A simple MLP trained on LAMOST spectra transfers to DESI surprisingly well—and modest fine-tuning makes it better.

Draft version July 29, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

arXiv:2602.15021v2 [astro-ph.SR] 28 Jul 2026

Generalization from Low- to Moderate-Resolution Spectra with Neural Networks for Stellar Parameter Estimation: A Case Study with DESI

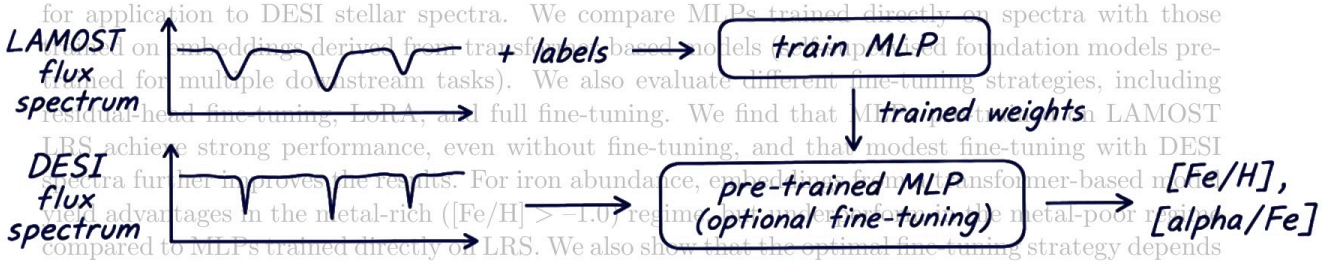
Xiaosheng Zhao [1], Yuan-Sen Ting [2,3,4], Rosemary F.G. Wyse [1], Alexander S. Szalay [1,5], Yang Huang [6,7], László Dobos [1,8], Tamás Budavári [9,1,5] and Viska Wei [1,5,9]

[1] *Department of Physics & Astronomy, The Johns Hopkins University, Baltimore, MD 21218, USA*
[2] *Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA*
[3] *Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA*
[4] *Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*
[5] *Department of Computer Science, The Johns Hopkins University, Baltimore, MD 21218, USA*
[6] *School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, People's Republic of China*
[7] *National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100012, People's Republic of China*
[8] *Department of Information Systems, Eötvös Loránd University, Budapest 1117, Hungary*
[9] *Department of Applied Mathematics & Statistics, Johns Hopkins University, Baltimore, MD 21218, USA*

ABSTRACT

Cross-survey generalization is a critical challenge in stellar spectral analysis, particularly in cases such as transfer from low- to moderate-resolution spectra. We investigate this problem using pre-trained models, focusing on simple neural networks such as multilayer perceptrons (MLPs), with a case study transferring from LAMOST low-resolution spectra (LRS) to DESI medium-resolution spectra (MRS). Specifically, we pre-train MLPs on either LRS or their embeddings and fine-tune them for application to DESI stellar spectra. We compare MLPs trained directly on spectra with those trained on embeddings derived from transformer-based models (spectral foundation models pre-trained for multiple downstream tasks). We also evaluate different fine-tuning strategies, including residual-head fine-tuning, LoRA, and full fine-tuning. We find that MLPs pre-trained on LAMOST LRS achieve strong performance, even without fine-tuning, and that modest fine-tuning with DESI spectra further improves the results. For iron abundance, embeddings from transformer-based models yield advantages in the metal-rich ([Fe/H] > -1.0) regime, but these models struggle in the metal-poor regime compared to MLPs trained directly on LRS. We also show that the optimal fine-tuning strategy depends on the specific stellar parameter under consideration. These results highlight that simple pre-trained MLPs can provide competitive cross-survey generalization, while the role of spectral foundation models for cross-survey stellar parameter estimation requires further exploration.

The challenge is keeping predictions reliable when moving between surveys.



A plain MLP gives foundation-model embeddings a meaningful baseline to beat.

Keywords: Galaxy stellar content (621), Fundamental parameters of stars (555), Astronomy data analysis (1858), Neural networks (1933)

1. INTRODUCTION

Large-scale spectroscopic surveys such as RAVE (M. Steinmetz et al. 2006), SEGUE (J. T. A. de Jong et al. 2010), APOGEE (S. R. Majewski et al. 2017), GALAH (G. M. De Silva et al. 2015), LAMOST (G. Zhao et al. 2012), and DESI (DESI Collaboration et al. 2016) have revolutionized our understanding of the Milky Way by delivering millions of stellar spectra. A central challenge is how to to generalize parameter-estimation pipelines effectively, i.e., maintain accuracy and consistency across heterogeneous datasets from surveys that differ in their wavelength coverage, spectral resolution, signal-to-noise ratio, and in the stellar populations they target (e.g., ranges of mass, evolutionary phase, and elemental abundances).

Conventional machine-learning models (M. Ness et al. 2015; Y.-S. Ting et al. 2017, 2019; M. Xiang et al. 2019; T. Różański et al. 2025; M. Zhang et al. 2024), al-

Corresponding author: Xiaosheng Zhao
xzhao113@jh.edu

though flexible and capable of learning to accommodate survey-specific systematics or noise, often suffer from domain mismatch when directly applied to a new survey. Their use is further limited by the small number of high-quality labels available in the target survey to train new models from scratch. Addressing this cross-survey generalization problem is therefore crucial for leveraging the full potential of current and upcoming spectroscopic datasets.

Resolution, wavelength coverage, noise, and target selection all shift together in real surveys.

Recent advances in machine learning have highlighted the promise of foundation models[1], which have achieved success in natural language processing and computer vision (A. Vaswani et al. 2017; J. Devlin et al. 2018; A. Radford et al. 2018, 2019; T. Brown et al. 2020). In stellar spectroscopy, transformer-based models trained on large datasets can in principle provide general-purpose spectral representations that are transferable across tasks (H. W. Leung & J. Bovy 2024; N. Koblischke & J. Bovy 2024; T. Buck & C. Schwarz 2024; J. Parker et al. 2024; M. Rizhko & J. Bovy 2024; A. J. Smith et al. 2024; F. Zhang et al. 2024; Euclid Collaboration et al. 2025; R. Pattle et al. 2025; X. Zhao et al. 2025b,a). However, the effectiveness of such spectral foundation models for cross-survey stellar parameter estimation remains a less-explored area. In parallel, simpler neural network architectures such as multilayer perceptrons (MLPs; M.-C. Popescu et al. 2009) may already capture much of the transferable information, if trained on sufficiently comprehensive datasets.

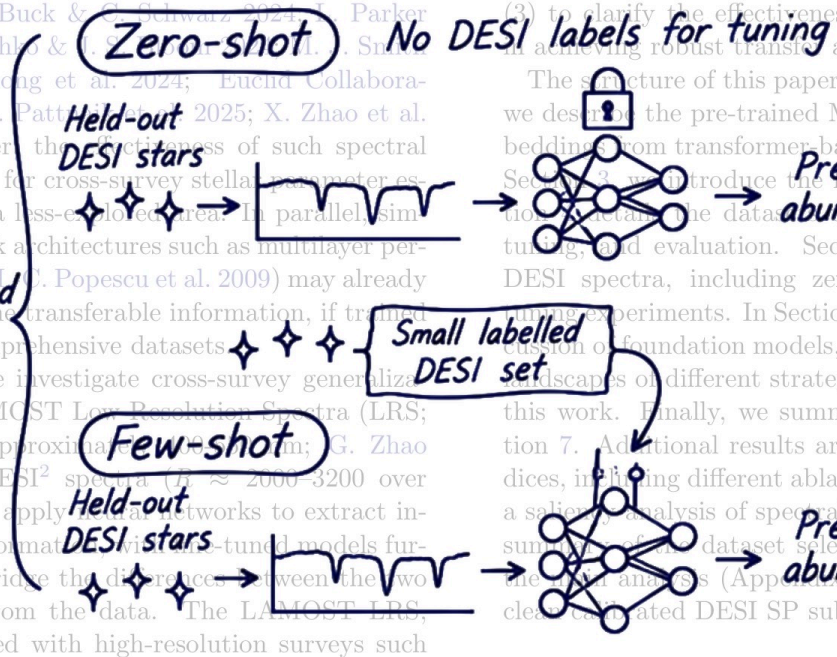
In this work, we investigate cross-survey generalization from the LAMOST Low-Resolution Spectra (LRS; $R \approx 1800$ over approximately 370–900 nm; G. Zhao et al. 2012) to DESI[2] spectra ($R \approx 2000$–3200 over 360–555 nm), and apply neural networks to extract intrinsic physical information. We fine-tuned models further learning to bridge the differences between the two surveys directly from the data. The LAMOST LRS, when cross-matched with high-resolution surveys such as APOGEE to get diverse labels from them, provides a rich training set. We pre-train MLPs on LAMOST LRS and evaluate their ability to generalize to DESI stellar parameter estimation, both in a *zero-shot* setting (directly applying a pre-trained model to a new survey without fine-tuning) and in a *few-shot* setting (using a small number of spectra from the new survey for fine-tuning). For comparison, we also explore transformer-based spectral embeddings as an example of a foundation-model approach, with and without contrastive alignment[3] to external surveys.

Zero-shot tests inherited knowledge; few-shot tests how cheaply we can repair the remaining mismatch.

Another goal of this study is the effectiveness of fine-tuning strategies. We compare full fine-tuning with light-weight approaches such as last-layer and head fine-tuning and Low-Rank Adaptation (LoRA, E. J. Hu et al. 2021; X. Zhao et al. 2025a), and assess how different strategies affect generalization performance.

The goals of this paper are three fold: (1) to evaluate the performance of pre-trained MLPs for cross-survey stellar parameter estimation, compared with the existing DESI pipelines; (2) to evaluate whether MLPs pre-trained directly on spectra, once fine-tuned, achieve performance comparable to or exceeding that of MLPs fine-tuned on foundation-model spectral embeddings; and (3) to clarify the effectiveness of fine-tuning strategies in achieving robust transfer across surveys.

The structure of this paper is as follows. In Section 2, we describe the pre-trained MLPs and the spectral embeddings from transformer-based foundation models. In Section 3, we introduce the fine-tuning strategies. Section 4 presents the datasets used for pre-training, fine-tuning, and evaluation. Section 5 presents results on DESI spectra, including zero-shot transfer and fine-tuning experiments. In Section 6, we present further discussion of foundation models, fine-tuning strategies, loss landscapes of different strategies, and the limitations of this work. Finally, we summarize our findings in Section 7. Additional results are presented in the Appendices, including different ablation studies (Appendix A), a saliency analysis of spectral features (Appendix B), a summary of the dataset selection workflow adopted in the main analysis (Appendix C), and comparison with cleaned and calibrated DESI SP subset (Appendix D).



2. THE MLP MODEL FOR LAMOST LRS

We pre-train simple multilayer perceptron (MLP, M.-C. Popescu et al. 2009) models on the LAMOST low-resolution spectra (LRS) dataset, which contains a moderate number ($\lesssim$ 100,000) of stellar spectra with labels obtained through cross-matching with high-resolution surveys such as APOGEE. One goal of this paper is to test whether such models, despite their simplicity, can provide robust cross-survey generalization.

[1] The "Foundation" model refer to a machine learning model that are pre-trained on a large base of datasets, with enough trainable parameters inside this model that are capable of capturing the full statistical relations in the datasets. Then this model is expected to be finetuned and do a good job for a specific domain or specific task, using a small-to-moderate amount of labeled datasets, assuming the successful pre-training in the previous stage.

[2] https://www.desi.lbl.gov/spectrograph/

[3] Contrastive training encourages similar samples, particularly spectra originating from the same source, to have nearby embeddings and dissimilar samples to be well separated, thereby producing more robust representations.

The experiment compares direct spectral learning with foundation-model embeddings under three ways of adapting to DESI.

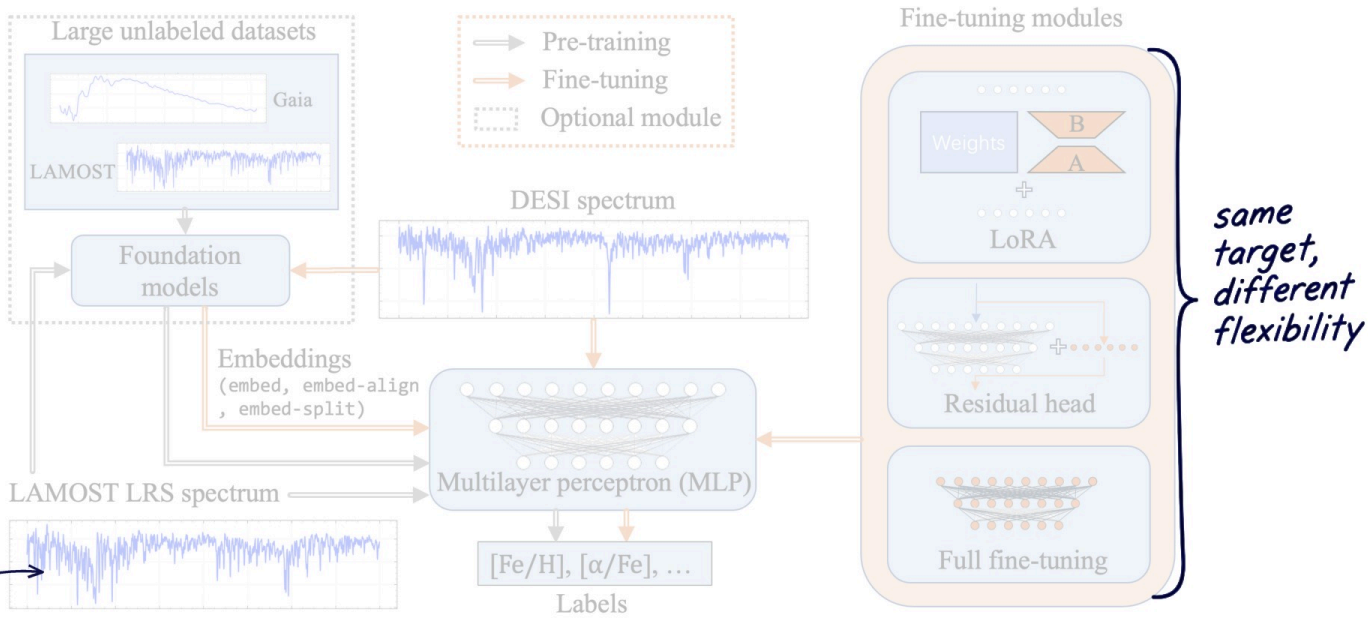


Figure 1. Sketch of the pre-training and fine-tuning workflow. Pre-training uses normalized spectra or foundation-model spectral embeddings to predict labels: [Fe/H] from APOGEE (> -2.0), supplemented at lower metallicities by PASTEL, SAGA, and other VMP/UMP datasets; [α/Fe] from APOGEE. The foundation model is trained on large unlabeled datasets. Fine-tuning then adapts the network using fewer labeled spectra, either from spectra or spectral embeddings, with different fine-tuning modules applied (see Section 3).

direct spectra, simplest route.

The MLP is designed to map spectral information to stellar parameters, with inputs given either as continuum-normalized spectra (hereafter referred to as `MLP-LRS`, or simply `lrs` in figures and tables) or as spectral embeddings (described in Section 2.1), where the spectral fluxes are restricted to the 1,462 logarithmically spaced wavelength bins between 400–560 nm. The network consists of several fully connected layers with ReLU (S. Blanco-Cuaresma 2019) activations, followed by an output layer that predicts the target parameter. In this work we adopt the architecture $[\text{input_dim}, 1024, 512, 64, 1]$, where $\text{input_dim} = 1462$ when the input is continuum-normalized LRS or 768 when the input is their embedding. The resulting model has ~ 2.06 million trainable parameters when applied directly to LRS fluxes.

We adopt the two pre-trained MLPs from X. Zhao et al. (2025b) to predict the logarithmic abundance ratios [Fe/H] and [α/Fe]. The [Fe/H] model uses 90,106 LAMOST LRS with labels drawn from APOGEE DR17 (Abdurro'uf et al. 2022) for stars with [Fe/H] > -2.0, and is supplemented at lower metallicities by the PASTEL and SAGA compilations (Y. Huang et al. 2024; T. Suda et al. 2008; C. Soubiran et al. 2016), the LAMOST/Subaru VMP sample (H. Li et al. 2022), and UMP datasets (F. Sestito et al. 2019). The [α/Fe] model uses 85,400 spectra with APOGEE DR17 labels. In both cases, the data are divided into training and validation sets at a 9:1 ratio. Training is performed with the AdamW (I. Loshchilov & F. Hutter 2019) optimizer (learning rate 1×10^{-3}, weight decay 1×10^{-4}, and batch size of 32) and is trained for 100 epochs, with the best checkpoint selected based on validation performance.

During DESI fine-tuning, only the downstream MLP is adapted; the embedding generators stay fixed.

2.1. *Embeddings from spectral foundation models*

To assess whether spectral foundation models can provide improved transferable features, we also train MLPs on embeddings from the SpecCLIP[4] (X. Zhao et al. 2025b; X. Zhao & Y. Huang 2025) framework. SpecCLIP is designed to align stellar spectra obtained from surveys with different wavelength ranges, resolutions, and noise properties by producing a shared representation space, while also retaining survey-specific information useful for downstream tasks.

The framework consists of modality-specific encoders trained on large spectroscopic datasets. For LAMOST LRS (DR11[5]), 966,082 high-quality spectra with $S/N_g > 50$ and $m_g < 15.8$ were selected. Continuum-normalized spectra are further standardized and tokenized into overlapping segments of length 20 (stride

[4] https://github.com/Xiaosheng-Zhao/SpecCLIP
[5] https://www.lamost.org/dr11/

APOGEE supplies most labels, but rare metal-poor stars require several specialist catalogs.

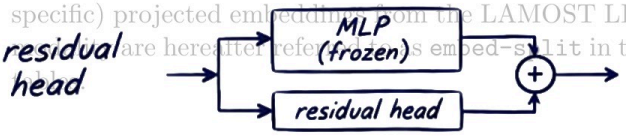
The three adaptation methods trade flexibility against the risk of overwriting useful pre-trained structure.

10), yielding 146 tokens per spectrum. In addition, the logarithmic standard deviation measured before standardization is provided as a separate special token. A 6-layer transformer encoder maps a masked version of these tokens into a sequence of features, from which a 768-dimensional embeddings are produced and used for reconstruction. During training, six token segments are randomly masked and reconstructed, encouraging the encoder to capture global spectral structure. MLPs pre-trained on these embeddings are hereafter referred to as `embed` in Tables 1 and A1.

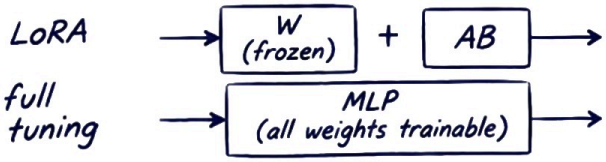
For Gaia XP spectra (F. De Angeli et al. 2023; Gaia Collaboration et al. 2023), which consist of 343 flux points spanning 336–1021 nm, we train an autoencoder with fully connected layers on 1 million normalized spectra (each divided by its flux at 550 nm), of which roughly 80% have matching LAMOST LRS counterparts, irrespective of signal-to-noise ratio. The XP and LRS encoders share the same embedding dimensionality (768) and parameter count (∼42.7 million), allowing their outputs to be compared directly.

For the base SpecCLIP variant used in this work, contrastive training is performed using 820,568 paired LRS–XP spectra to align modalities. Projection heads map the encoder outputs into a shared latent space, with each head including a cross-attention mechanism that projects the modality-specific embeddings into another 768-dimensional vector (one per modality). The model is trained such that paired spectra from the same star are close in this space, while spectra from different stars are pushed apart. The resulting embeddings are designed to support a wide range of downstream tasks, including parameter estimation and similarity search. MLPs pre-trained on the LAMOST LRS embeddings from this alignment are hereafter referred to as `embed-align` in the tables.

Another SpecCLIP variant explicitly partitions the projected 768-dimensional embedding space into a shared 512-dimensional subspace and a modality-specific 256-dimensional subspace through separate projection networks. The shared subspaces of the two modalities are aligned via contrastive training, while for each modality both shared and modality-specific subspaces are used to reconstruct its input, thereby maximizing the mutual information (R. Devon Hjelm et al. 2018) between the input and projected embeddings. MLPs pre-trained on the combined (shared + modality-specific) projected embeddings from the LAMOST LRS [illegible] are hereafter referred to as `embed-split` in the tables.



3. FINE-TUNING STRATEGIES FOR MLPS



In this work, we explore fine-tuning strategies only for the MLPs. For MLPs pre-trained on different spectral embeddings, the components responsible for generating the embeddings are kept fixed. End-to-end fine-tuning of the full models may require substantially larger training sets, and we leave a systematic exploration of that direction to future work. The overall workflow of pre-training and fine-tuning is illustrated in Figure 1.

3.1. *LoRA fine-tuning*

To adapt pre-trained models to DESI spectra with limited labeled data, we explore parameter-efficient fine-tuning using Low-Rank Adaptation (LoRA, E. J. Hu et al. 2021; X. Zhao et al. 2025a). Instead of updating all network parameters, LoRA inserts small low-rank matrices into selected linear layers. This allows the original pre-trained weights to remain frozen while a modest number of additional parameters capture survey-specific adjustments.

Formally, a weight update $\Delta W \in \mathbb{R}^{m\times n}$ is decomposed as $\Delta W = AB$, where $A \in \mathbb{R}^{m\times r}$ and $B \in \mathbb{R}^{r\times n}$ with rank $r \ll \min(m, n)$. This greatly reduces the number of [illegible] while preserving expressive capacity. In our implementation, LoRA modules are applied to all linear layers of the MLP. We experiment with LoRA configurations using a rank of 128 and a scaling factor of $\alpha = 256$, introducing $\sim 5.89 \times 10^5$ additional parameters (about 22% of the MLP size).

Residual heads correct outputs; LoRA alters internal mappings; full tuning can move every MLP weight.

3.2. *Residual-head fine-tuning*

As an alternative, we use a residual-head strategy in which a small MLP, operating in parallel with the frozen pre-trained MLP, is added to correct its outputs. The residual head has the architecture $[\text{input_dim}, 384, 1]$, with ReLU activation. This introduces $\sim 5.62 \times 10^5$ trainable parameters (∼21% of the base MLP size). By restricting adaptation to the residual head, this method balances flexibility with a reduced risk of overfitting.

3.3. *Full fine-tuning*

For completeness, we also apply full fine-tuning, updating all (100%) parameters of the pre-trained MLP. This strategy provides maximal adaptation flexibility but carries an increased risk of overfitting when training data are limited.

3.4. *Fine-tuning details*

All fine-tuning is performed with the AdamW (I. Loshchilov & F. Hutter 2017) optimizer using a learning rate of 1×10^{-5}, weight decay of 1×10^{-4}, and a batch size of 32. Models are trained for up to 200 epochs, and the best checkpoint is chosen based on the

LoRA and residual-head adaptation use similar parameter counts: about 0.6 million each.

A carefully separated DESI–APOGEE cross-match provides 2,069 adaptation stars and 5,539 untouched test stars.

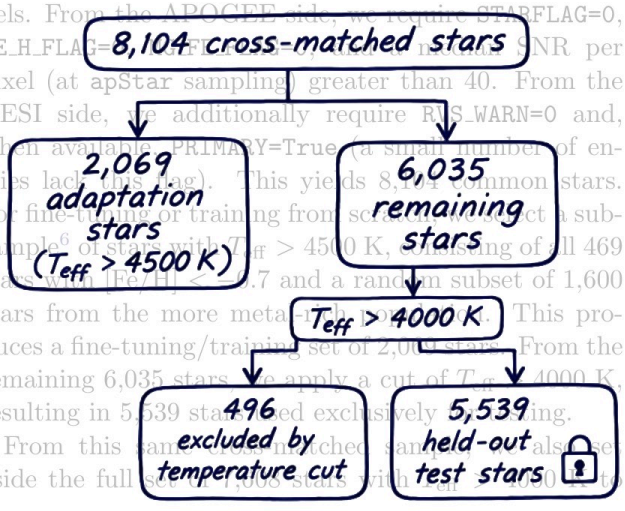
validation performance. Experiments are run on a single NVIDIA V100 GPU; training time depends on the fine-tuning strategy and the size of the fine-tuning dataset; full fine-tuning with the maximum sample size (about 2,000), including testing, takes about 20 s.

4. DATASETS

Identical regridding and normalization remove avoidable preprocessing differences before testing transfer.

We use the Data Release 1 (DR1; S. E. Koposov et al. 2025) of the SPectra Analysis & Retrievable Catalog Lab (SPARCL; S. Juneau et al. 2024; S. Juneau et al. 2021) at NOIRLab's Astro Data Lab (M. J. Fitzpatrick et al. 2014; R. Nikutta et al. 2020). DESI DR1 is used for fine-tuning (or for training from scratch as a baseline) and for testing the MLP models, owing to its substantially larger cross-match with APOGEE compared to DESI EDR. DESI EDR is used exclusively to evaluate the resulting [α/Fe]–[Fe/H] distribution as an initial demonstration, as this dataset has also been examined by independent groups (M. Zhang et al. 2024). All spectra are re-gridded to the same wavelength bins as the LAMOST LRS using linear interpolation and continuum-normalized in the same manner (see Section 2.1). A summary of the sample selection cuts is provided in Appendix C. In the main text, we select samples to retain most of the cross-matched sources, with additional conditions outlined in Appendix D.

For DESI DR1, we first cross-match the catalog with APOGEE DR17 using the Gaia DR3 `source_id` provided in the DESI catalog to obtain high-quality labels. From the APOGEE side, we require `STARFLAG=0`, `FE_H_FLAG=0`, and a combined S/N per pixel (at `apStar` sampling) greater than 40. From the DESI side, we additionally require `RVS_WARN=0` and, when available, `PRIMARY=True` (a small number of entries lack this flag). This yields 8,104 common stars. For fine-tuning or training from scratch, we select a subsample[6] of stars with $T_{\rm eff} > 4500$ K, consisting of all 469 stars with [Fe/H] < -0.7 and a random subset of 1,600 stars from the more metal-rich portion. This produces a fine-tuning/training set of 2,069 stars. From the remaining 6,035 stars, we apply a cut of $T_{\rm eff} > 4000$ K, resulting in 5,539 stars used exclusively for testing.



From this same cross-matched sample, we also set aside the full set of 1,605 stars with [Fe/H] < -2.0 to construct the [α/Fe]–[Fe/H] diagram using APOGEE labels, providing a direct comparison to MLP-based predictions from DESI EDR.

We included all 469 eligible metal-poor stars plus 1,600 metal-richer stars for adaptation.

For DESI EDR, we apply the following quality cuts: `RVS_WARN=0`, `PRIMARY=True`, `RR_SPECTYPE=STAR`, and `MEDIAN_COADD_SNR_B>20`. This yields 134137 stars with unique `TARGETID`. We further restrict the sample to stars with $T_{\rm eff} > 4000$ K (based on DESI SP pipeline estimates), remove sources with missing [Fe/H] or [α/Fe] values, and require successful flux normalization. The final EDR sample contains 126,253 stars, which we use to evaluate the model's consistency with DESI SP predictions. Because only a small fraction of these stars have cross-matched high-quality labels, this sample is used primarily for visual inspection and qualitative comparison with APOGEE DR17 labels cross-matched to DESI DR1.

EDR is used mainly for population structure, not star-by-star accuracy, because few labels overlap.

5. RESULTS

5.1. *Testing the MLP model with DESI spectra*

We first evaluate the zero-shot performance of the MLP models pre-trained on LAMOST LRS, applying them directly to DESI DR1 and EDR spectra. We focus on iron abundance ([Fe/H]) and α-elemental abundance ([α/Fe]), and compare the results with the DESI stellar parameter pipeline and with APOGEE labels for cross-matched stars. As an additional baseline, we also train MLPs from scratch using the same number of DESI spectra as in the fine-tuning experiments.

[α/Fe] *and* [Fe/H] *abundances*—Figure 2 shows that the zero-shot performance of the pre-trained MLP (`MLP-LRS`) provides improved agreement with APOGEE reference labels compared to both the DESI SP pipeline[7] (`DESI SP`) and models trained from scratch (`MLP-Scratch`) on DESI data, underscoring the value of pre-training even for simple architectures. We note that for [α/Fe], only Ca and Mg in the DESI estimates are considered reliable (S. E. Koposov et al. 2025), and the increased scatter around [α/Fe] $\lesssim 0$ is largely due to low-temperature stars with $T_{\rm eff}$ around 4000 K. The figure legend reports the coefficient of determination and robust scatter (standard deviation of the residuals) for each method. For further reference, the comparison us-

[6] This cut avoids low-temperature stars, for which APOGEE ASPCAP (A. E. García Pérez et al. 2016) abundances are reported to be less reliable (see also the SDSS DR17 abundance documentation: https://www.sdss4.org/dr17/irspec/abundances/). In the testing stage, we loosen this constraint to $T_{\rm eff} > 4000$ K to enable broader comparison.

[7] Note that the SP pipeline uses `FERRE` (C. Allende-Prieto & Apogee Team 2023), a template-fitting method, in both the EDR and DR1 releases. The RVS pipeline performs comparably to, or worse than, the SP pipeline, and is therefore not shown here; see S. E. Koposov et al. (2025). Our analysis includes the full metallicity range available in the sample, including the metal-rich regime that is sparsely represented in some previously published comparisons.

Pre-training beats both the DESI pipeline and scratch training, while fine-tuning gives the strongest agreement with APOGEE.

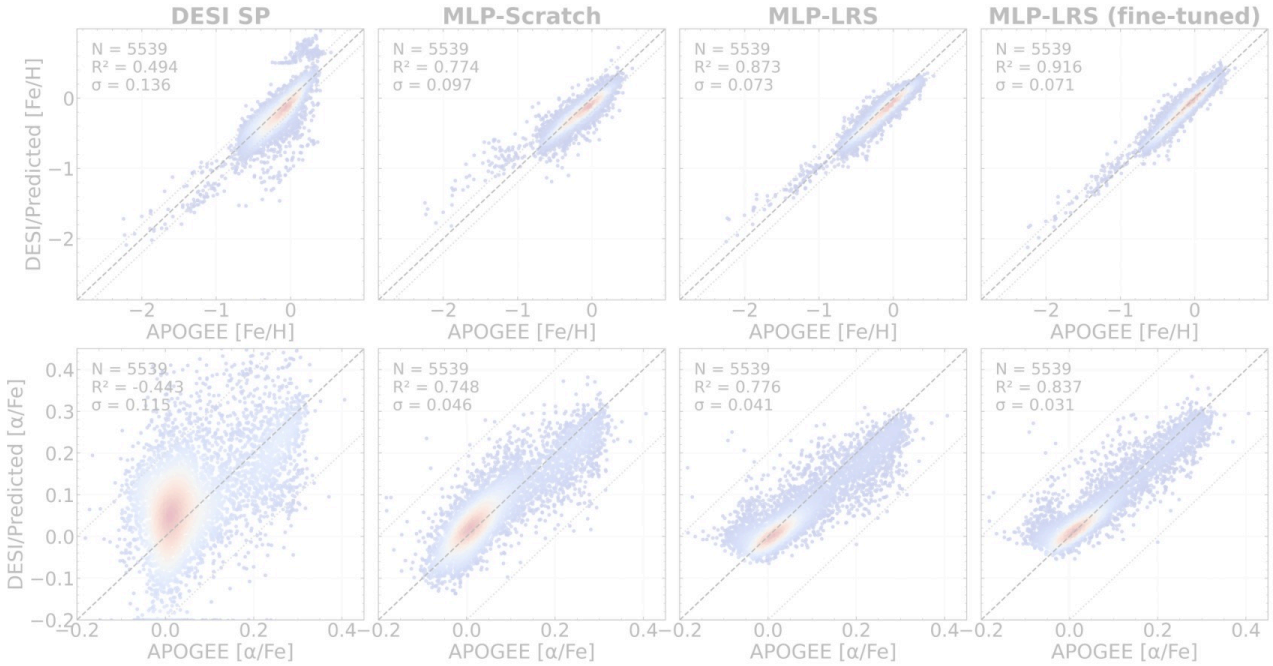


Figure 2. Comparison of [Fe/H] and [α/Fe] estimates between the DESI SP pipeline and MLP-based models, referenced against APOGEE DR17 labels. From left to right: `DESI SP`–DESI SP pipeline, `MLP-Scratch`–MLP trained from scratch on the same number of DESI spectra used for fine-tuning, `MLP-LRS`–MLP pre-trained on LAMOST LRS (zero-shot application), and `MLP-LRS (fine-tuned)`–MLP pre-trained on LAMOST LRS and fine-tuned with DESI spectra. The first row shows [Fe/H] results; the second row shows [α/Fe]. Legends list the number of test stars (N), coefficient of determination (R^2), and the robustly estimated standard deviation of the residuals (σ, computed after 3σ clipping with `sigma_clip` in `astropy`). All sources are restricted to $T_{\rm eff} > 4000$ K. The dashed black line indicates the ideal 1:1 relation, and the gray dotted lines mark ± 0.2 dex deviations. The fine-tuning set consists of 2,069 stars. A comparison using the clean, calibrated DESI SP subset is provided in Appendix D.

The recovered thin-thick disk split is a demanding population-level sanity check, not merely a low residual.

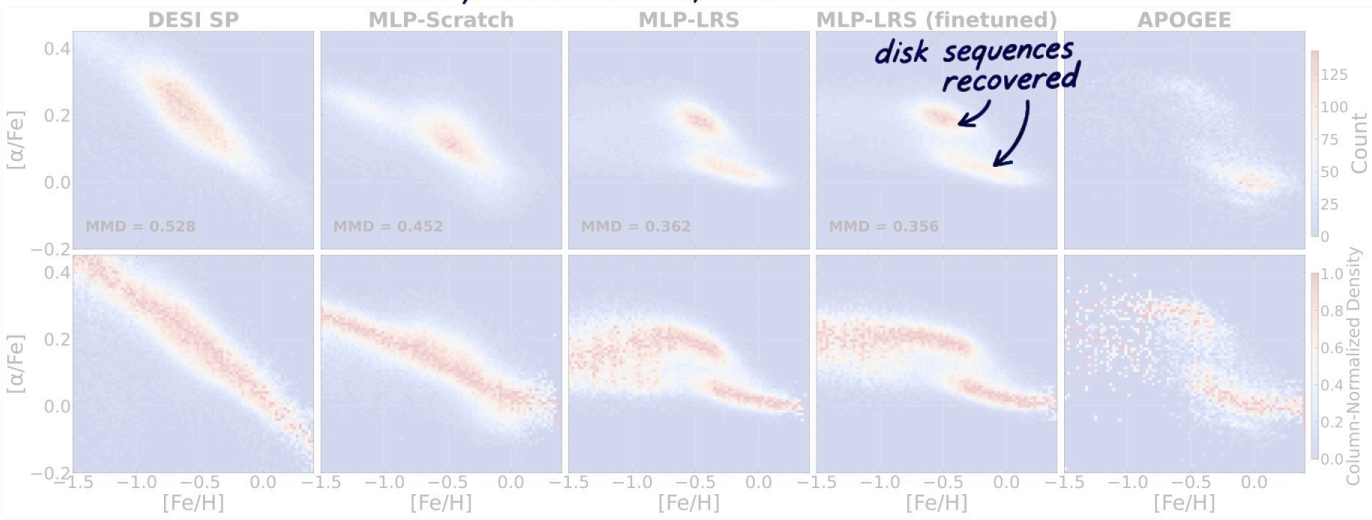


Figure 3. [α/Fe]–[Fe/H] diagrams for DESI EDR and APOGEE reference sample. From left to right: (1) `DESI SP`: DESI SP pipeline, (2) `MLP-Scratch`: MLP trained from scratch on the same number of DESI spectra used for fine-tuning, (3) `MLP-LRS`: MLP pre-trained on LAMOST LRS (zero-shot application), (4) `MLP-LRS (fine-tuned)`: MLP pre-trained on LAMOST LRS and fine-tuned with DESI spectra, and (5) `APOGEE`: DESI DR1–APOGEE DR17 cross-matched sample (7,608 stars). The first row shows raw counts, and the second row shows column-normalized densities, where each [Fe/H] column is divided by its maximum count. The legend reports the Maximum Mean Discrepancy (MMD) between each sample and the APOGEE reference.

Zero-shot uses no DESI training. Fine-tuning lowers MMD from 0.362 to 0.356; DESI SP has 0.528. Lower is closer to APOGEE.

A few hundred DESI labels capture most of the gain, although reducing [Fe/H] scatter needs more data.

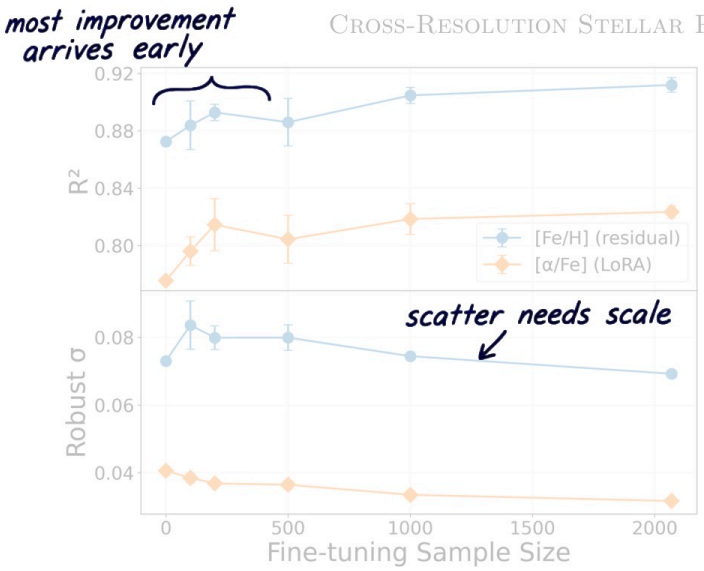


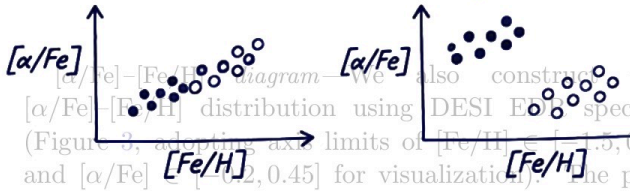
Figure 4. Effect of fine-tuning sample size (0, 100, 200, 500, 1000, and 2069, where 0 denotes zero-shot performance), evaluated on the testing sample. Shown are R^2 and the robustly estimated standard deviation of the residuals (σ) for [Fe/H] and [α/Fe], obtained after fine-tuning (residual-head fine-tuning for [Fe/H] and LoRA fine-tuning for [α/Fe]). For both metrics, we plot the mean and 1σ error bars from five independent runs. The scatter plots for the different sample sizes are shown in Figure A1 and A2.

The initial [Fe/H] scatter bump shows that tiny adaptation sets can briefly worsen residual scatter.

ing the clean, calibrated DESI SP subset is provided in Appendix D.

low MMD

high MMD



[α/Fe]–[Fe/H] diagram.—We also construct the [α/Fe]–[Fe/H] distribution using DESI EDR spectra (Figure 3, adopting axis limits of [Fe/H] $\in [-1.5, 0.4]$ and [α/Fe] $\in [-0.2, 0.45]$ for visualization). The pre-trained MLP reproduces the characteristic separation between the thin and thick disks, as also demonstrated in M. Zhang et al. (2024), whereas the DESI SP pipeline and models trained from scratch show no clear bimodality (see Appendix A.3 for the effect of learning rate on the latter). One reason for our choice of sample selection is to retain as many sources as possible, so that the larger sample better highlights the contrast between the high- and low-[α/Fe] sequences; the same [α/Fe]–[Fe/H] comparisons using a clean, calibrated DESI SP subset are provided in Appendix D.

We quantify the similarity between stellar chemical abundance distributions using the Maximum Mean Discrepancy (MMD; A. Gretton et al. 2012), a non-parametric, kernel-based metric. The unbiased estimator of MMD^2 is computed as

$$\widehat{\mathrm{MMD}}^2 = \frac{1}{m(m-1)} \sum_{i \neq j}^{m} k(x_i, x_j) + \frac{1}{n(n-1)} \sum_{i \neq j}^{n} k(y_i, y_j) - \frac{2}{mn} \sum_{i=1}^{m} \sum_{j=1}^{n} k(x_i, y_j)$$

MMD compares whole abundance distributions, complementing star-by-star R^2 and residual scatter.

where $\{x_i\}_{i=1}^{m}$ and $\{y_i\}_{i=1}^{n}$ are samples drawn from two distributions in the two-dimensional [α/Fe]–[Fe/H] space. One distribution corresponds to one of the first four columns in Figure 3, while the other corresponds to the APOGEE reference sample shown in the last column.

We adopt a Gaussian (RBF) kernel, $k(x, y) = \exp(-\gamma \|x - y\|^2)$, with bandwidth parameter $\gamma = 2.0$. Lower MMD values indicate greater similarity between the two distributions. As shown in Figure 3, the pre-trained MLP models yield smaller MMD values than both the DESI SP pipeline and models trained from scratch, indicating improved consistency with the APOGEE abundance distribution.

Despite these improvements, Figure 2 shows that the pre-trained MLP systematically underestimates [Fe/H] in the metal-rich regime while slightly overestimating it in the metal-poor regime. For [α/Fe], there is an overall tendency toward underestimation—particularly evident in Figure 3, where the underestimation bias of [α/Fe] is most pronounced in the metal-poor regime when compared with the DESI DR1–APOGEE DR17 cross-matched sample shown in the last column. These discrepancies likely reflect differences in resolution and signal-to-noise between LAMOST LRS and DESI spectra, and suggest that modest fine-tuning on DESI data may further improve performance.

5.2. *Fine-tuning the MLP model*

We next apply fine-tuning, using LoRA fine-tuning for [α/Fe] and a residual-head strategy for [Fe/H] (See Section 6.2 for comparisons of fine-tuning strategies for the two parameters). The one-to-one comparisons in Figure 2 (`MLP-LRS (fine-tuned)`) show clear gains relative to the zero-shot models, largely correcting the anomalous trend noted above, which implies that fine-tuning on DESI helps compensate for differences in resolution and spectral properties that are not fully captured by pre-training on LAMOST LRS.

The [α/Fe]–[Fe/H] diagram also improves after fine-tuning (Figure 3, `MLP-LRS (fine-tuned)`). In particular, the underestimation of [α/Fe] in the metal-poor regime is reduced, bringing the distribution into closer agreement with that of the DESI DR1–APOGEE DR17 cross-matched sample. The MMD value decreases fur-

After fine-tuning, embeddings help metal-rich [Fe/H], while direct spectra perform better for metal-poor stars and [alpha/Fe].

Table 1. Coefficient of determination (R^2) for MLPs with different fine-tuning strategies and input types in [Fe/H] and [α/Fe] estimation on DESI spectra.

Negative R^2 means some scratch models are worse than predicting the subset mean.

[Fe/H]				
(Full/MP/MR)	lrs	embed	embed-align	embed-split
residual	0.923 / 0.636 / 0.904	0.922 / 0.196 / 0.913	0.922 / 0.003 / 0.918	0.933 / 0.219 / 0.927
LoRA	0.922 / 0.571 / 0.904	0.921 / 0.031 / 0.916	0.928 / 0.131 / 0.922	0.933 / 0.133 / 0.930
full fine-tune	0.923 / 0.329 / 0.911	0.926 / 0.173 / 0.918	0.927 / 0.085 / 0.922	0.933 / 0.079 / 0.931
zero-shot	0.873 / 0.678 / 0.833	0.833 / 0.694 / 0.779	0.851 / 0.465 / 0.810	0.839 / 0.537 / 0.791
from scratch	0.880 / -0.736 / 0.878	0.877 / -1.246 / 0.886	0.647 / -0.072 / 0.542	0.794 / -0.649 / 0.759
[α/Fe]				
residual	0.799 / 0.366 / 0.802	0.790 / 0.396 / 0.792	0.796 / 0.232 / 0.801	0.796 / 0.335 / 0.800
LoRA	0.829 / 0.418 / 0.831	0.807 / 0.370 / 0.810	0.769 / 0.290 / 0.772	0.733 / 0.205 / 0.736
full-finetune	0.816 / 0.381 / 0.819	0.810 / 0.315 / 0.813	0.784 / 0.247 / 0.787	0.773 / 0.231 / 0.777
zero-shot	0.776 / 0.318 / 0.778	0.730 / 0.183 / 0.733	0.772 / 0.202 / 0.776	0.681 / 0.105 / 0.685
from scratch	0.761 / 0.067 / 0.766	0.729 / 0.401 / 0.730	0.611 / 0.144 / 0.612	0.749 / 0.374 / 0.751

Note. `lrs` refers to an MLP trained directly on normalized LAMOST LRS. `embed` uses embeddings from the LAMOST LRS pre-trained encoder of SpecCLIP, `embed-align` uses projected embeddings of LAMOST LRS from the alignment between LAMOST LRS and Gaia XP spectra, and `embed-split` uses combined (shared + modality-specific) projected embeddings. `Residual`, `LoRA`, and `full fine-tune` denote different fine-tuning strategies applied to pre-trained MLPs. `zero-shot` indicates applying pre-trained MLPs directly to DESI spectra without fine-tuning, and `scratch` indicates MLPs trained from scratch on the same number of DESI spectra as used for fine-tuning. Reported values are R^2 formatted as full/metal-poor/metal-rich (Full/MP/MR), where metal-poor corresponds to [Fe/H] < -1.0, and metal-rich corresponds to [Fe/H] > -1.0. Top three entries of R^2 within each subset ("full", "poor", and "rich") are highlighted with decreasing shades of blue. The numbers are reported as the average over 5 independent runs. Similar to Figure 2, the number of test samples is 5,539, with a fine-tuning set of 2,069 stars.

The metal-poor test subset has only 79 stars. Zero-shot is an exception: embed slightly beats direct LRS there.

the individual abundance [illegible] ment between the model-predicted abundance distribution and the APOGEE reference sample.

Figure 4 shows the effect of fine-tuning sample size on estimation performance after averaging over five independent runs[8]. Notably, even with only 200–500 fine-tuning samples, the performance metrics remain robust and approach those obtained with the maximum sample size. However, for the residual scatter σ of [Fe/H], larger fine-tuning sample sizes (1,000–2,000) are needed to achieve overall performance comparable to or better than the zero-shot case. Further insight into the scatter plots at different sample sizes is provided in Figures A1 and A2.

These results demonstrate that small to moderate amounts of fine-tuning data can enhance cross-survey generalization, even for simple MLP models.

[8] Independence is controlled by five random seeds that are fixed across experiments. These seeds determine the train/validation split and neural network configurations (e.g., weight initialization and dropout). For all sample sizes except 2,069, the seeds also determine which stars are selected for training and validation out of the total 2,069 available, which can lead to larger variation in these experiments.

6. DISCUSSION

6.1. *Foundation models for LAMOST LRS*

We first assess the role of foundation models in stellar parameter estimation. We compare three MLP variants: one trained directly on the normalized LAMOST LRS (`lrs`), one trained on embeddings from the LAMOST LRS pre-trained encoder (`embed`), one trained on projected embeddings from the alignment of LAMOST LRS and Gaia XP spectra via contrastive training (`embed-align`), and one trained on combined (shared + modality-specific) projected embeddings (`embed-split`). The latter three models rely on embeddings derived from SpecCLIP pre-trained in a self-supervised way (see Section 2.1).

Table 1 (see, e.g., the `residual` row for [Fe/H] and `LoRA` row for [α/Fe]) summarizes the results. For [Fe/H], the zero-shot performance of the embedding-based MLPs (`embed`, `embed-align`, and `embed-split`) is not superior to that of MLPs pre-trained on LAMOST LRS itself, except in the metal-poor regime, where `embed` performs better–although these comparisons may be limited by the small testing sample size (79). Fine-tuning the embedding-based MLPs improves performance in the metal-rich regime ([Fe/H] > -1), as ex-

emplified by the best metal-rich results achieved by `embed-split`, but the fine-tuning also degrades performance in the metal-poor regime. For [α/Fe], embedding-based MLPs generally show lower accuracy. These results suggest that, at least for [Fe/H] and [α/Fe], embeddings from SpecCLIP pre-trained on LAMOST LRS and Gaia XP capture some information that downstream MLPs can readily exploit for cross-survey transfer to DESI. However, they may be less informative than the normalized LRS spectra themselves in certain parameter regimes when transferring to DESI, such as the metal-poor regime. A more detailed analysis of sample-size effects on different MLP input types (normalized spectra vs. embeddings) is provided in Appendix A.1.

6.2. *Fine-tuning strategies*

Table 1 also compares different fine-tuning strategies on DESI spectra, including LoRA, residual-head, and full fine-tuning (see, e.g., the `lrs` column). Residual-head fine-tuning yields the best performance among the three strategies for [Fe/H] in the metal-poor regime, close to that of the zero-shot model. It shows slightly worse performance in the metal-rich regime compared to the other two strategies. Nevertheless, its metal-rich performance remains superior to the zero-shot results and corrects the underestimation trend, particularly around [Fe/H] ~ -0, seen in Figure 2 (`MLP-LRS (fine-tuned)` vs. `MLP-LRS`). This suggests that residual heads are effective when spectral features are relatively clean and already well captured by the pre-trained MLP. This includes the metal-poor regime, since the pre-training set contained a relatively well-constructed metal-poor subset (X. Zhao et al. 2025b), so that only small output-level corrections are needed. In contrast, both LoRA and full fine-tuning tend to improve the metal-rich regime but degrade performance in the metal-poor regime, consistent with the LoRA behavior reported by X. Zhao et al. (2025a). This likely reflects the fact that LoRA and full fine-tuning can shift internal representations, which may overfit to the dominant metal-rich ([Fe/H] > -1.0) population at the expense of performance in the minority metal-poor regime. Overall, LoRA and full fine-tuning perform best. Our [Fe/H] estimates are a composite quantity combining multiple spectral abundances, and its reliable estimation is more sensitive to differences in signal-to-noise and resolution between LAMOST LRS and DESI. Moreover, the distribution of [α/Fe] is relatively balanced, which reduces the risk of overfitting to any particular regime. In such cases, more flexible adaptation, such as full fine-tuning, is beneficial. Since LoRA fine-tuning achieves slightly better performance across experiments, we adopt this strategy for [α/Fe] in the main results of this paper. A more detailed analysis of sample-size effects on these strategies is provided in Appendix A.1.

6.3. *Loss landscapes of fine-tuning strategies*

To better understand the behaviors of different fine-tuning strategies, we examine their loss landscapes for [Fe/H] in the metal-poor regime, where fine-tuning is particularly challenging. Following H. Li et al. (2017), we project the high-dimensional trainable parameter space–consisting of only the adapter weights for residual-head and LoRA, or all weights for full fine-tuning–onto a two-dimensional plane defined by two random directions δ and η sampled from a random Gaussian distribution, and evaluate the loss function $f(\alpha, \beta) = L(\theta^* + \alpha\delta + \beta\eta)$, where θ^* is the set of pre-trained parameters and α and β are scale coefficients along the two directions. We use a 21×21 grid with $\alpha, \beta \in [-3, 3]$ (corresponding to a maximum perturbation of three times the filter-wise L2 norm of weights), and normalize by filter-wise scaling, where a "filter" corresponds to the weights associated with one output neuron[9] in the MLP.

We adopt mean squared error (MSE) loss between predicted and true [Fe/H], computed on 79 DESI spectra (from the testing sample) with [Fe/H] < -1, labels taken from APOGEE DR17, cross-matched with the LRS experiment. The resulting loss landscapes along two example random directions are shown in Figure 5. Compared with residual-head fine-tuning, LoRA yields a higher minimum in the landscape over the perturbations, and this minimum lies farther from the pre-trained solution, suggesting overfitting to the metal-rich regime, where the globally optimized weights are suboptimal for metal-poor stars. Full fine-tuning yields its minimum away from the pre-trained weights as well, but with an overall shallower loss landscape around the pre-trained solution. These results are consistent across random seeds (ten random directions), with only one exception in which LoRA attains the smallest minimum–though this minimum occurs much farther from the pre-trained solution than in the residual-head case. These loss landscape visualizations help explain the empirical differences in the loss strategies. A more physically interpretable saliency analysis of spectral features, together with a comparison of different fine-tuning strategies in this regime, is provided in Appendix B. In particular, the residual-head

[9] For LoRA layers, the "filter" in our implementation corresponds to one rank vector in the LoRA-A matrix, or one output-dimension vector in the LoRA-B matrix.

The best adaptation strategy depends on whether the target needs a small correction or a deeper representation change.

Possible explanation, not proof

Schematic

Metal-rich majority

Metal-poor minority

Metal-rich majority

Metal-poor minority

Frozen backbone

Internal mapping can shift

Small output correction

For [Fe/H], LoRA and full tuning may overfit the dominant metal-rich population at the expense of rarer metal-poor stars.

This is why we use residual tuning for [Fe/H] and LoRA for [α/Fe] in the main results.

One plausible explanation: [alpha/Fe] needs more flexible adaptation to changes in resolution and signal-to-noise.

Metal-poor loss landscapes help explain the stronger transfer from residual-head tuning; they are diagnostic slices, not proof.

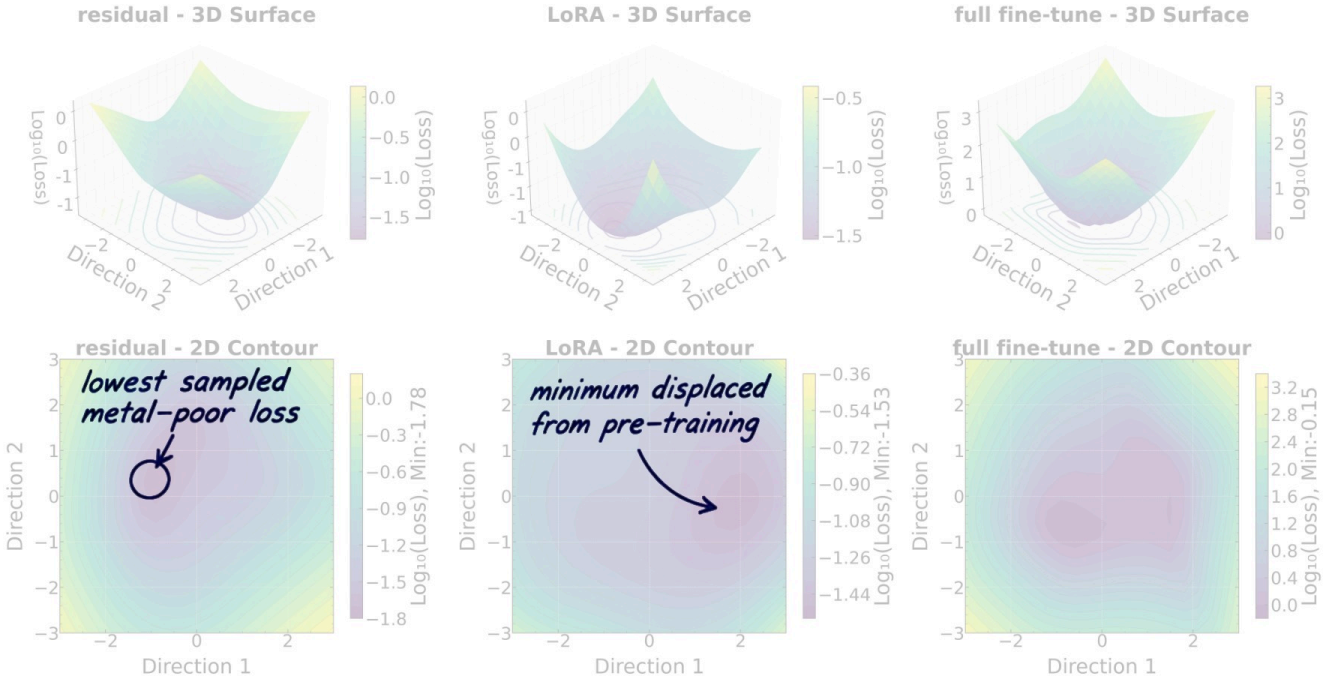


Figure 5. Loss landscapes for different fine-tuning strategies, evaluated on the metal-poor regime of [Fe/H] (79 stars). The first row shows 3D surfaces, and the second row shows 2D contours of the logarithmic MSE loss, plotted as a function of perturbations around the pre-trained model parameters along two random directions. The minimum losses are indicated in the colorbar of the 2D contours.

broader local basin

fine-tuning shows higher sensitivity of [Fe/H] to the relevant lines in the spectra. A detailed mathematical analysis of these discrepancies among fine-tuning strategies is left for future work.

The relevant test contains only 79 metal-poor stars, so we treat landscape geometry as diagnostic, not proof.

6.4. *Limitations*

Several limitations of this work should be noted. First, the spectral foundation models used here (SpecCLIP) are pre-trained on blue-arm spectra (400–560 nm) of LAMOST low-resolution spectra (LRS). While it contains many informative lines, it does not capture the full information content of stellar spectra. Second, the encoders we employ are relatively modest in scale (∼42.7 million parameters for the LAMOST LRS model and ∼0.1 billion when including alignment with the Gaia XP model), whereas larger models may be required to represent more complex spectral features. Third, we demonstrate our analysis on only two parameters, [Fe/H] and [α/Fe]; other quantities such as $T_{\rm eff}$ and $\log g$ may require different modeling choices. Fourth, the fine-tuning datasets are limited in size; larger samples may be necessary to fully evaluate the potential of foundation models, especially when fine-tuning backbone transformer layers rather than only downstream MLPs. In such cases, approaches like LoRA–shown to be effective in large language models and requiring only a limited number of trainable parameters–may prove more advantageous for spectral applications. Finally, while our results highlight the surprising strength of simple MLPs in cross-survey generalization, a foundation model trained directly on DESI spectra—with sufficient high-quality labels—may ultimately surpass them. Indeed, X. Zhao et al. (2025b) has shown that such models can achieve state-of-the-art performance, suggesting that future survey-specific foundation models remain a promising direction.

The colour scales differ, so compare numerical loss values, not darkness across panels.

Our foundation model is modest, blue-only, and tested on just two labels; broader claims would be premature.

7. SUMMARY

We summarize the main findings of this work as follows. This paper investigates the effectiveness of pre-trained multilayer perceptron (MLP) models for low-to-moderate-resolution cross-survey generalization, using the transfer from LAMOST low-resolution spectra (LRS) to DESI medium-resolution spectra (MRS) as a case study.

First, we show that pre-trained MLPs produce estimates that are more consistent with APOGEE labels than the DESI stellar parameter pipeline for both iron abundance ([Fe/H]) and α-elemental abundance ([α/Fe]), even in a zero-shot setting without fine-tuning. Notably, the model recovers the two thin-thick disk sequences in the DESI DR1 sample–features that do not appear in S. E. Koposov et al. (2025). Modest fine-tuning (∼2,000 spectra) with DESI further improves performance, suggesting that fine-tuning helps mitigate

For LAMOST-to-DESI transfer, a pre-trained MLP plus modest fine-tuning provides a practical baseline.

differences in resolution and data quality between surveys.

A DESI-trained foundation model may ultimately win; this paper establishes the baseline it must beat.

Second, we find that spectral foundation models, while embeddings from these models can be useful for some applications, do not outperform supervised spectral foundation models trained on LRS data. Despite their larger model sizes and pre-training on millions of spectra, they underperform in the metal-rich regime and do not improve [Fe/H] performance in the metal-rich regime, they degrade performance in the metal-poor regime and do not improve [α/Fe] predictions. This indicates that, for cross-survey transfer between LAMOST and DESI, simple MLPs trained directly on LRS remain competitive with more complex foundation models.

Third, we find that fine-tuning strategies have parameter-dependent effects: residual-head fine-tuning is most effective for [Fe/H], while LoRA yields better results for [α/Fe] given our fine-tuning sample size. These findings highlight the importance of tailoring fine-tuning approaches to the specific stellar parameter of interest.



Our experimental results demonstrate that pre-trained MLPs offer a strong baseline for cross-survey generalization in stellar spectroscopic analyses. Even with modest labelled DESI sets, we show even lightweight architectures can capture robust information transfer across surveys, while the role of larger spectral foundation models remains an open direction for exploration–especially as larger and more diverse fine-tuning datasets become available.

The demonstrated transfer covers [Fe/H] and [alpha/Fe]; other labels and surveys still need validation.

From a survey design perspective, these findings suggest a practical path for upcoming large-scale stellar spectroscopic programs such as the stellar program with DESI, SDSS-V, and the Galactic Archaeology survey of the Subaru Strategic Program with the Prime Focus Spectrograph on the Subaru telescope (M. Chiba et al. 2026). Simple pre-trained MLPs, when initialized on existing low-resolution surveys and fine-tuned with modest samples of new survey data, can already deliver competitive parameter estimates and recover key Galactic structures such as the thin–thick disk separation. This demonstrates the immediate utility of lightweight models for rapid deployment.

ACKNOWLEDGEMENTS

X.Z., R.F.G.W., A.S.S., L.D., T.B., and V.W. have been supported by a grant from the Schmidt Sciences. Y.S.T is supported by the National Science Foundation under Grant No. AST-2406729. Y.H. acknowledges the support from the National Science Foundation of China (NSFC grant No. 12422303), the Fundamental Research Funds for the Central Universities (grant Nos. 118900M122, E5EQ3301X2, and E4EQ3301X2), and the National Key R&D Programme of China (grant No. 2023YFA1608303).

This work made use of the data from LAMOST (Large Sky Area Multi-Object Fiber Spectroscopic Telescope, also known as the Guoshoujing Telescope) (https://cstr.cn/31118.02.LAMOST). LAMOST is a Chinese national mega-science facility, operated by National Astronomical Observatories, Chinese Academy of Sciences.

This research used data obtained with the Dark Energy Spectroscopic Instrument (DESI). DESI construction and operations is managed by the Lawrence Berkeley National Laboratory. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High-Energy Physics, under Contract No. DE–AC02–05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions: www.desi.lbl.gov/collaborating-institutions. The DESI collaboration is honored to be permitted to conduct scientific research on I'oligam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. National Science Foundation, the U.S. Department of Energy, or any of the listed funding agencies.

Funding for the Sloan Digital Sky Survey V has been provided by the Alfred P. Sloan Foundation, the Heising-Simons Foundation, the National Science Foundation, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. SDSS telescopes are located at Apache Point Observatory, funded by the Astrophysical Research Consortium and operated by New Mexico State University, and at Las Campanas Observatory, operated by the Carnegie Institution for Science. The SDSS web site is www.sdss.org.

SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration, including the Carnegie Institution for

The appendix stress-tests how sample size, representation choice, and optimization affect the headline result.

Science, Chilean National Time Allocation Committee (CNTAC) ratified researchers, Caltech, the Gotham Participation Group, Harvard University, Heidelberg University, The Flatiron Institute, The Johns Hopkins University, L'Ecole polytechnique fédérale de Lausanne (EPFL), Leibniz-Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Extraterrestrische Physik (MPE), Nanjing University, National Astronomical Observatories of China (NAOC), New Mexico State University, The Ohio State University, Pennsylvania State University, Smithsonian Astrophysical Observatory, Space Telescope Science Institute (STScI), the Stellar Astrophysics Participation Group, Universidad Nacional Autónoma de México (UNAM), University of Arizona, University of Colorado Boulder, University of Illinois at Urbana-Champaign, University of Toronto, University of Utah, University of Virginia, Yale University, and Yunnan University.

DATA AVAILABILITY

All observational data used in this work are publicly available from the archives. No proprietary data were used. The code and trained weights are available upon request.

Software: SpecCLIP (X. Zhao & Y. Huang 2025), PyTorch (A. Paszke et al. 2019), Astropy (Astropy Collaboration et al. 2013, 2018), SPARCL (S. Juneau et al. 2024)

Figures A1/A2 evaluate all unused stars; the sample size tables use the fixed 5,539-star test sample.

Metal-rich [Fe/H] bias begins correcting with about 100 examples; [α/Fe] needs roughly 200.

APPENDIX

A. ABLATION STUDIES

In this section, we systematically examine the impact of key fine-tuning factors, including the number of training samples, the number of trainable parameters, and the learning rate for training from scratch. The results are presented in Appendices A.1–A.3.

A.1. *Scatter plots and performance of embedding types and fine-tuning strategies across sample sizes*

Here we present additional results on how fine-tuning sample size impacts one-to-one comparisons with APOGEE labels, and on the performance of different embedding types and fine-tuning strategies.

Figure A1 and A2 show the one-to-one comparisons between the predictions and the APOGEE labels and the [α/Fe]–[Fe/H] diagrams across different fine-tuning sample sizes. The evaluation is performed on all remaining samples not used for fine-tuning (in contrast to all other tables and figures in this paper, which use only the 5,539-star testing sample), in order to maximize the testing coverage in each case. For [Fe/H], the underestimation trend in the metal-rich regime is already corrected with as few as 100 fine-tuning samples, although the overall scatter is initially larger and decreases as the fine-tuning sample size increases. For [α/Fe], the underestimation trend begins to be corrected with about 200 fine-tuning samples and continues to improve with larger sample sizes. Note that R^2 fluctuates because it is affected by both the testing sample size and random-seed variance. Comparisons using the fixed 5,539-star testing sample, with results averaged over the same five seeds, are shown in Table A1.

Table A1 compares MLPs trained on normalized spectra (`lrs`), SpecCLIP embeddings (`embed`, `embed-align`, and `embed-split`), and models trained from scratch across a range of fine-tuning sample sizes, for both [Fe/H] and [α/Fe]. The ratio of metal-poor to metal-rich stars is kept fixed in each subset. Using the `lrs` model with as few as 100 (12:88) fine-tuning samples, the pre-trained MLP achieves better performance in the majority of cases (metal-rich stars for [Fe/H] and α-rich stars for [α/Fe]) compared to the zero-shot case (last column). This demonstrates that the pre-trained MLP has learned transferable features from the LAMOST LRS dataset that can be effectively adapted to DESI with only small amounts of fine-tuning data. Embedding-based MLPs perform better in the metal-rich regime across sample sizes, but underperform in the metal-poor regime for [Fe/H] and consistently underperform for [α/Fe]. Models trained from scratch perform poorly at all sample sizes, underscoring the importance of pre-training.

Table A2 evaluates fine-tuning strategies under the same sample-size variations. For [Fe/H], residual-head fine-tuning consistently yields the best results among the three strategies, particularly in the metal-poor regime, and remains robust with small sample sizes. It also shows better performance in the metal-rich regime when the fine-tuning sample size is limited. LoRA and full fine-tuning perform competitively in the metal-rich regime with moderate fine-tuning sample sizes, but perform less well with smaller sample sizes and in the metal-poor regime. For [α/Fe], LoRA fine-tuning performs best overall, with full fine-tuning remaining competitive and achieving comparable performance, while

More adaptation data can reduce abundance scatter while retaining the thin-thick disk structure.

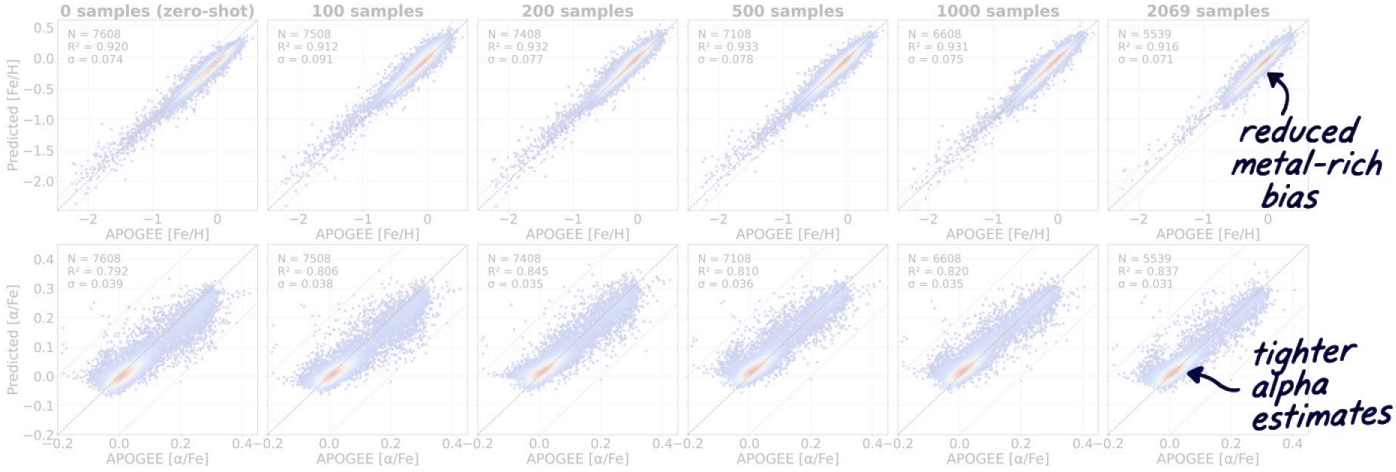


Figure A1. Effect of fine-tuning sample size on the one-to-one comparisons, evaluated on all remaining samples not used for fine-tuning (See Table A1 for comparisons using a fixed testing sample). Results are shown for residual-head fine-tuning of [Fe/H] and LoRA fine-tuning of [α/Fe]. Other conventions follow Figure 2.

These panels use changing test samples and random splits; their R-squared values are not a controlled fixed-test comparison.

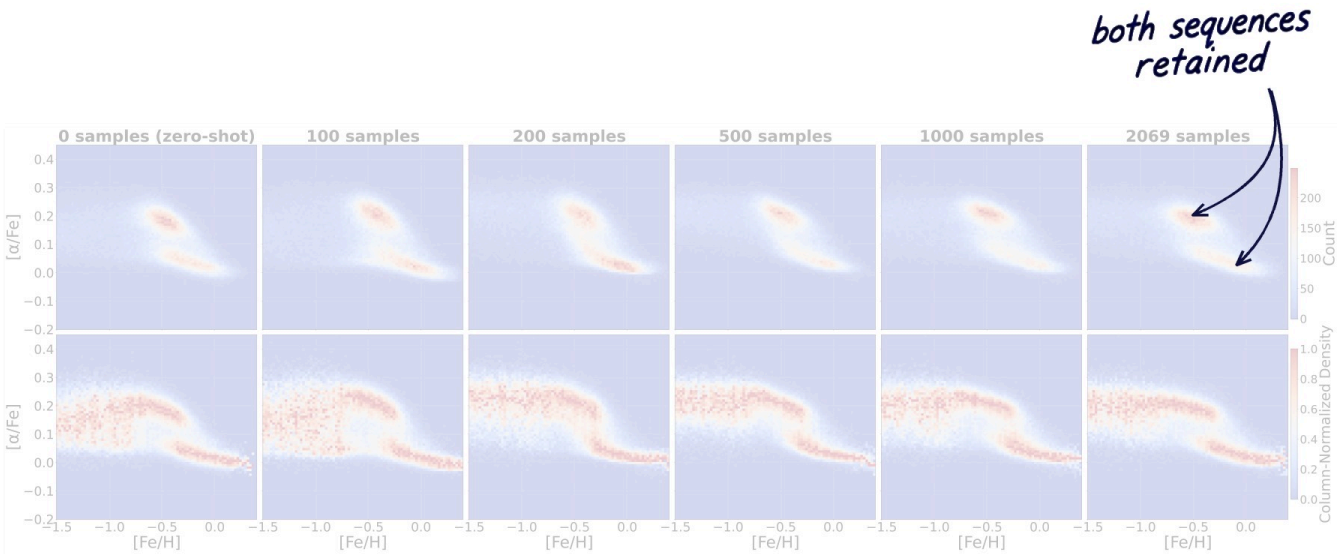


Figure A2. Effect of fine-tuning sample size on the predicted [α/Fe]–[Fe/H] diagrams, evaluated on all remaining samples not used for fine-tuning. Results are shown for residual-head fine-tuning of [Fe/H] and LoRA fine-tuning of [α/Fe]. Other conventions follow Figure 3.

The branch geometry stabilizes before every individual prediction becomes precise.

trained from scratch perform poorly, again highlighting the benefits of pre-training.

With fine-tuning, direct LRS inputs outperform embeddings for metal-poor [Fe/H] and for [alpha/Fe].

Table A1. Coefficient of determination (R^2) for MLPs with different input types as a function of fine-tuning sample size.

[Fe/H]			
M-poor:M-rich	243:1826	117:883	59:441
lrs	0.923 / 0.636 / 0.904	0.916 / 0.619 / 0.894	0.910 / 0.600 / 0.886
embed	0.922 / 0.196 / 0.913	0.920 / 0.160 / 0.911	0.913 / 0.216 / 0.900
embed-align	0.922 / 0.003 / 0.918	0.922 / -0.026 / 0.918	0.910 / -0.233 / 0.907
embed-split	0.933 / 0.219 / 0.927	0.927 / 0.058 / 0.924	0.926 / 0.213 / 0.917
from scratch	0.880 / -0.736 / 0.878	0.826 / -1.663 / 0.826	0.433 / -3.575 / 0.336
M-poor:M-rich	23:177	12:88	0:0
lrs	0.909 / 0.616 / 0.885	0.906 / 0.643 / 0.881	0.873 / 0.678 / 0.833
embed	0.916 / 0.232 / 0.904	0.915 / 0.245 / 0.902	0.833 / 0.694 / 0.779
embed-align	0.914 / -0.247 / 0.913	0.914 / -0.150 / 0.911	0.851 / 0.465 / 0.810
embed-split	0.925 / -0.098 / 0.924	0.916 / -0.130 / 0.913	0.839 / 0.537 / 0.791
from scratch	0.098 / -4.912 / -0.090	-0.575 / -8.525 / -0.924	
[α/Fe]			
M-poor:M-rich	243:1826	117:883	59:441
lrs	0.829 / 0.418 / 0.831	0.828 / 0.363 / 0.831	0.820 / 0.287 / 0.824
embed	0.807 / 0.370 / 0.810	0.797 / 0.184 / 0.802	0.786 / 0.266 / 0.790
embed-align	0.769 / 0.290 / 0.772	0.762 / 0.212 / 0.765	0.779 / 0.194 / 0.784
embed-split	0.733 / 0.205 / 0.736	0.758 / 0.084 / 0.764	0.714 / 0.028 / 0.719
from scratch	0.761 / 0.067 / 0.766	0.647 / -0.042 / 0.651	-0.333 / -0.880 / -0.342
M-poor:M-rich	23:177	12:88	0:0
lrs	0.827 / 0.329 / 0.831	0.808 / 0.297 / 0.812	0.776 / 0.318 / 0.778
embed	0.793 / 0.298 / 0.796	0.785 / 0.265 / 0.789	0.730 / 0.183 / 0.733
embed-align	0.777 / 0.221 / 0.781	0.774 / 0.244 / 0.778	0.772 / 0.202 / 0.776
embed-split	0.725 / 0.089 / 0.730	0.725 / 0.121 / 0.729	0.681 / 0.105 / 0.685
from scratch	-0.381 / -0.867 / -0.392	-1.631 / -2.117 / -1.657	

Note. Conventions follow Table 1, but with varying numbers of fine-tuning samples (keeping the same ratio of metal-poor (M-poor) to metal-rich (M-rich) stars, as listed in the first row for [Fe/H] and [α/Fe]). The last column shows the zero-shot results. Top three entries of R^2 within each subset ("full", "poor", and "rich") are highlighted with decreasing shades of blue.

strong metal-poor transfer

scratch failure

useful with 100 labels

Zero-shot metal-poor [Fe/H] is an exception: embed slightly exceeds direct LRS. Fine-tuning changes that comparison.

Scratch training collapses first when the adaptation set becomes small.

Residual tuning leads metal-poor iron estimates after adaptation; the best alpha strategy depends on sample size.

Table A2. Coefficient of determination (R^2) for MLPs with different fine-tuning strategies as a function of fine-tuning sample size.

[Fe/H]			
M-poor:M-rich	243:1826	117:883	59:441
residual	0.923 / 0.636 / 0.904	0.916 / 0.619 / 0.894	0.910 / 0.600 / 0.886
LoRA	0.922 / 0.571 / 0.904	0.912 / 0.496 / 0.891	0.914 / 0.590 / 0.892
full fine-tune	0.923 / 0.329 / 0.911	0.916 / 0.401 / 0.900	0.915 / 0.516 / 0.895
scratch	0.880 / -0.736 / 0.878	0.826 / -1.663 / 0.826	0.433 / -3.575 / 0.336
M-poor:M-rich	23:177	12:88	0:0
residual	0.909 / 0.616 / 0.885	0.906 / 0.643 / 0.881	0.873 / 0.678 / 0.833
LoRA	0.897 / 0.562 / 0.869	0.874 / 0.499 / 0.840	0.873 / 0.678 / 0.833
full fine-tune	0.904 / 0.494 / 0.881	0.887 / 0.501 / 0.858	0.873 / 0.678 / 0.833
from scratch	0.098 / -4.912 / -0.090	-0.575 / -8.525 / -0.924	
[α/Fe]			
M-poor:M-rich	243:1826	117:883	59:441
residual	0.799 / 0.366 / 0.802	0.824 / 0.285 / 0.828	0.840 / 0.347 / 0.844
LoRA	0.829 / 0.418 / 0.831	0.828 / 0.363 / 0.831	0.820 / 0.287 / 0.824
full-finetune	0.816 / 0.381 / 0.819	0.816 / 0.356 / 0.819	0.812 / 0.264 / 0.816
from scratch	0.761 / 0.067 / 0.766	0.647 / -0.042 / 0.651	-0.333 / -0.880 / -0.342
M-poor:M-rich	23:177	12:88	0:0
residual	0.813 / 0.292 / 0.817	0.830 / 0.206 / 0.835	0.776 / 0.318 / 0.778
LoRA	0.827 / 0.329 / 0.831	0.808 / 0.297 / 0.812	0.776 / 0.318 / 0.778
full-finetune	0.826 / 0.329 / 0.830	0.807 / 0.290 / 0.810	0.776 / 0.318 / 0.778
scratch	-0.381 / -0.867 / -0.392	-1.631 / -2.117 / -1.657	

best with 2,069 examples

residual leads with 500 examples

Note. Conventions follow Table 1, but with varying numbers of fine-tuning samples (keeping the same ratio of metal-poor (M-poor) to metal-rich (M-rich) stars, as listed in the first row for [Fe/H] and [α/Fe]). The last column shows the zero-shot results. Top three entries of R^2 within each subset ("full", "poor", and "rich") are highlighted with decreasing shades of blue.

A.2. *Number of trainable parameters for different fine-tuning strategies*

H384 residual-head and R128 LoRA each use about 0.6 million trainable parameters in the main comparison.

We explore the impact of the number of trainable parameters on fine-tuning performance, using the residual-head and LoRA. Figure A3 shows the results for [Fe/H] estimation using the full set of 2,069 fine-tuning samples, averaged over five independent runs. LoRA achieves better performance in the metal-rich regime, whereas residual-head fine-tuning performs better in the metal-poor regime, across a range from a few thousand up to ∼1M trainable parameters. For the results reported in the main text, we adopt a hidden layer size of 384 (`H384`) for residual-head fine-tuning and a rank of 128 (`R128`) for LoRA, corresponding to ∼0.6 million trainable parameters for both strategies, ensuring a fair comparison.

A.3. *Learning rate for training from scratch*

We also test the effect of learning rate on models trained from scratch with DESI spectra. Four different learning rates are explored (1e-2, 1e-3, 1e-4, and 1e-5), using the same number of DESI spectra as in the fine-tuning experiments. As shown in Figure A4, a relatively large learning rate (1e-3) yields better performance for both [Fe/H] and [α/Fe] compared to the smaller rate (1e-5) used for LAMOST LRS pre-training and DESI fine-tuning. The [α/Fe]–[Fe/H] diagram in Figure A5 also indicates that a learning rate of 1e-3 better recovers the thin–thick disk separation, whereas 1e-5 fails to capture these trends. However, even this selected learning rate introduces systematic effects: [Fe/H] is overestimated in the metal-poor regime (Figure A4), and [α/Fe] is likely overestimated for metal-poor stars (Figure A5).

B. SALIENCY ANALYSIS OF SPECTRAL FEATURES

Here we examine the sensitivity of the output iron abundance–in particular for the metal-poor subset–to the input DESI spectra. The saliency maps (K. Simonyan et al. 2013) used here are defined as the gradients $\partial\widehat{[\mathrm{Fe/H}]}/\partial F_\lambda$, i.e.

The zero-shot metal-poor iron score stays higher than these adapted scores: overall gains need not benefit every subset.

Residual often wins on metal-poor iron, while scratch-model scores depend on learning rate.

Residual is stronger at most tested capacities, but the metal-poor curves cross at several choices.

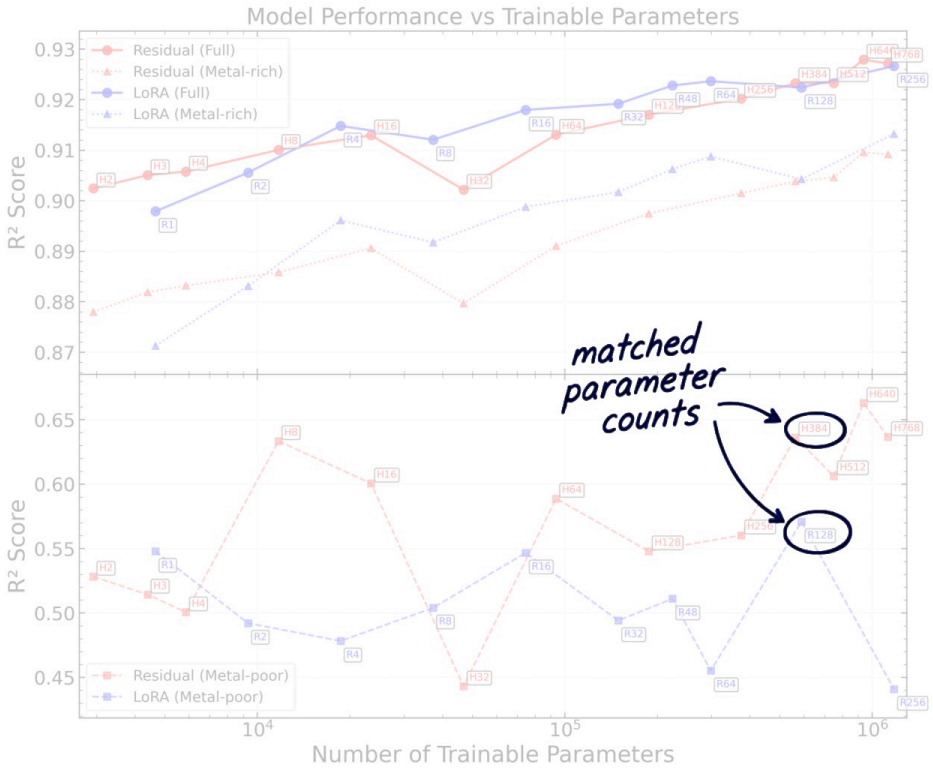


Figure A3. Effect of the number of trainable parameters on [Fe/H] fine-tuning performance, comparing LoRA and residual-head strategies. The y-axis shows the coefficient of determination (R^2) for [Fe/H] estimation using the full sample, metal-rich subset ([Fe/H] > -1.0), and metal-poor subset ([Fe/H] < -1.0). The number of test samples is 5,539, with a fine-tuning set of 2,069 stars. Annotations indicate the hidden-layer size for residual-head models (e.g., H384) and the rank for LoRA models (e.g., R128).

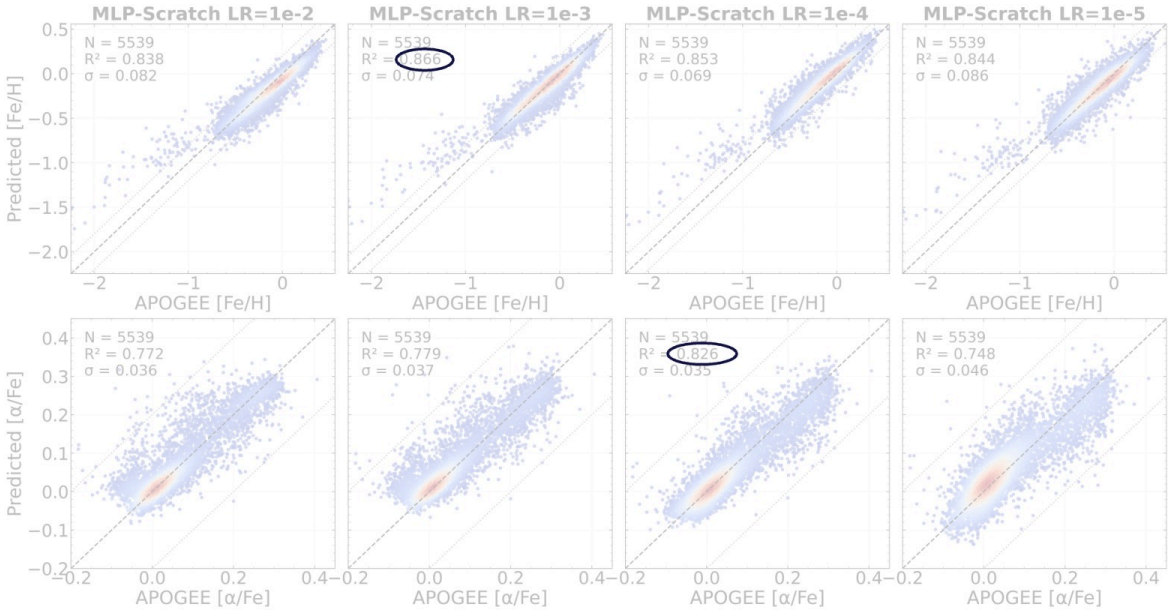


Figure A4. Effect of learning rate (LR) on models trained from scratch for [Fe/H] and [α/Fe], referenced against APOGEE. Conventions follow Figure 2.

The highest displayed scratch R-squared uses LR=1e-3 for iron and LR=1e-4 for alpha; no single rate wins both.

Scratch population structure depends on learning rate; saliency links transferred predictions to familiar spectral features.

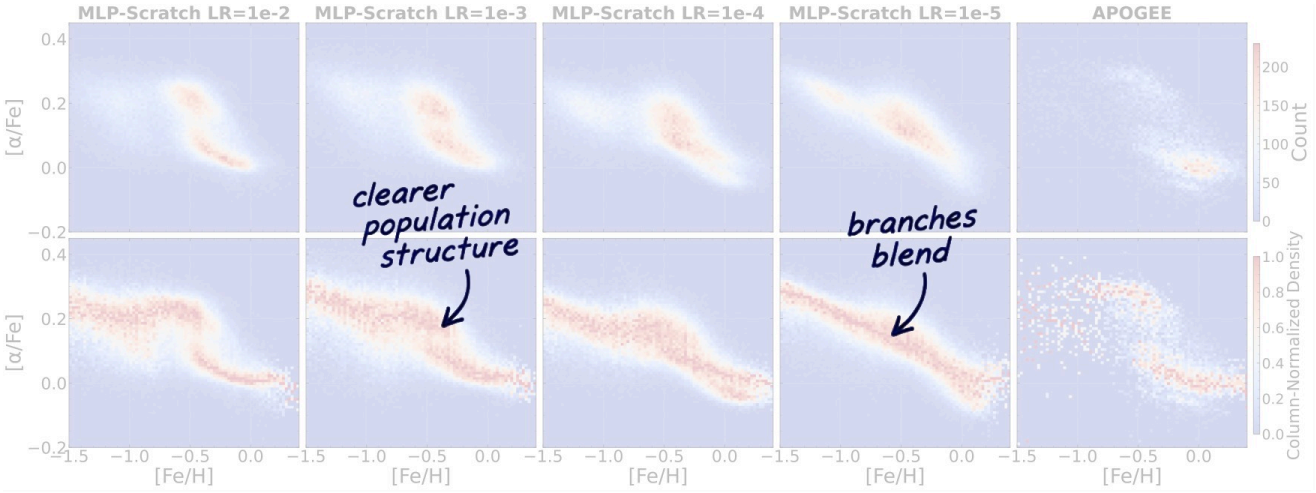


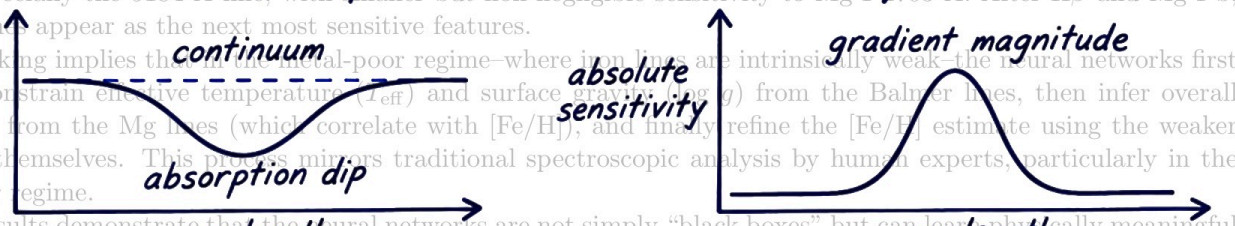
Figure A5. Effect of learning rate (LR) on models trained from scratch for the [α/Fe]–[Fe/H] diagram. Conventions follow Figure 3.

the derivative of the predicted abundance $\widehat{[\mathrm{Fe/H}]}$ with respect to the input spectrum F_λ. We take the absolute value of the gradients as a first step in visualization. The results are shown in Figure B6, where all spectra and saliency maps have been shifted to the rest frame using radial velocities reported by the `DESI SP` pipeline.

Comparing the upper and middle panels, we find that the neural networks after residual-head fine-tuning are sensitive to both relatively clean and relatively blended iron lines. The lower panel shows that zero-shot testing with the MLP trained on LAMOST LRS already captures the correct sensitivity to iron lines in DESI spectra, even without ever having seen DESI data during training. Fine-tuning further enhances the sensitivity to most iron lines, which may partly explain its superior performance in [Fe/H] estimation (Sections 6.2 and 6.3).

Balmer and Mg sensitivity suggests stellar-parameter context matters when iron lines are weak.

An additional finding from the saliency maps is that the networks show the strongest sensitivity to the Hβ Balmer line, with somewhat weaker but still notable sensitivity to Hδ and Hγ. The next strongest signal is to the Mg I b triplet, especially the 5184 Å line, with smaller but non-negligible sensitivity to Mg I 8703 Å. After Hβ and Mg I b, the iron lines appear as the next most sensitive features.



This ranking implies that in the metal-poor regime—where iron lines are intrinsically weak—the neural networks first learn to constrain effective temperature ($T_{\rm eff}$) and surface gravity (log g) from the Balmer lines, then infer overall metallicity from the Mg lines (which correlate with [Fe/H]), and finally refine the [Fe/H] estimate using the weaker iron lines themselves. This process mirrors traditional spectroscopic analysis by human experts, particularly in the metal-poor regime.

These results demonstrate that the neural networks are not simply "black boxes" but can learn physically meaningful strategies for label inference such as [Fe/H]. They also suggest that joint training on additional labels (e.g., $T_{\rm eff}$, log g, and [Mg/Fe]) could reduce degeneracies in the [Fe/H] gradient maps, so that saliency maps would primarily highlight true iron lines, leading to more interpretable and reliable abundance estimations. A more detailed discussion of these topics, including the sign of the gradients, is left to future work.

Sensitivity is not emitted light.

C. DATASET SELECTION WORKFLOW

A summary of the cuts for DESI DR1 and DESI EDR samples and resulting sample sizes, adopted in the main analysis, is provided in Table C3 (see also Section 4).

D. COMPARISON WITH CLEAN CALIBRATED DESI SP SUBSET

In this section, we present the [Fe/H] and [α/Fe] scatter plots obtained after applying a more restrictive set of quality cuts to the DESI SP sample. Following the quality cut process of S. E. Koposov et al. (2025) for comparisons between DESI and external catalogs, we impose the additional selections: `BESTGRID` $\neq$ `s_rdesi1`, `RR_SPECTYPE` = `STAR`, `SN_R` > 10, and `FEH_ERR` < 0.1, where `FEH_ERR` is the [Fe/H] uncertainty reported by the DESI RVS pipeline.

Saliency is a sensitivity diagnostic, not proof that the network follows a sequential human reasoning process.

Abundance sensitivity peaks align with Balmer, magnesium, and iron features.

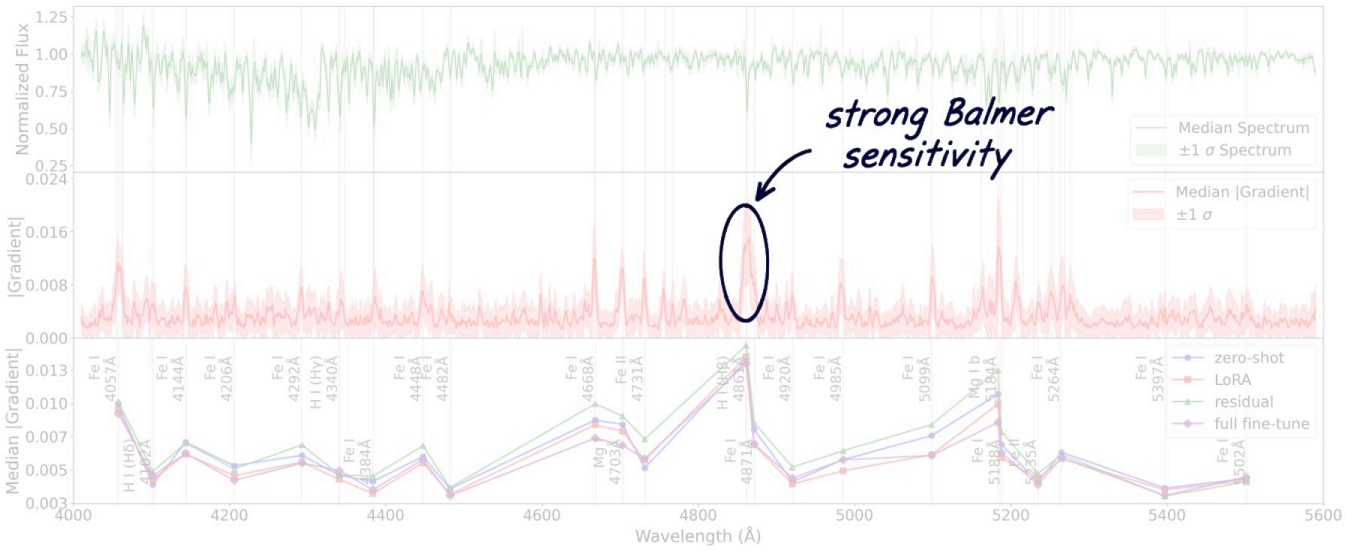


Figure B6. Saliency maps of the output iron abundance with respect to the input spectra in the rest frame. The upper panel shows the median and 1σ dispersion of the normalized spectra for the testing metal-poor subset ([Fe/H]). The middle panel presents the saliency maps (absolute gradients; median and 1σ) for the residual-head fine-tuning, indicating the sensitivity of the output iron abundance to different spectral regions. The lower panel highlights selected spectral lines with prominent sensitivity, comparing all fine-tuning strategies and the zero-shot results. Lines (from A. Kramida et al. 2024) corresponding to major spikes are shown in grey (iron) and red (non-iron) in the upper and middle panels, and selectively plotted in the lower panel.

Absolute gradients measure local sensitivity, not its sign or a causal reasoning sequence.

Table C3. Summary of DESI DR1 and DESI EDR samples adopted in the main analysis

Sample	Selection Criteria	Size	Notes
DESI DR1			
Initial cross-match	APOGEE: {STARFLAG,FE_H_FLAG,MG_FE_FLAG}=0; SNR> 40		
	DESI: RVS_WARN=0; when available, PRIMARY=True	8,104	
Training sample	$T_{\rm eff} > 4500$ K, 469 metal-poor + 1,600 metal-rich	2,069	For fine-tuning or scratch training
Testing sample	Remaining stars with $T_{\rm eff} > 4000$ K	5,539	For testing only
[α/Fe]–[Fe/H] reference	$T_{\rm eff} > 4000$ K after initial cross-match	7,608	For APOGEE [α/Fe]–[Fe/H] reference
DESI EDR			
Initial quality cuts	RVS_WARN=0, PRIMARY=True, RR_SPECTYPE=STAR		
	MEDIAN_COADD_SNR_B>20	134,137	
Final EDR sample	$T_{\rm eff} > 4000$ K; valid [Fe/H],[α/Fe]; good normalization	126,253	For evaluating [α/Fe]–[Fe/H]

The ledger separates 2,069 training stars from 5,539 held-out test stars.

Following S. E. Koposov et al. (2025), we further apply the recommended calibration to the DESI SP [Fe/H] values, noting that metallicity biases with respect to APOGEE and GALAH in the original DESI catalog show temperature-dependent trends that are well described by quadratic functions of $T_{\rm eff}$, with distinct relations for dwarfs and giants. The resulting comparisons are shown in Figure D7. Relative to Figure 2, approximately 2,100 stars are removed by these stricter quality cuts. In addition, some stars have [Fe/H] reported as NaN after the calibration step and therefore do not appear in the scatter plot.

For the [α/Fe]–[Fe/H] diagrams, we also provide a version based on cross-matching the selected EDR sample from the main text with DR1 sources that have PRIMARY=True, in order to obtain the most up-to-date [Fe/H] and [α/Fe] values (noting that EDR and DR1 use the same SP pipeline). We then apply the additional requirement BESTGRID

Fine-tuned MLPs still lead on the cleaner, calibrated DESI comparison subset.

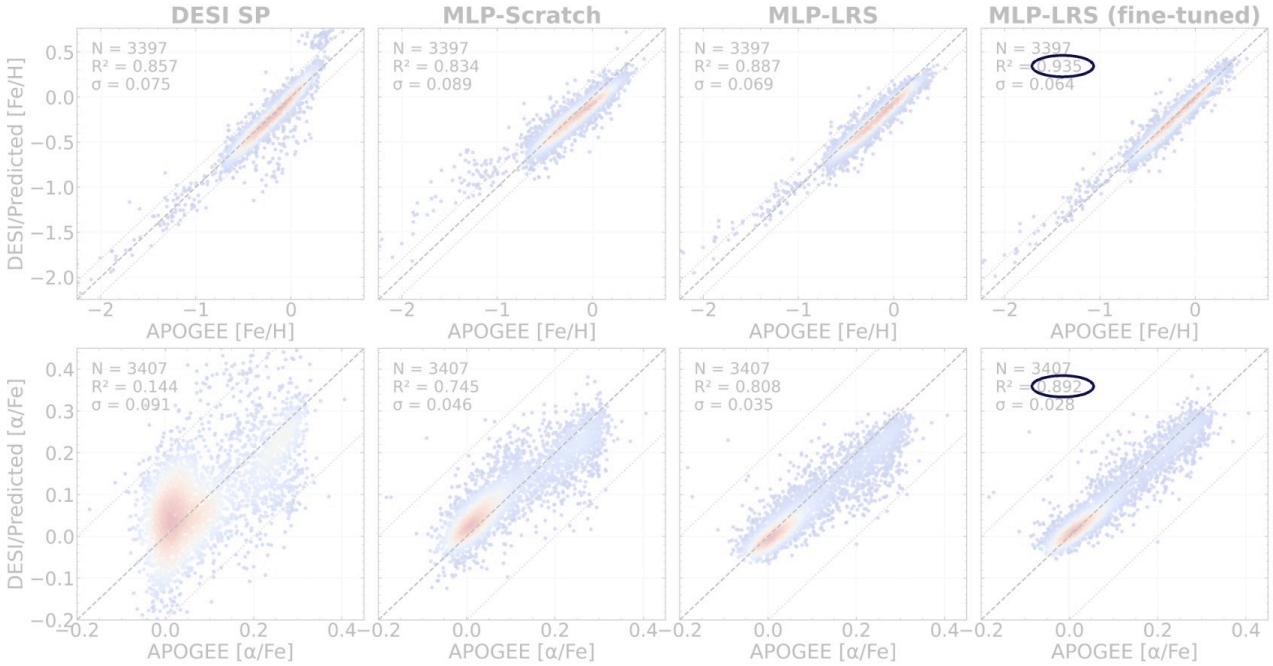


Figure D7. Comparison of [Fe/H] and [α/Fe] estimates between the clean, calibrated DESI SP subset and MLP-based models. Plotting conventions follow Figure 2.

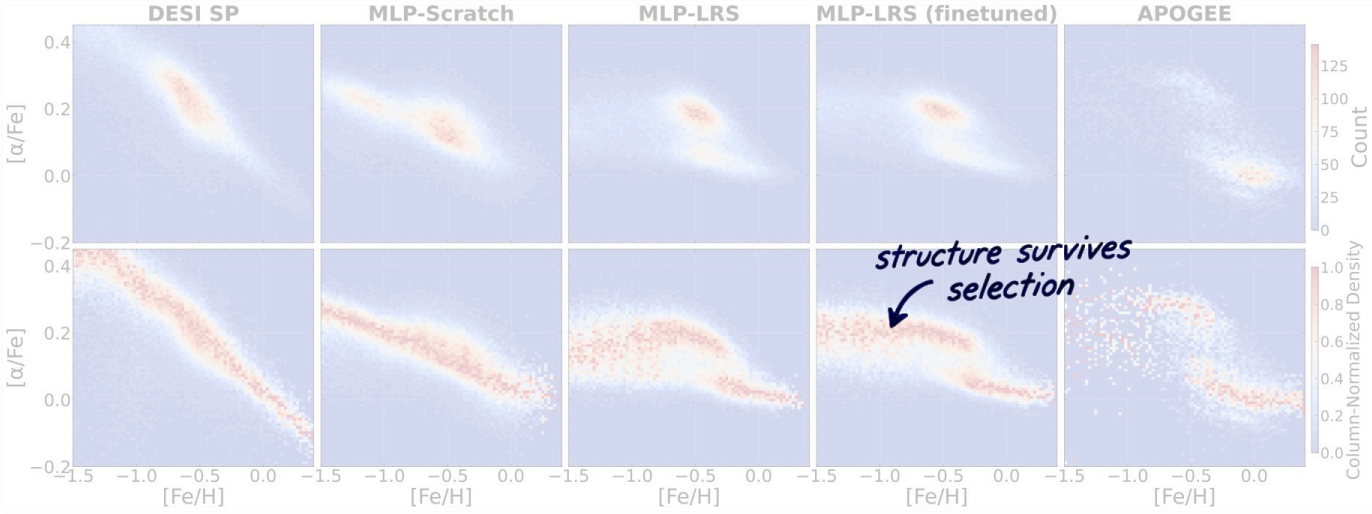


Figure D8. [α/Fe]–[Fe/H] diagrams for the DESI EDR and the APOGEE reference sample. A clean, calibrated DESI SP subset is used, with common sources shown in the second, third, and fourth columns. Plotting conventions follow Figure 3.

$\neq$ `s_rdesi1` and adopt the same [Fe/H] calibration as above. After these selections, 86,412 sources with valid stellar parameters remain from the originally selected EDR sample. The resulting comparison is shown in Figure D8.

Population structure complements the abundance scatter test: a tight relation alone does not establish correct disk sequences.

Stricter cuts remove roughly 2,100 comparison stars; the benchmark now covers a more selective sample.

REFERENCES

Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, The Seventeenth Data Release of the Sloan Digital Sky Surveys: Complete Release of MaNGA, MaStar, and APOGEE-2 Data, ApJS, 259, 35, doi: 10.3847/1538-4365/ac4414

Allende-Prieto, C., & Apogee Team. 2023, FERRE: Match physical models to measurements,, Astrophysics Source Code Library, record ascl:2301.016 http://ascl.net/2301.016

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, Astropy: A community Python package for astronomy, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Blanco-Cuaresma, S. 2019, Modern stellar spectroscopy caveats, MNRAS, 486, 2075, doi: 10.1093/mnras/stz549

Brown, T., Mann, B., Ryder, N., et al. 2020, Language Models are Few-Shot Learners, in Advances in Neural Information Processing Systems, ed. H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin, Vol. 33 (Curran Associates, Inc.), 1877–1901

Buck, T., & Schwarz, C. 2024, Deep Multimodal Representation Learning for Stellar Spectra, arXiv e-prints, arXiv:2410.16081, doi: 10.48550/arXiv.2410.16081

Chiba, M., Wyse, R. F. G., Kirby, E. N., et al. 2026, Galactic Archaeology with the Subaru 'Ōnohi'ula Prime Focus Spectrograph Strategic Program, arXiv e-prints, arXiv:2604.09875, doi: 10.48550/arXiv.2604.09875

De Angeli, F., Weiler, M., Montegriffo, P., et al. 2023, Gaia Data Release 3. Processing and validation of BP/RP low-resolution spectral data, A&A, 674, A2, doi: 10.1051/0004-6361/202243680

de Jong, J. T. A., Yanny, B., Rix, H.-W., et al. 2010, Mapping the Stellar Structure of the Milky Way Thick Disk and Halo Using SEGUE Photometry, ApJ, 714, 663, doi: 10.1088/0004-637X/714/1/663

De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, The GALAH survey: scientific motivation, MNRAS, 449, 2604, doi: 10.1093/mnras/stv327

DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, The DESI Experiment Part I: Science,Targeting, and Survey Design, arXiv e-prints, arXiv:1611.00036, doi: 10.48550/arXiv.1611.00036

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. 2018, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, arXiv e-prints, arXiv:1810.04805, doi: 10.48550/arXiv.1810.04805

Devon Hjelm, R., Fedorov, A., Lavoie-Marchildon, S., et al. 2018, Learning deep representations by mutual information estimation and maximization, arXiv e-prints, arXiv:1808.06670, doi: 10.48550/arXiv.1808.06670

Euclid Collaboration, Siudek, M., Huertas-Company, M., et al. 2025, Euclid Quick Data Release (Q1) Exploring galaxy properties with a multi-modal foundation model, arXiv e-prints, arXiv:2503.15312, doi: 10.48550/arXiv.2503.15312

Fitzpatrick, M. J., Olsen, K., Economou, F., et al. 2014, The NOAO Data Laboratory: a conceptual overview, in Observatory Operations: Strategies, Processes, and Systems V, ed. A. B. Peck, C. R. Benn, & R. L. Seaman, Vol. 9149, International Society for Optics and Photonics (SPIE), 91491T, doi: 10.1117/12.2057445

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, Gaia Data Release 3. Summary of the content and survey properties, A&A, 674, A1, doi: 10.1051/0004-6361/202243940

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, ASPCAP: The APOGEE Stellar Parameter and Chemical Abundances Pipeline, AJ, 151, 144, doi: 10.3847/0004-6256/151/6/144

Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., & Smola, A. 2012, A Kernel Two-Sample Test, Journal of Machine Learning Research, 13, 723

Hu, E. J., Shen, Y., Wallis, P., et al. 2021, LoRA: Low-Rank Adaptation of Large Language Models, arXiv e-prints, arXiv:2106.09685, doi: 10.48550/arXiv.2106.09685

Huang, Y., Beers, T. C., Xiao, K., et al. 2024, J-PLUS: Beyond Spectroscopy. III. Stellar Parameters and Elemental-abundance Ratios for Five Million Stars from DR3, ApJ, 974, 192, doi: 10.3847/1538-4357/ad6b94

Juneau, S., Olsen, K., Nikutta, R., Jacques, A., & Bailey, S. 2021, Jupyter-Enabled Astrophysical Analysis Using Data-Proximate Computing Platforms, Computing in Science & Engineering, 23, 15, doi: 10.1109/MCSE.2021.3057097

Juneau, S., Jacques, A., Pothier, S., et al. 2024, SPARCL: SPectra Analysis and Retrievable Catalog Lab, arXiv e-prints, arXiv:2401.05576, doi: 10.48550/arXiv.2401.05576

Koblischke, N., & Bovy, J. 2024, SpectraFM: Tuning into Stellar Foundation Models, arXiv e-prints, arXiv:2411.04750, doi: 10.48550/arXiv.2411.04750

Koposov, S. E., Allende Prieto, C., Cooper, A. P., et al. 2024, DESI Early Data Release Milky Way Survey value-added catalogue, MNRAS, 533, 1012, doi: 10.1093/mnras/stae1842

Koposov, S. E., Li, T. S., Allende Prieto, C., et al. 2025, DESI Data Release 1: Stellar Catalogue, arXiv e-prints, arXiv:2505.14787, doi: 10.48550/arXiv.2505.14787

Kramida, A., Ralchenko, Y., Reader, J., & NIST ASD Team. 2024, NIST Atomic Spectra Database,, https://physics.nist.gov/asd Gaithersburg, MD: National Institute of Standards and Technology, doi: 10.18434/T4W30F

Leung, H. W., & Bovy, J. 2024, Towards an astronomical foundation model for stars with a transformer-based model, MNRAS, 527, 1494, doi: 10.1093/mnras/stad3015

Li, H., Xu, Z., Taylor, G., Studer, C., & Goldstein, T. 2017, Visualizing the Loss Landscape of Neural Nets, arXiv e-prints, arXiv:1712.09913, doi: 10.48550/arXiv.1712.09913

Li, H., Aoki, W., Matsuno, T., et al. 2022, Four-hundred Very Metal-poor Stars Studied with LAMOST and Subaru. II. Elemental Abundances, ApJ, 931, 147, doi: 10.3847/1538-4357/ac6514

Loshchilov, I., & Hutter, F. 2017, Decoupled Weight Decay Regularization, arXiv e-prints, arXiv:1711.05101, doi: 10.48550/arXiv.1711.05101

Loshchilov, I., & Hutter, F. 2019, Decoupled Weight Decay Regularization, in International Conference on Learning Representations. https://openreview.net/forum?id=Bkg6RiCqY7

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, The Apache Point Observatory Galactic Evolution Experiment (APOGEE), AJ, 154, 94, doi: 10.3847/1538-3881/aa784d

Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015, The Cannon: A data-driven approach to Stellar Label Determination, ApJ, 808, 16, doi: 10.1088/0004-637X/808/1/16

Nikutta, R., Fitzpatrick, M., Scott, A., & Weaver, B. 2020, Data Lab—A community science platform, Astronomy and Computing, 33, 100411, doi: https://doi.org/10.1016/j.ascom.2020.100411

Parker, L., Lanusse, F., Golkar, S., et al. 2024, AstroCLIP: a cross-modal foundation model for galaxies, MNRAS, 531, 4990, doi: 10.1093/mnras/stae1450

Paszke, A., Gross, S., Massa, F., et al. 2019, PyTorch: An Imperative Style, High-Performance Deep Learning Library, in Advances in Neural Information Processing Systems 32 (Curran Associates, Inc.), 8024–8035

Pattnaik, R., Kartaltepe, J. S., & Binu, C. 2025, SpecPT (Spectroscopy Pre-trained Transformer) Model for Extragalactic Spectroscopy: I. Architecture and Automated Redshift Measurement, arXiv e-prints, arXiv:2501.01070, doi: 10.48550/arXiv.2501.01070

Popescu, M.-C., Balas, V. E., Perescu-Popescu, L., & Mastorakis, N. 2009, Multilayer perceptron and neural networks, WSEAS Transactions on Circuits and Systems, 8, 579

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. 2018, Improving Language Understanding by Generative Pre-Training, Tech. rep., OpenAI. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

Radford, A., Wu, J., Child, R., et al. 2019, Language Models are Unsupervised Multitask Learners, Tech. rep., OpenAI. https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf

Rizhko, M., & Bloom, J. S. 2024, AstroM^3: A self-supervised multimodal model for astronomy, arXiv e-prints, arXiv:2411.08842, doi: 10.48550/arXiv.2411.08842

Różański, T., Ting, Y.-S., & Jabłońska, M. 2025, TransformerPayne: Enhancing Spectral Emulation Accuracy and Data Efficiency by Capturing Long-range Correlations, ApJ, 980, 66, doi: 10.3847/1538-4357/ad9b99

Sestito, F., Longeard, N., Martin, N. F., et al. 2019, Tracing the formation of the Milky Way through ultra metal-poor stars, MNRAS, 484, 2166, doi: 10.1093/mnras/stz043

Simonyan, K., Vedaldi, A., & Zisserman, A. 2013, Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps, arXiv e-prints, arXiv:1312.6034. https://arxiv.org/abs/1312.6034

Smith, M. J., Roberts, R. J., Angeloudi, E., & Huertas-Company, M. 2024, AstroPT: Scaling Large Observation Models for Astronomy, arXiv e-prints, arXiv:2405.14930, doi: 10.48550/arXiv.2405.14930

Soubiran, C., Le Campion, J.-F., Brouillet, N., & Chemin, L. 2016, The PASTEL catalogue: 2016 version, A&A, 591, A118, doi: 10.1051/0004-6361/201628497

Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, The Radial Velocity Experiment (RAVE): First Data Release, AJ, 132, 1645, doi: 10.1086/506564

Suda, T., Katsuta, Y., Yamada, S., et al. 2008, Stellar Abundances for the Galactic Archeology (SAGA) Database — Compilation of the Characteristics of Known Extremely Metal-Poor Stars, PASJ, 60, 1159, doi: 10.1093/pasj/60.5.1159

Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2019, The Payne: Self-consistent ab initio Fitting of Stellar Spectra, ApJ, 879, 69, doi: 10.3847/1538-4357/ab2331

Ting, Y.-S., Rix, H.-W., Conroy, C., Ho, A. Y. Q., & Lin, J. 2017, Measuring 14 Elemental Abundances with R = 1800 LAMOST Spectra, ApJL, 849, L9, doi: 10.3847/2041-8213/aa921c

Vaswani, A., Shazeer, N., Parmar, N., et al. 2017, Attention is All you Need, in Advances in Neural Information Processing Systems, ed. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett, Vol. 30 (Curran Associates, Inc.)

Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, Abundance Estimates for 16 Elements in 6 Million Stars from LAMOST DR5 Low-Resolution Spectra, ApJS, 245, 34, doi: 10.3847/1538-4365/ab5364

Zhang, M., Xiang, M., Ting, Y.-S., et al. 2024, Determining Stellar Elemental Abundances from DESI Spectra with the Data-driven Payne, ApJS, 273, 19, doi: 10.3847/1538-4365/ad51dd

Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., & Deng, L.-C. 2012, LAMOST spectral survey — An overview, Research in Astronomy and Astrophysics, 12, 723, doi: 10.1088/1674-4527/12/7/002

Zhao, X., & Huang, Y. 2025, SpecCLIP v1.0.0: Aligning and Translating Spectroscopic Measurements for Stars, v1.0.0 Zenodo, doi: 10.5281/zenodo.17824840

Zhao, X., Ting, Y.-S., Szalay, A. S., & Huang, Y. 2025a, Finetuning Stellar Spectra Foundation Models with LoRA, arXiv e-prints, arXiv:2507.20972. https://arxiv.org/abs/2507.20972

Zhao, X., Huang, Y., Xue, G., et al. 2025b, SpecCLIP: Aligning and Translating Spectroscopic Measurements for Stars, arXiv e-prints, arXiv:2507.01939, doi: 10.48550/arXiv.2507.01939

Zhong, F., Napolitano, N. R., Heneka, C., et al. 2024, Galaxy Spectra Networks (GaSNet). III. Generative pre-trained network for spectrum reconstruction, redshift estimate and anomaly detection, arXiv e-prints, arXiv:2412.21130, doi: 10.48550/arXiv.2412.21130