, highlighted by the difference in the red and black curves.

Figure 7. Distributions of the *Cannon*-measured elemental abundances of our final sample. We mark the mean GALAH-reported error ($\mu_{\mathrm{err},G}$) in each abundance with a dashed navy line and the new mean *Cannon*-reported abundance error (μ_{err}, C) in lighter purple. For the majority of elements, *The Cannon* achieves either comparable or lower abundance uncertainties compared to the reported GALAH uncertainties.

Figure 8. [Fe/H] as a function of T_{eff} for all 122 open cluster stars in our sample. Background shows the full sample's distribution in this plane. The top panel presents the GALAH DR3 abundances, and the bottom panel presents the *Cannon* results. The 86 stars with S/N > 40 and unflagged abundances in all elements (see Section 3.1.4) serve as the reference point for the chemical homogeneity of stars born together, an ingredient in our measurement of the doppelganger rate.

Cannon is very high. In Figure 7, we show the abundance distribution of our sample (black histogram) and compare it to *The Cannon*'s mean abundance uncertainty (solid purple line) and the original GALAH DR3 uncertainty (dashed navy line). Elements such as Zr, Y, and Ce, for example, previously had uncertainties comparable to the width of the abundance distribution of the full sample. Thus, any potential distinguishing power in these elements was thwarted by the GALAH precision. With our *Cannon*-enhanced abundance precision, these elements can now be used to potentially distinguish between doppelgangers.

3.2. Final Catalog

Our final catalog consists of 28,120 stars with newly inferred values of v_broad, T_{eff}, log g, [Fe/H], and [X/Fe] for O, Al, Mg, Si, Ca, Cr, Cu, Zn, Mn, Zr, Y, Ba, Ce, Nd, and Eu for each star that populates the polygon in Figure 1. In Figure 8, we present the [Fe/H] versus T_{eff} distributions for the 14 open clusters in our sample, showing the GALAH-reported distribution in the top panel and the *Cannon*-inferred distribution in the bottom panel. The table schema for our final catalog is included in Appendix B Table 1 and the full table is available. As mentioned in Section 3.1.4, we flag as unreliable all stars with global χ^2 and individual element χ^2 goodness-of-fit values that exceed two times the degrees of freedom of the spectrum or relevant line mask. We hereafter only consider stars with unflagged global and individual element abundances.

To assess the chemical homogeneity of the open clusters in our sample in light of our re-derived abundances, we draw all possible intracluster pairs with $\Delta T_{\mathrm{eff}} < 100$ K, $\Delta\log g < 0.1$ dex, and unflagged abundances in all elements and present χ^2 distributions (see Equation (1)) for all intracluster pairs in each element in Figure 9, with median and mean values marked with a dashed and solid line, respectively. If all clusters in our sample

uncertainties (red crosses) as a function of S/N. Our resulting abundance precision is improved relative to the GALAH-reported precision by up to a factor of 3. We find that for most elements, the dispersion in the difference between GALAH-reported labels and our *Cannon*-inferred labels from our cross-validation (solid gray line) is similar to the GALAH uncertainty, indicating that the internal precision of *The*
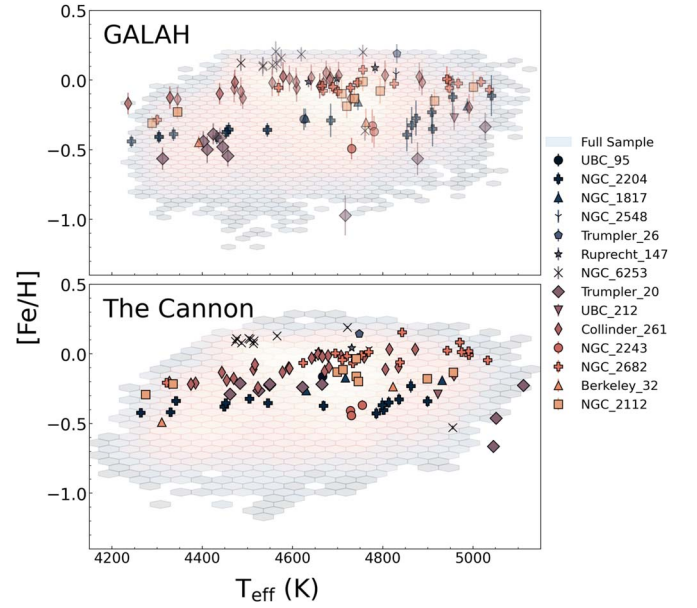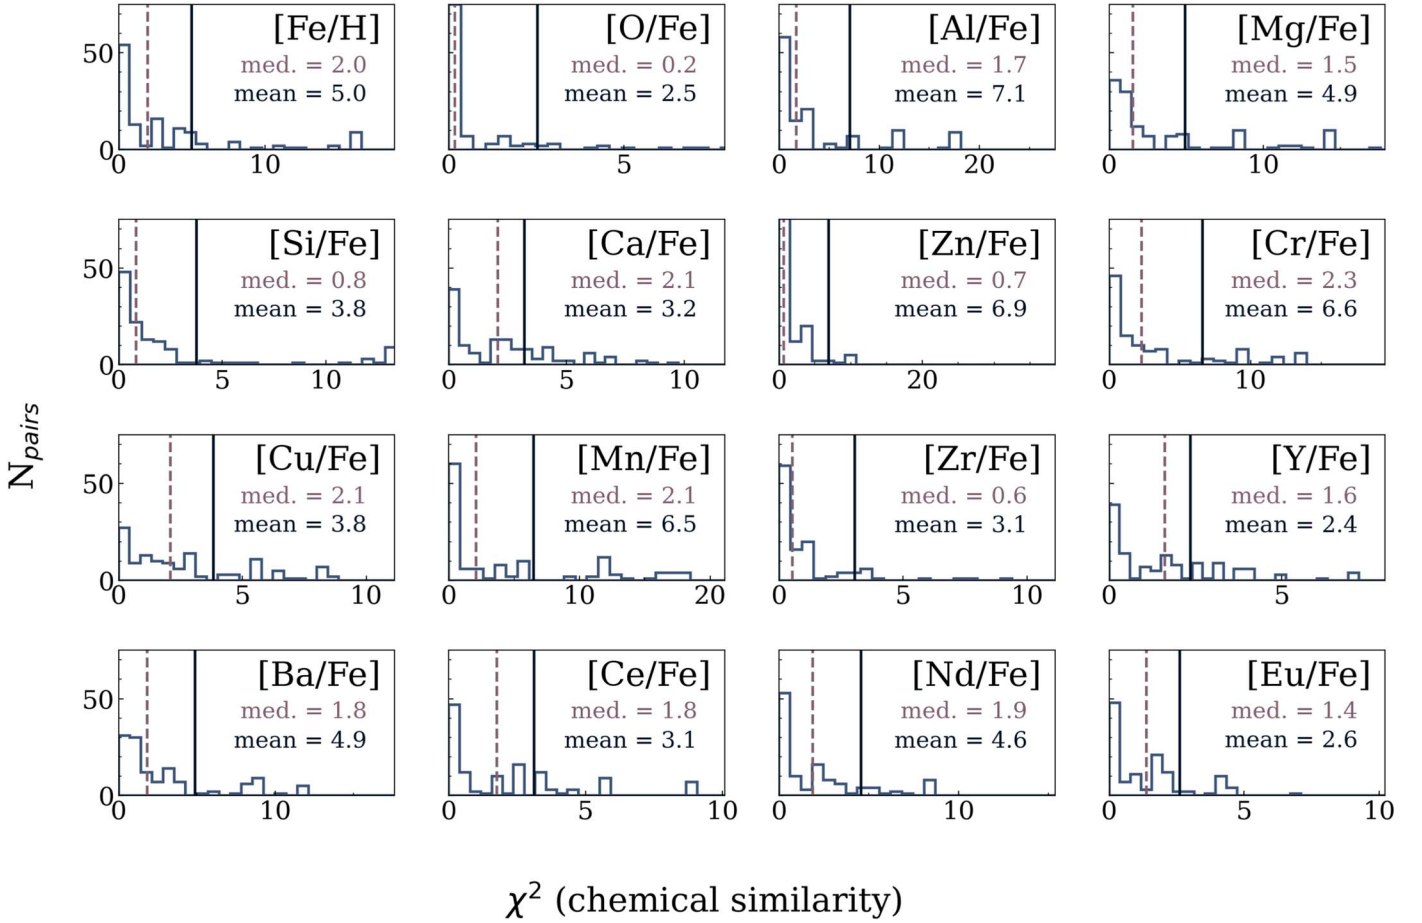
Figure 9. Distributions of χ^2 values for 122 intracluster pairs with $\Delta T_{\rm eff} < 100$ K, $\Delta \log g < 0.1$ in each element. The median and mean χ^2 values are marked by purple dashed and solid navy lines, respectively. If all sampled clusters were perfectly chemically homogeneous in an element, we would expect median and mean χ^2 values of 0.45 and 1, respectively (Ness et al. 2018). While most intracluster pairs appear chemically indistinguishable within measurement uncertainties ($\chi^2 \lesssim 1$), several pairs show high χ^2 values (>3) indicative of significant abundance differences.

were perfectly chemically homogeneous, we would expect the median and mean χ^2 values to be 0.45 and 1, respectively, for each element (Ness et al. 2018). The majority of intracluster pairs have χ^2 values near 1, indicating that cluster stars are chemically indistinguishable at the current precision level. However, several intracluster pairs have χ^2 values much greater than 1, suggesting that some clusters have abundance dispersions exceeding the abundance uncertainty.

3.3. Measuring the Doppelganger Rate in GALAH DR3

Instead of measuring the doppelganger rate in our full sample, we measure it in a higher-quality subset of our data. This is because the doppelganger rate is sensitive to abundance precision and our choice of open cluster reference pairs. In general, in the low-S/N regime, where abundance uncertainties are high, pairs of stars will tend to look more chemically similar within uncertainties. We find that snr_c2_iraf > 40 is an ideal S/N lower limit for measuring a meaningful doppelganger rate in our sample, as it ensures we maximize our abundance precision. This S/N cutoff enables us to sample 47 stars across five clusters (NGC 2112, NGC 6253, NGC 2204, Collinder 261, and M67) and 13,375 stars in the field.

To build our reference sample of open cluster stars, we draw all possible combinations of intracluster pairs where stars in the pair have $\Delta T_{\rm eff} < 100$ K and $\Delta \log g < 0.1$ dex. The $\Delta T_{\rm eff}$ and

$\Delta \log g$ requirement ensures that we avoid potential systematic trends between abundance and $T_{\rm eff}$, $\log g$ to artificially enhance or minimize the abundance similarity of stars in a pair. We ultimately build a population of 122 intracluster pairs that serve as a reference point for the chemical similarity of stars born together in open clusters. For our field sample, we draw one million unique pairs of stars that are not members of Spina et al. (2021) open clusters and that satisfy the same $\Delta T_{\rm eff}$, $\Delta \log g$ requirements as those for open cluster pairs. These field pairs serve to sample the chemical diversity (or lack thereof) of the phase-mixed Galactic disk population.

After building our intracluster and field pair samples, we measure the doppelganger rate using the method of Ness et al. (2018), computing a global χ^2 value using Equation (1) for each pair of stars and defining doppelgangers to be field pairs with χ^2 less than the median χ^2 value of intracluster pairs in the considered elements. As in Ness et al. (2018), we use Fe with respect to H ([Fe/H]) and the remaining elements with respect to Fe ([X/Fe]). Thus, our results are heavily influenced by the [Fe/H] similarity of the stars in the field pair. When drawing intracluster pairs, the [Fe/H] abundances of stars will likely be similar. Hence, when randomly sampling the entire field, our doppelganger rate is driven primarily by the likelihood of drawing two stars with similar [Fe/H], which depends on the [Fe/H] distribution of our sample. For this reason, we measure the doppelganger rate twice: once for the entire field and again
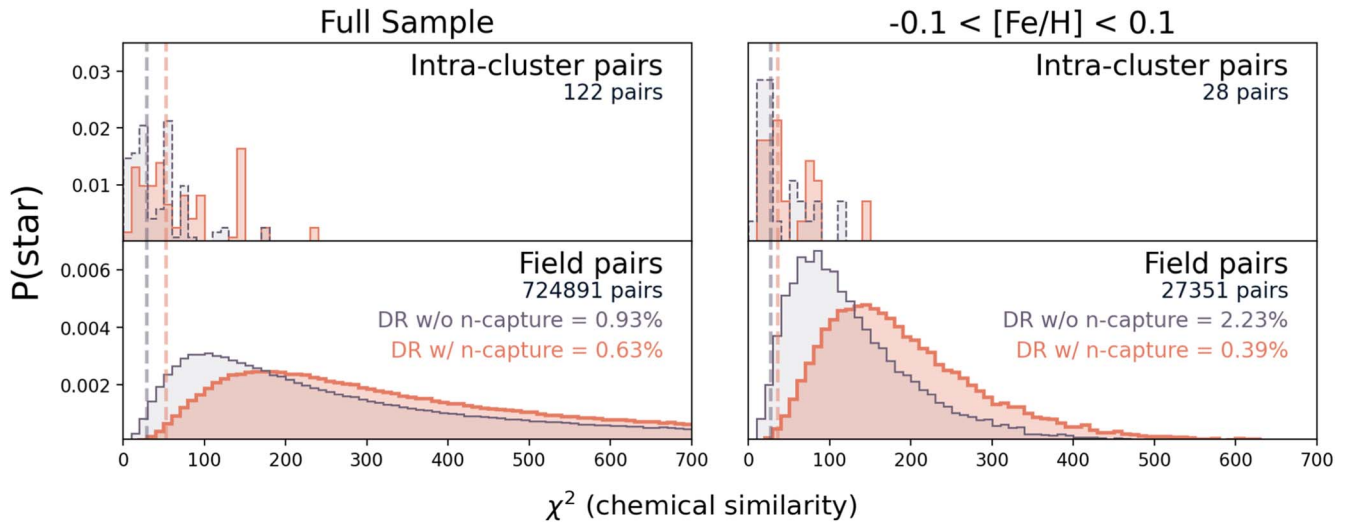
Figure 10. χ^2 (see Equation (1)) distributions for intracluster pairs (top panels) and those drawn from the field (bottom panels). Gray histograms present the χ^2 distributions when considering the light (here, Al), α (Mg, Si, Ca), and iron-peak (Fe, Cr, Cu, Zn, Mn) elements only. The orange histograms present the χ^2 distributions when considering the aforementioned elements as well as the neutron-capture elements (Zr, Y, Ba, Ce, Nd, Eu). Vertical lines mark the median χ^2 value for open cluster pairs without (gray) and with (orange) the consideration of neutron-capture elements. Left panels contain the results for our full S/N > 40 sample while the right panels contain results for stars within the $-0.1 < [Fe/H] < 0.1$ range. The addition of neutron-capture elements reduces the doppelganger rate by a factor of 1/3 for all field pairs and nearly 6 for field pairs in the narrow [Fe/H] range.

for a narrow [Fe/H] range to control for the dependence of the doppelganger rate on the [Fe/H] distribution of our sample.

4. Results

4.1. Impact of Neutron-capture Elements on the DR

Our primary question in this work asks whether the neutron-capture elements affect the doppelganger rate. In Figure 10, we use Equation (1) (considering [Fe/H] and [X/Fe]; see Section 3.3) to determine the χ^2 distributions, a proxy for chemical similarity, of intracluster pairs (top panels) and field pairs (bottom panels). We compare the χ^2 distribution of the field pairs with that of the intracluster pairs when excluding (gray) and including (orange) the neutron-capture elements. Faint dashed lines correspond to the median χ^2 value for intracluster pairs when excluding (gray) and including (orange) neutron-capture elements. The left panel presents the aforementioned for the full S/N > 40 sample, while the right panel presents it for a subset of stars with $-0.1 < [Fe/H] < 0.1$. As found in Ness et al. (2018), stars born together in open clusters are far more chemically similar than random field pairs. However, there exist field star pairs that are just as, if not more, chemically alike than intracluster pairs, as quantitatively shown in their χ^2 values. These field pairs are deemed doppelgangers. As expected, adding neutron-capture elements (Zr, Y, Ba, Ce, Nd, and Eu) shifts the χ^2 distributions for each population due to adding degrees of freedom to the χ^2 calculation. However, the χ^2 distribution of open cluster pairs shifts comparatively *less* than that of random field pairs. If neutron-capture elements had no further distinguishing power compared to the lighter elements, then the doppelganger rate would remain constant when including them. However, including neutron-capture elements decreases the doppelganger rate from 0.9% to 0.6% in our full sample and from 2.2% to 0.4% in our sample with a narrow [Fe/H] range. Figure 10 illustrates that the neutron-capture elements have a subtle but nonnegligible affect on the doppelganger rate.

4.2. Which Elements Matter Most in Distinguishing Disk Stars?

In Figure 11, we present the impact of each successive elemental family on the doppelganger rate via a *cumulative doppelganger rate* (CDR) computed in the narrower $-0.1 < [Fe/H] < 0.1$ range. In the top panel, we report the measured CDR as a function of each elemental family and the elemental families lighter than it. For example, when we plot the CDR associated with the iron-peak elements, we plot the percentage of pairs that have χ^2 values computed using Fe, Cr, Cu, Mn, and Zn less than the median of intracluster pairs. Next, for the light elements, we plot the same as for the iron-peak elements but this time also considering Al in our χ^2 calculations. Next, for α elements, we plot the same but considering Al, Fe, Cr, Cu, Mn, Zn, Mg, Ca, and Si. This continues on as one moves rightward, ultimately terminating with the s-process elements. In the bottom panel, we illustrate the practical discriminating power of each elemental family by reporting the multiplicative factor with which the CDR changes upon the addition of each new family. We find that once doppelgangers are identified using the iron-peak elements, the s-process elements possess the greatest additional distinguishing power, followed by the α-elements, the light elements, and the r-process elements. We note that this narrow [Fe/H] range contains primarily low-α disk stars, and this may play a role in the relatively low distinguishing power of the α-elements. As in Figure 10, we find that the doppelganger rate considering light, α, and iron-peak elements is 2.23% for stars with $-0.1 < [Fe/H] < 0.1$. With the addition of the neutron-capture elements (Eu, Zr, Y, Ba, Ce, and Nd), the doppelganger rate reduces to 0.39%, by a factor of 5.75 (with r-process element Eu reducing it to 1.73% and s-process elements reducing it to the final 0.39%).

4.3. Partial versus Complete Doppelgangers

Upon discovering that neutron-capture elements affect the doppelganger rate, we isolate pairs of field stars that are *partial* doppelgangers (that is, they satisfy doppelganger requirements
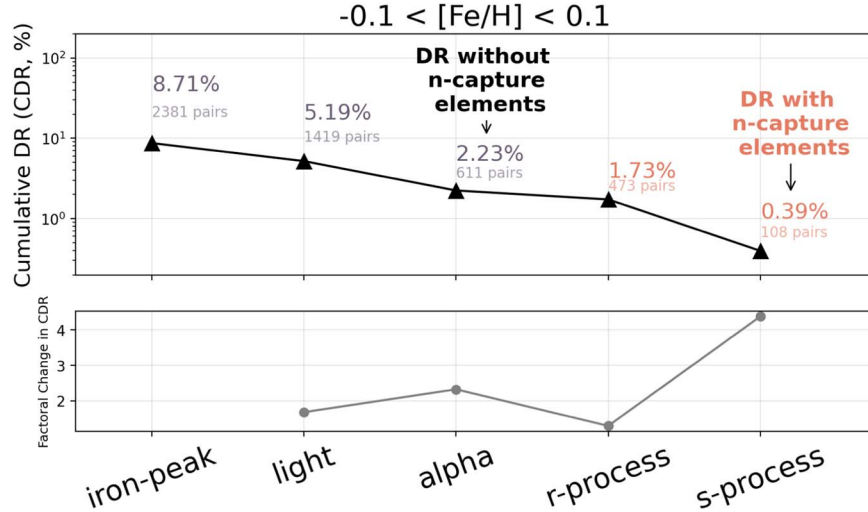
Figure 11. Top panel: the cumulative doppelganger rate (CDR) as a function of the addition of each new elemental family. The CDR represents the total doppelganger rate when considering all elemental elemental families successively. That is, on the far left, we determine the doppelganger rate considering just the iron-peak elements while on the far right, we determine the doppelganger rate with the light, α, iron-peak, r-, and s-process elements. The introduction of neutron-capture elements reduces the doppelganger rate by a factor of 5.75. It is possible that s-process elements have a greater influence on the CDR than the r-process elements, but we remind the reader that we only use one r-process element, Eu, but five s-process elements (Zr, Y, Ba, Ce, and Nd), so the comparison is not straightforward. Bottom panel: the factor with which the CDR changes upon adding each new elemental family. The y-axis reports the factor with which the doppelganger rate changes as a function of each added elemental family. The bottom panel illustrates that once pairs are selected by chemical similarity in iron-peak elements, the s-process elements possess the greatest additional distinguishing power, followed by the alpha elements, the light elements, and the r-process elements.


Figure 12. Violin plots showing the absolute difference in abundance (Δ [Fe/H] for Fe, Δ [X/Fe] for the remaining elements) between stars in doppelganger pairs for each element. The abundance difference distributions for partial doppelgangers (those that are doppelganger exclusively in light, α, and iron-peak elements) are represented by the empty violin plots. The same distributions for complete doppelgangers (doppelganger in all measured elements) are represented by the filled violin plots. The median abundance difference in each element for intracluster pairs is represented by the orange crosses. This figure illustrates that random pairs of field stars can appear as chemically similar as open cluster stars in the lighter elements but show strong deviations in the heavier elements.

in the light, α, and iron-peak elements) from pairs of field stars that are *complete* doppelgangers (that is, they satisfy doppelganger requirements in all elements). In Figure 12, we show the distributions in absolute difference in abundance for partial (open violin plots) and complete (filled violin plots) doppelgangers and also include as reference the median absolute difference in abundance for intracluster pairs (yellow crosses). We note that this plot reports abundance differences between stars in a pair, not χ^2 values, which we use in the actual computation of the doppelganger rate. This figure illustrates that there exist pairs of field stars that are as chemically similar in the light, α, and iron-peak elements as stars born together but deviate in the neutron-capture elements by up to 0.6 dex. In Figure 13, we show example spectra of partial doppelgangers (black) and complete doppelgangers (red). Note that despite both pairs of stars satisfying doppelganger requirements in the lighter elements, the partial doppelganger spectra show deviations in the neutron-capture lines (highlighted in faint orange) while the complete doppelgangers do not. We note here that since neutron-capture elements can distinguish partial from complete doppelgangers, the results of Ness et al. (2018) may overestimate the doppelganger rate due to the exclusion of

this elemental family. We explore this possibility in C. Manea et al. (2024, in preparation).

It is of physical interest to investigate why some pairs are partial doppelgangers while others are complete doppelgangers. Adopting the dynamical parameters and age estimates from the associated GALAH value-added catalogs (see Section 2), we compare the similarity in these characteristics for pairs of stars that are partial versus complete doppelgangers. In Figure 14, we plot the running mean (left column) and standard deviation (right column) in the (absolute for the left column) difference in age (top panel) and a series of dynamical characteristics (J_R, L_Z, J_Z, eccentricity, $z_{\max}$, and energy) for pairs as a function of their chemical similarity, which we measure using a reduced χ^2, defined by the formula in Equation (1) but further divided by the number of degrees of freedom (e.g., number of elements considered, which is 9 when ignoring the neutron-capture elements, marked by the gray curves, and 15 when including them, marked by the orange curves.) Dashed lines represent the median reduced χ^2 value for intracluster pairs. Shading represents the uncertainty, which for the left panels is $+/-\sigma/(N)^{1/2}$ and for the right panel is $+/-\sigma/(2N)^{1/2}$ where N is the number of stars in the bin. χ^2 correlates strongly with similarity in age or dynamical characteristics. However, we do

Figure 13. Example spectra of the two types of doppelganger pairs: partial doppelgangers (black) and complete doppelgangers (red). In each panel, we highlight spectral regions around neutron-capture element lines, with the faint orange line marking the line core of the labeled element. Note how the complete doppelganger spectra match in the highlighted neutron-capture features while the partial doppelgangers show deviations, as expected given our definitions of these two populations.

not find clear evidence that complete doppelgangers (orange curve leftward of the dashed lines) possess more similar age and dynamical parameters relative to partial doppelgangers (gray curve leftward of the dashed lines) using the χ^2 metric. However, these results should be followed up using higher-precision ages and dynamical characteristics as any potential signature may be concealed beneath large uncertainties in age and orbital properties. As mentioned in Section 2, our choice to use giant branch and clump stars in this analysis means that ages in particular are poorly constrained. Furthermore, our definition of χ^2, which is a sum of chemical similarity in all elements, may also complicate this analysis. Treating each element individually may yield different results, but we leave this for a future investigation.

5. Discussion

Our investigations into the doppelganger rate in GALAH DR3 conclude that the neutron-capture elements possess subtle but nonnegligible distinguishing power in the Milky Way disk that is important to consider when chemically characterizing stars. As captured in Figures 10 and 11, these results suggest that the neutron-capture elements contain information distinct

from that of the lighter elements and thus add to the dimensionality of Milky Way abundance space. For example, when ignoring the neutron-capture elements, we identify over 6700 pairs of apparently unrelated field stars that are as chemically similar as stars in open clusters. However, when introducing the neutron-capture elements, we find that one-third of those pairs are not in fact as chemically similar as stars in open clusters when considering the neutron-capture elements. When restricting ourselves to a narrow [Fe/H] range ($-0.1 <$ [Fe/H] < 0.1), we identify over 600 pairs of doppelgangers when omitting the neutron-capture elements, and upon introducing the neutron-capture elements, we are left with just $\sim$108 pairs (a reduction by a factor of 5.75). These results directly support those of Griffith et al. (2022), who found that introducing a chemical dimension corresponding to AGB star nucleosynthesis is required to reproduce the chemical distribution of the GALAH survey. Additionally, these results support those of Lambert & Reddy (2016), which found that open clusters deviate most significantly in the neutron-capture elements even when sharing light ($Z < 30$) element compositions, suggesting that neutron-capture elements contain additional distinguishing power not captured by the lighter elements.
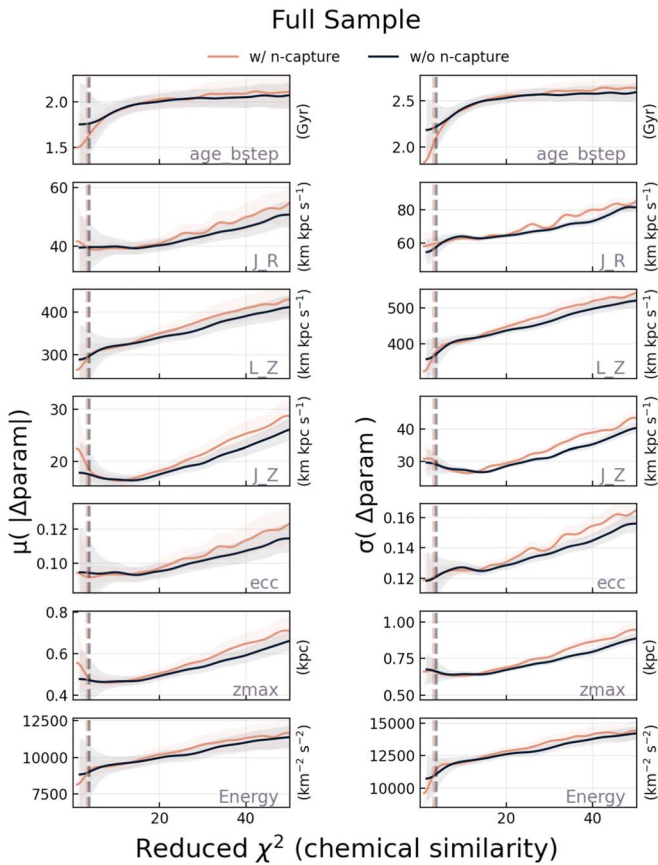
Figure 14. Running mean (left column) and standard deviation (right column) in (absolute, for left column) age and dynamical property difference for random pairs of field stars as a function of their reduced χ^2, a measure of their chemical similarity normalized by the number of elements considered (i.e., the degrees of freedom). The gray curve omits the neutron-capture elements while the orange curve includes the neutron-capture elements, and the thick dashed line represents the median reduced χ^2 for intracluster pairs. While chemical similarity globally correlates with age or dynamical similarity, complete doppelgangers (orange curve that lies leftward of the dashed line) do not appear to be more similar in age and dynamical parameters than partial doppelgangers (gray curve that lies leftward of dashed line) using our defined χ^2 metric.

5.1. Possible Physical Origins of Enhanced Distinguishing Power of Neutron-capture Elements

There are potential physical explanations for why neutron-capture elements could add to the chemical dimensionality of Milky Way stars. Their unique production and dispersal sites may embed neutron-capture elemental abundances with temporal and spatial information that is distinct from that contained in the lighter ($Z < 30$) elements. For example, several works have identified that s-process element abundances, when compared to α-element abundances, can be effective chemical clocks that probe stellar age both within and outside the Milky Way (e.g., Feltzing et al. 2017; Skúladóttir et al. 2019; Ratcliffe et al. 2023). This is due to their primary nucleosynthetic origins in AGB stars, which are lower mass ($M < 8$–$10 \, M_\odot$) than the high-mass ($M > 10 \, M_\odot$) stars that produce the α elements (e.g., Karakas & Lattanzio 2014). Because stellar mass strongly influences stellar lifetime, s-process elements have a longer delay time for enrichment in the Galaxy with respect to the α-elements. When stars form from gas that is enriched in products of both AGB star and core collapse supernova nucleosynthesis, this delay-time difference enables s-process elements to trace the star's age when compared to α elements, though the

relationship between the [s-process/alpha] ratio and stellar age is neither universal nor simple (e.g., Casali et al. 2020). In addition to age information, s-process elements may also add information about stellar birth position that is not captured by supernova-produced elements. Simulations suggest that the s-process elements have shorter correlation lengths in the Milky Way interstellar medium relative to the lighter elements due to the nature of their dispersal (e.g., Armillotta et al. 2018; Krumholz & Ting 2018; Emerick et al. 2020). The s-process elements are dispersed via relatively gentle, localized AGB star winds, contrary to the highly energetic dispersal of the lighter elements via supernovae (e.g., Cox et al. 2012). The more localized nature of their dispersal may thus allow for greater variation in s-process composition as a function of location within the interstellar medium. Thus, stars born together in pockets of this enriched gas will be chemically similar in these heavier elements, but stars across different birth groups, even at fixed metallicity, may differ in their s-process composition. We indeed see evidence of this when studying open clusters (e.g., Lambert & Reddy 2016). This could imply that either (a) the timescale between the enrichment of s-process products into the interstellar medium and the formation of stars from this material is shorter than the timescale with which the interstellar medium mixes away gas-phase abundance variations, or (b) the interstellar medium is less efficient at mixing s-process products (again potentially due to their less-energetic dispersal compared to supernova products; Krumholz & Ting 2018; Emerick et al. 2020).

The r-process elements also have physical reasons to add a unique dimension to Milky Way abundance space. The r-process elements are believed to be formed in stochastic events such as magnetorotational supernovae (e.g., Siegel & Metzger 2017; Halevi & Mösta 2018; Siegel et al. 2019) and compact object mergers (e.g., Korobkin et al. 2012). The stochastic nature of their synthesis could allow them to carry additional information about stellar birth position or age. We note that the origin of the r-process elements is still a major open question in the field (e.g., Kobayashi et al. 2023; Lian et al. 2023).

The results of our analysis of the CDR (Figure 11) appear to suggest that the r-process elements hold less distinguishing power than the s-process elements. We caution that Eu is the only pure r-process element in our analysis that represents this nucleosynthetic family, whereas we have five elements representing the s-process family, one of which also has significant r-process contribution (Nd; e.g., Kobayashi et al. 2020). It is thus possible that much of the distinguishing power of r-process elements is overshadowed by the numerous s-process elements we sample. However, the s-process elements possessing greater distinguishing power than the r-process elements also have physical support. It is believed that r-process elements are dispersed via energetic supernovae, and this may cause them to behave differently than s-process elements in the interstellar medium. Further work into the dependence of the doppelganger rate on r-process elements must be conducted to clarify whether they truly carry less distinguishing power than s-process elements, or if this is just a consequence of the choice of elements considered.

Given the neutron-capture elements that we consider in this work, the doppelganger rate reduces from $\sim 2.2\%$ to 0.4%. While this is a substantial drop in the doppelganger rate, and informative for nucleosynthetic sources and mixing, this still

prohibitive for the prospect of strong chemical tagging in the Milky Way. This rate, which measures the probability of which random stars are chemically as similar as stars born together, is still a factor of $\approx$1000–10000 times greater than the expected rate of recovering true birth pairs in a disk. This is assuming that clusters form with a typical mass of $1 \times 10^4 \, M_\odot$–$1 \times 10^6 \, M_\odot$ (e.g., Bland-Hawthorn et al. 2010).

5.2. Analysis Limitations and Looking Ahead to GALAH DR4

We emphasize that the doppelganger rate is influenced by several factors, some intrinsic to the Milky Way (e.g., mixing efficiency in the interstellar medium, variations in nucleosynthetic yields) and some caused by the available data. The results of this work are limited by the precision of derived elemental abundances. With increased abundance precision, we may find that the doppelganger rate decreases even further. However, the results of this work suggest that at the abundance precision of GALAH DR3 combined with *The Cannon*, we are able to harness the distinguishing power of neutron-capture elements. We note that if we repeat this experiment using the provided GALAH DR3 abundances of a high-S/N (S/N > 100), high-precision subsample of the catalog, we obtain the same qualitative results. This work is also limited by the number of elements that we consider in our analysis. For the purpose of building a reasonably sizeable training set, we could not consider all $\sim$30 elements reported by GALAH, as it is rare for stars to have high-fidelity abundance measurements in all elements. As such, we only considered between one and five elements from each nucleosynthetic family.

The upcoming fourth data release of GALAH is an ideal environment for expanding on this experiment. GALAH DR4 will have a larger sample size overall and enhanced sampling of open cluster stars. Enhancing the reference open cluster sample would allow for a more granular exploration of the doppelganger rate as a function of metallicity. In this work, we are limited by the small number of open cluster stars sampled by GALAH DR3, and separating the data into metallicity bins would make the number of reference open cluster stars per metallicity bin prohibitively small. In DR4, the number of sampled open cluster stars will be doubled, and the sampled open clusters will span $-2 <$ [Fe/H] < 1, whereas DR3 only reasonably samples open clusters $-0.5 <$ [Fe/H] < 0.5. Furthermore, GALAH DR4 will have improved abundance precision, further minimizing its extrinsic effect on the measured doppelganger rate. Finally, GALAH DR4 will provide a larger set of stars sampled in a wider range of elements, enabling the consideration of a greater number of elements and thus allowing for a result that better reflects the intrinsic doppelganger rate of our Galaxy.

6. Conclusions

In this work, we measure the doppelganger rate among red clump and red giant stars in GALAH DR3. The doppelganger rate measures the rate at which randomly drawn pairs of apparently unrelated field stars appear to be as chemically similar as stars born together. It probes the chemical diversity of Milky Way stars, the chemical dimensionality of Milky Way abundance space, and the complexity with which Galactic chemical evolution operates. After re-deriving stellar parameters and abundances with *The Cannon*, we measure the chemical doppelganger rate. We find that 0.9% of random pairs

of fields stars are doppelgangers in the light-, α-, and iron-peak elements. This number increases to $\sim$2.2% when we restrict ourselves to stars in the $-0.1 <$ [Fe/H] < 0.1 dex range. However, we find that including neutron-capture elements Zr, Y, Ba, Ce, Nd, and Eu in our analysis decreases the doppelganger rate significantly. When considering our full sample, the neutron-capture elements reduce the doppelganger rate to 0.6%, and when restricting to the $-0.1 <$ [Fe/H] < 0.1 dex range, the doppelganger rate drops to 0.4%, by nearly a factor of 6 relative to the rate measured considering only the lighter ($Z < 30$) elements. In other words, up to 85% of stars that are highly chemically similar in the lighter elements deviate in their neutron-capture element abundances. Chemical similarity strongly correlates with similarity in age or dynamics. However, we do not identify any clear signatures that complete doppelgangers, pairs of stars that are doppelgangers in the light, α, iron-peak, *and* neutron-capture elements, are more similar in age and dynamical characteristics than partial doppelgangers, those that are doppelgangers in the lighter elements but show deviations in the neutron-capture elements. However, these results are not conclusive, so additional follow-up work should be done to further explore this. Finally, our results suggest that the *s*-process elements may carry greater distinguishing power in our sample than the *r*-process elements, though we urge additional follow-up to confirm this.

This work highlights that the neutron-capture elements carry unique information that is distinct from that found in the light-, α-, and iron-peak elements and thus constitutes important tools in the chemical characterization of Milky Way stars. Despite their enhanced distinguishing power, our final doppelganger rate of $\sim$0.4% suggests that neutron-capture elements measured at the precision of this work are likely not sufficient to satisfy the requirements for strong chemical tagging. However, our results illustrate that neutron-capture elements can distinguish between 85% of stars that appear chemically similar in the light, α, and iron-peak elements, suggesting that these heavy elements, particularly the *s*-process elements, are potentially important tools for the weak chemical tagging of stars to known clusters and stellar populations. Our results motivate the need for continued work to improve atomic data for the heavy elements and enhance our ability to extract precise and accurate neutron-capture elemental abundances from stellar spectra.

Data Availability

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 work has also made use of GALAH DR3, based on data acquired through the Australian Astronomical Observatory, under programs: A/2013B/13 (The GALAH pilot survey); A/2014A/25, A/2015A/19, A2017A/18 (The GALAH survey).

We acknowledge the traditional owners of the land on which the AAT stands, the Gamilaraay people. We pay our respects to elders past and present. The GALAH DR3 data underlying this work are available in the Data Central at https://cloud.datacentral.org.au/teamdata/GALAH/public/GALAH_DR3/ and can be accessed with the unique identifier galah_dr3 for this release and sobject_id for each spectrum.

Appendix A

We present the first-order coefficients of our *Cannon* model as a function of wavelength in Figures 15 and 16. These Figures illustrate the importance of each pixel when inferring various stellar parameters and abundances. Each panel corresponds to a different label, and the dashed red and gray lines mark strong lines of the element (where applicable) and Balmer lines. While *The Cannon* uses the entire spectrum to infer labels, strong lines of each element tend to have the largest coefficient values, suggesting that our model draws label information primarily from the expected strong lines.

Table 1 presents the table schema for our final catalog, which contains *Cannon*-re-derived, higher-precision parameters and abundances for 28,120 giant stars sampled by GALAH DR3. We emphasize that these abundances are not necessarily accurate in the absolute sense. Instead, they are accurate in the *relative* sense: stars with the same spectra have the same labels. We do not apply any global abundance offsets, as is common in surveys (e.g., Buder et al. 2021; Abdurro'uf et al. 2022), to
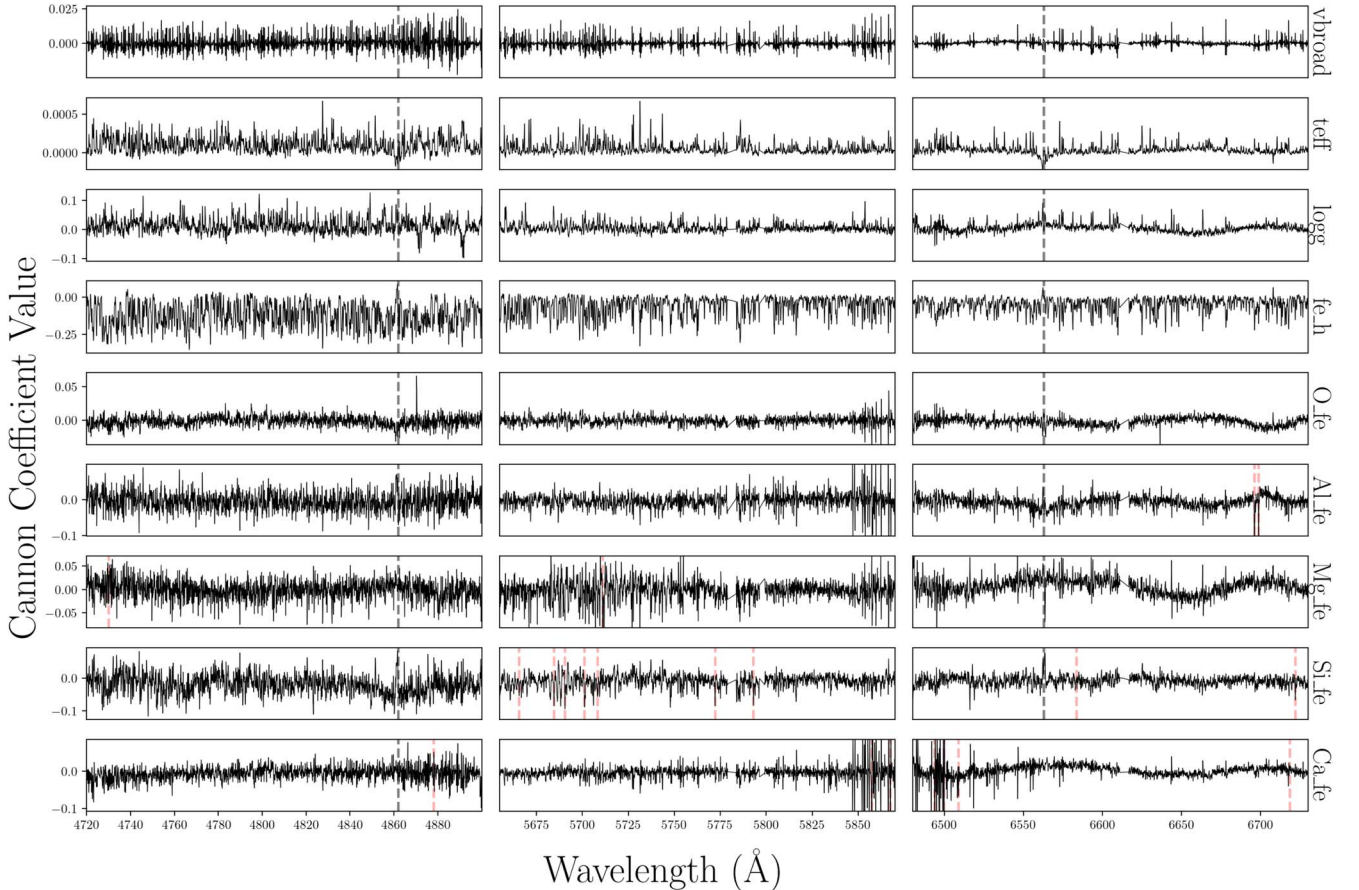

Figure 15. First-order *Cannon* model coefficients for the first nine labels as a function of wavelength. The dashed gray lines mark Hβ and Hα. The dashed red lines mark known strong lines of the relevant element corresponding to each panel. Note how the model is using the entire spectrum to derive abundance information, which includes both strong lines and the continuum, as opposed to just a subset of known strong lines. See the next page for the remaining eight labels.
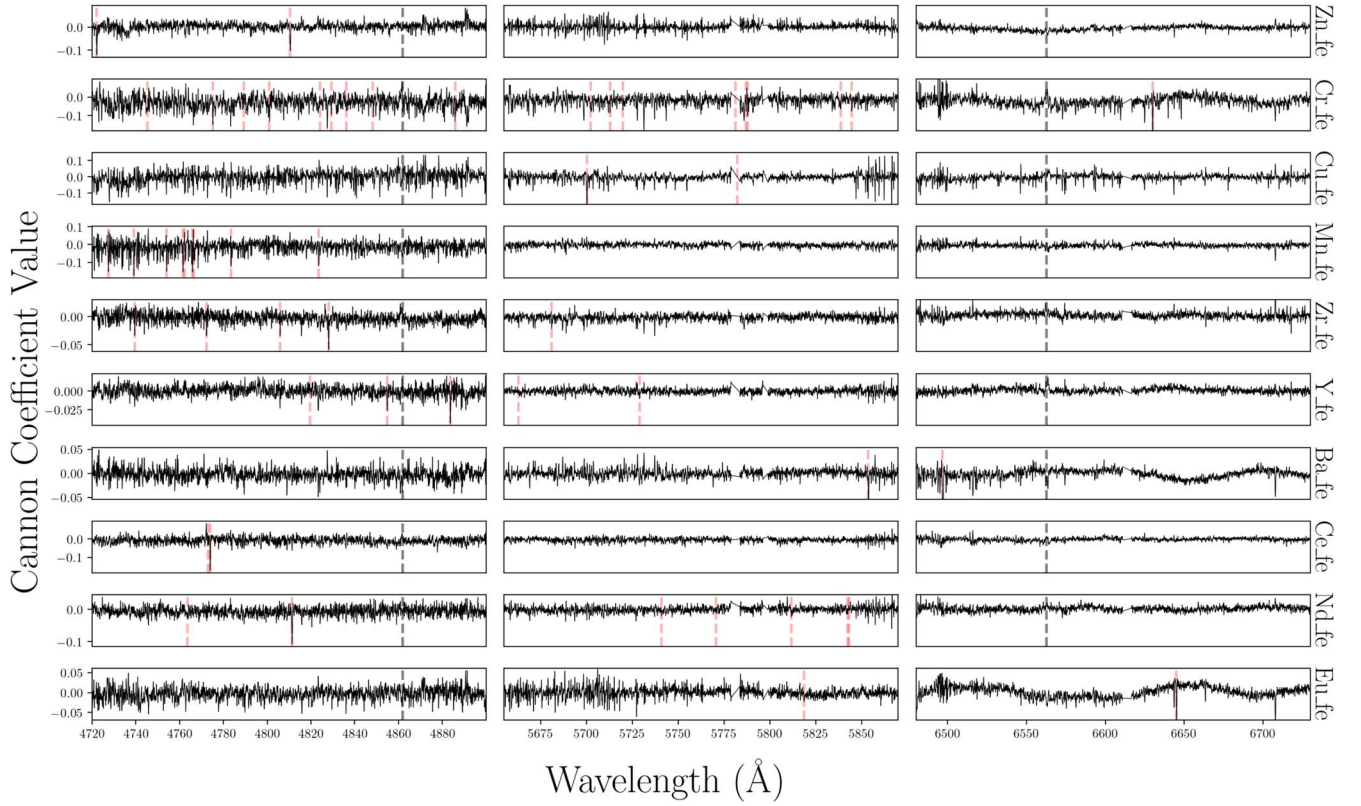
Figure 16. First-order *Cannon* model coefficients for the remaining eight labels. See the previous page for more details.

shift the scales of abundances. In a practical sense, this catalog is useful for relative abundance comparisons, such as in studies of chemical homogeneity, but not for absolute comparison to theoretical yield models unless the comparison is relative in nature. When using this catalog for most typical applications, we highly recommend only considering abundances with `chisq` and `X_fe_cannon_chisq < 2`. However, stars with `chisq > 2` and abundances with `X_fe_cannon_chisq > 2` may be useful for identifying stars with unusual spectral features or anomalous abundances.

Appendix B
Censoring Test

As described in Section 3.1.4, when running *The Cannon* without censoring, it is important to run several tests confirming that the resulting abundances are informed by absorption lines associated with the element as opposed to underlying non-physical trends in the training data. While we confirm that *The Cannon* draws its primary abundance information from strong lines of each element, it is of interest to explore the impact of censoring on our final results.

We run *The Cannon* with censoring, adopting the GALAH SME line masks presented in (Buder et al. 2018) for all elements except for [Fe/H], which we allow to be derived from the full spectrum. As expected, our *Cannon*-derived abundance uncertainties are higher with censoring, though they remain lower than the GALAH-reported uncertainties for all elements. We then measure the doppelganger rate as before and find that in general, the doppelganger rate increases with censoring. This is expected given that larger abundance uncertainties can cause a greater percentage of stars to qualify as doppelgangers. Qualitatively, the main results remain the same: the doppelganger rate reduces when including the neutron-capture elements. For the full sample, the neutron-capture elements have a similar impact on the doppelganger rate as in our non-censoring run. For the narrow sample $(-0.1 < [\text{Fe/H}] < 0.1)$, the neutron-capture elements have a smaller impact on the doppelganger rate than before, cutting it in half as opposed to by a factor of 5. Overall, we find that the neutron-capture elements have a detectable effect on the doppelganger rate whether or not we apply censoring, suggesting that our results are robust against our approach to using *The Cannon*. For reader interest, we provide our *Cannon*-derived abundances both without and with censoring applied in Table 1.

Table 1
Schema for Our *Cannon*-inferred Parameters and Abundances for the 28,120 Stars in Our Sample

Column	Type	Unit	Description
GaiaDR3_source_id	int	...	Gaia DR3 source identifier (adopted from GALAH DR3)
sobject_id	int	...	Unique GALAH identifier (adopted from GALAH DR3)
chisq	float	...	Global χ^2 fit of *The Cannon* model to GALAH spectrum
vbroad_cannon	float	km s^{-1}	*The Cannon*-reported broadening velocity
teff_cannon	float	K	*The Cannon*-inferred T_{eff}
logg_cannon	float	...	*The Cannon*-inferred log g
fe_h_cannon	float	...	*The Cannon*-inferred [Fe/H]
e_fe_h_cannon	float	...	Final *The Cannon* [Fe/H] uncertainty
X_fe_cannon	float	...	*The Cannon*-inferred [X/Fe] (where X = O, Al, Mg, Ca, Si, Zn, Cu, Cr, Mn, Zr, Y, Ba, Ce, Nd, or Eu)
e_X_fe_cannon	float	...	Final *The Cannon* [X/Fe] uncertainty (where X = O, Al, Mg, Ca, Si, Zn, Cu, Cr, Mn, Zr, Y, Ba, Ce, Nd, or Eu)
X_fe_cannon_chisq	float	...	Element-specific median χ^2 fit of *The Cannon* model to GALAH spectrum around strong lines of element

Note. The full table can be found on Zenodo: doi:10.5281/zenodo.10864201. We include two data tables: manea_et_al_2023.fits and manea_et_al_2023_censoring.fits. The former contains the Cannon results used throughout this work, which were computed *without* censoring (see Section 3.1.4). The latter file is provided solely for reader interest and contains our Cannon results *with* censoring.

ORCID iDs

Catherine Manea ⓘ https://orcid.org/0000-0002-0900-6076
Keith Hawkins ⓘ https://orcid.org/0000-0002-1423-2174
Melissa K. Ness ⓘ https://orcid.org/0000-0001-5082-6693
Sven Buder ⓘ https://orcid.org/0000-0002-4031-8553
Sarah L. Martell ⓘ https://orcid.org/0000-0002-3430-4163
Daniel B. Zucker ⓘ https://orcid.org/0000-0003-1124-8477

References

Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, ApJS, 259, 35
Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, ApJS, 249, 3
Armillotta, L., Krumholz, M. R., & Fujimoto, Y. 2018, MNRAS, 481, 5000
Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Bedell, M., Bean, J. L., Meléndez, J., et al. 2018, ApJ, 865, 68
Beeson, K. L., Kos, J., de Grijs, R., et al. 2024, MNRAS, 529, 2483
Bell, R. A., Eriksson, K., Gustafsson, B., & Nordlund, A. 1976, A&AS, 23, 37
Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, MNRAS, 496, 1922
Bland-Hawthorn, J., Krumholz, M. R., & Freeman, K. 2010, ApJ, 713, 166
Bovy, J. 2015, ApJS, 216, 29
Bovy, J. 2016, ApJ, 817, 49
Bovy, J., Bahmanyar, A., Fritz, T. K., & Kallivayalil, N. 2016, ApJ, 833, 31
Brown, A. G. A., Vallenari, A., Prusti, T., et al. 2021, A&A, 649, A1
Buder, S., Asplund, M., Duong, L., et al. 2018, MNRAS, 478, 4513
Buder, S., Sharma, S., Kos, J., et al. 2021, MNRAS, 506, 150
Cantat-Gaudin, T., Jordi, C., Vallenari, A., et al. 2018, A&A, 618, A93
Casali, G., Spina, L., Magrini, L., et al. 2020, A&A, 639, A127
Cheng, C. M., Price-Jones, N., & Bovy, J. 2021, MNRAS, 506, 5573
Conroy, C., Naidu, R. P., Zaritsky, D., et al. 2019, ApJ, 887, 237
Cox, N. L. J., Kerschbaum, F., van Marle, A. J., et al. 2012, A&A, 537, A35
Crameri, F. 2023, Scientific Colour Maps v. 8.0.1, Zenodo, doi:10.5281/zenodo.1243862
Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, RAA, 12, 1197
de Mijolla, D., Ness, M. K., Viti, S., & Wheeler, A. J. 2021, ApJ, 913, 12
De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, MNRAS, 449, 2604
Emerick, A., Bryan, G. L., & Mac Low, M.-M. 2020, ApJ, 890, 155
Feltzing, S., Howes, L. M., McMillan, P. J., & Stonkutė, E. 2017, MNRAS, 465, L109
Fischer, D. A., & Valenti, J. 2005, ApJ, 622, 1102
Freeman, K., & Bland-Hawthorn, J. 2002, ARA&A, 40, 487
Gao, X., Lind, K., Amarsi, A. M., et al. 2018, MNRAS, 481, 2666

Gilmore, G., Randich, S., Worley, C. C., et al. 2022, A&A, 666, A120
Griffith, E. J., Hogg, D. W., Dalcanton, J. J., et al. 2024, AJ, 167, 98
Griffith, E. J., Weinberg, D. H., Buder, S., et al. 2022, ApJ, 931, 23
Gustafsson, B., Bell, R. A., Eriksson, K., & Nordlund, A. 1975, A&A, 500, 67
Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, A&A, 486, 951
Halevi, G., & Mösta, P. 2018, MNRAS, 477, 2366
Hawkins, K., Lucey, M., Ting, Y.-S., et al. 2020, MNRAS, 492, 1164
Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, ApJ, 836, 5
Holtzman, J. A., Hasselquist, S., Shetrone, M., et al. 2018, AJ, 156, 125
Hunter, J. D. 2007, CSE, 9, 90
Jofré, P., Heiter, U., & Soubiran, C. 2019, ARA&A, 57, 571
Jofré, P., Heiter, U., Worley, C. C., et al. 2017, A&A, 601, A38
Karakas, A. I., & Lattanzio, J. C. 2014, PASA, 31, e030
Kobayashi, C., Karakas, A. I., & Lugaro, M. 2020, ApJ, 900, 179
Kobayashi, C., Mandel, I., Belczynski, K., et al. 2023, ApJL, 943, L12
Korobkin, O., Rosswog, S., Arcones, A., & Winteler, C. 2012, MNRAS, 426, 1940
Krumholz, M. R., & Ting, Y.-S. 2018, MNRAS, 475, 2236
Lambert, D. L., & Reddy, A. B. S. 2016, ApJ, 831, 202
Lian, J., Storm, N., Guiglion, G., et al. 2023, MNRAS, 525, 1329
Liu, F., Asplund, M., Yong, D., et al. 2016, MNRAS, 463, 696
Liu, F., Asplund, M., Yong, D., et al. 2019, A&A, 627, A117
Marigo, P., Girardi, L., Bressan, A., et al. 2017, ApJ, 835, 77
Nandakumar, G., Hayden, M. R., Sharma, S., et al. 2022, MNRAS, 513, 232
Nelson, T., Ting, Y.-S., Hawkins, K., et al. 2021, ApJ, 921, 118
Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015, ApJ, 808, 16
Ness, M., Rix, H. W., Hogg, D. W., et al. 2018, ApJ, 853, 198
Ness, M. K., Johnston, K. V., Blancato, K., et al. 2019, ApJ, 883, 177
Ness, M. K., Wheeler, A. J., McKinnon, K., et al. 2022, ApJ, 926, 144
Nielsen, J., Gent, M. R., Bergemann, M., Eitner, P., & Johansen, A. 2023, A&A, 678, A74
Piskunov, N., & Valenti, J. A. 2017, A&A, 597, A16
Poovelil, V. J., Zasowski, G., Hasselquist, S., et al. 2020, ApJ, 903, 55
Price-Jones, N., Bovy, J., Webb, J. J., et al. 2020, MNRAS, 496, 5101
Ratcliffe, B., Minchev, I., Anders, F., et al. 2023, MNRAS, 525, 2208
Reback, J., McKinney, W., Jbrockmendel, et al. 2020, pandas-dev/pandas: Pandas v1.0.3, Zenodo, doi:10.5281/zenodo.3715232
Sharma, S., Hayden, M. R., Bland-Hawthorn, J., et al. 2022, MNRAS, 510, 734
Sharma, S., Stello, D., Buder, S., et al. 2018, MNRAS, 473, 2004
Siegel, D. M., Barnes, J., & Metzger, B. D. 2019, Natur, 569, 241
Siegel, D. M., & Metzger, B. D. 2017, PhRvL, 119, 231102
Skúladóttir, Á., Hansen, C. J., Salvadori, S., & Choplin, A. 2019, A&A, 631, A171
Spina, L., Ting, Y. S., De Silva, G. M., et al. 2021, MNRAS, 503, 3279
Steinmetz, M., Matijevič, G., Enke, H., et al. 2020, AJ, 160, 82
Souto, D., Cunha, K., Smith, V. V., et al. 2018, ApJ, 857, 14
Souto, D., Cunha, K., & Smith, V. V. 2021, ApJ, 917, 11

Taylor, M. B. 2005, in ASP Conf. Ser. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert (San Francisco, CA: ASP), 29

Ting, Y.-S., Conroy, C., Rix, H.-W., & Asplund, M. 2018, ApJ, 860, 159

Ting, Y.-S., Freeman, K. C., Kobayashi, C., De Silva, G. M., & Bland-Hawthorn, J. 2012, MNRAS, 421, 1231

Ting, Y.-S., & Weinberg, D. H. 2022, ApJ, 927, 209

Valenti, J. A., & Piskunov, N. 1996, A&AS, 118, 595

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261

Vogrinčič, R., Kos, J., Zwitter, T., et al. 2023, MNRAS, 521, 3727

Walsen, K., Jofré, P., Buder, S., et al. 2024, MNRAS, 529, 2946

Weinberg, D. H., Holtzman, J. A., Johnson, J. A., et al. 2022, ApJS, 260, 32

Wheeler, A., Ness, M., Buder, S., et al. 2020, ApJ, 898, 58