

---

# Finetuning Stellar Spectra Foundation Models with LoRA

---

Xiaosheng Zhao<sup>1</sup> Yuan-Sen Ting<sup>2</sup> Alexander S. Szalay<sup>1</sup> Yang Huang<sup>3</sup>

## Abstract

Foundation models are beginning to impact stellar spectroscopy, where spectra encode rich physical information in a structured, language-like form. A key challenge is adapting these models across heterogeneous surveys with differing resolution and coverage. We apply Low-Rank Adaptation (LoRA) to fine-tune SpecCLIP—a contrastively pre-trained model on LAMOST and Gaia XP spectra—for downstream tasks on DESI Early Data Release (EDR) spectra. We show that LoRA enables few-shot learning on DESI, with performance varying by fine-tuned module and benefiting from Gaia XP knowledge embedded in the pre-trained model. Our results demonstrate that LoRA provides a lightweight and effective strategy for extending spectral foundation models to new instruments and survey domains.

## 1. Introduction

In recent decades, large-scale spectroscopic surveys have transformed our understanding of the Milky Way’s formation and evolution. This progress has been driven by two parallel advances: the continual expansion of spectroscopic datasets and the development of increasingly sophisticated methods to infer stellar properties from spectra. These methods range from traditional template-matching algorithms like UlySS (Koleva et al., 2009) and LSP3 (Xiang et al., 2015) to machine learning techniques such as The Cannon (Ness et al., 2015), The Payne (Ting et al., 2019; Xiang et al., 2019; Zhang et al., 2024) and TransformerPayne (Różański et al., 2025). Despite their success, most of these approaches rely on explicit supervision and are constrained either by

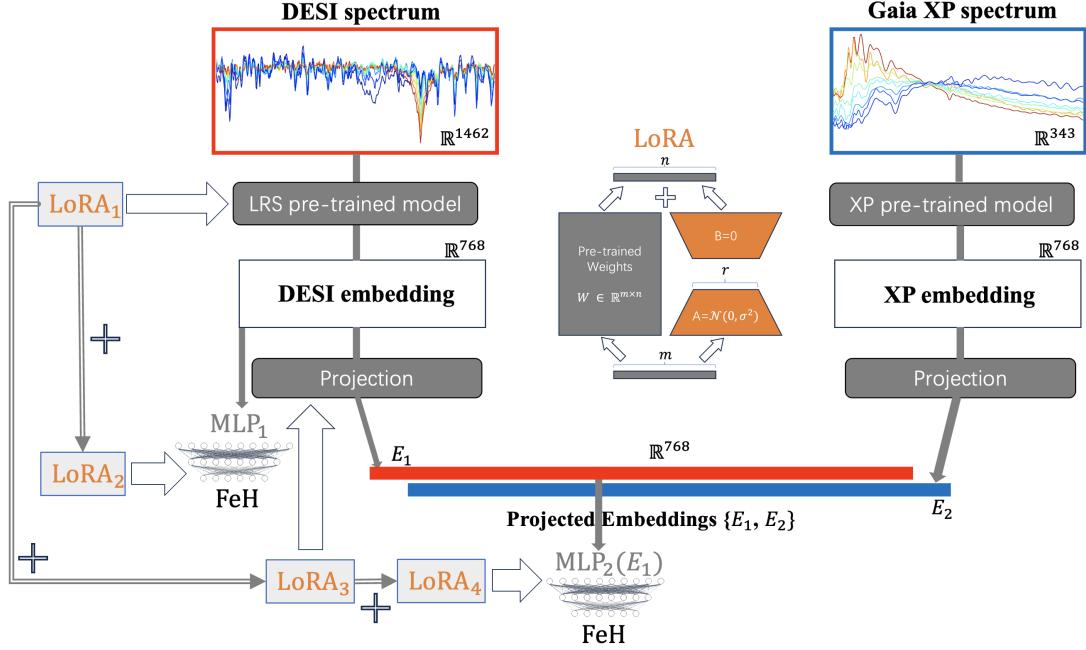
limited empirical templates or imperfect synthetic models. The challenge becomes more severe when combining heterogeneous spectra from surveys with different wavelength coverages, resolutions, and signal-to-noise ratios, making it difficult to achieve consistent, survey-independent parameter estimates.

Motivated by the recent success of large language models (LLMs) in learning general-purpose representations from large-scale data, we previously introduced SpecCLIP (Zhao et al., 2025), a spectral foundation model trained on LAMOST low-resolution spectra (LRS) and Gaia XP spectra using contrastive learning and some extensions to learn/arrange both shared and non-shared information in the embeddings. Inspired by Contrastive Language-Image Pre-training (CLIP) (Radford et al., 2021) and AstroCLIP (Parker et al., 2024), the first introduction of CLIP into astronomy, SpecCLIP aligns representations across surveys and enables cross-modal applications such as similarity search, parameter estimation, and cross-modal prediction. The pre-trained models exhibit strong few-shot performance, and they provide a unified framework for homogenizing stellar parameters across instruments.

In this follow-up study, we address a critical challenge in foundation model deployment: how to adapt pre-trained models to new spectroscopic domains with minimal supervision. Specifically, we explore the use of Low-Rank Adaptation (LoRA, Hu et al., 2021) to fine-tune different components of SpecCLIP on the DESI Early Data Release (EDR, Koposov et al., 2024) stellar spectra. To our knowledge, this is the first application of LoRA in stellar spectroscopy. We find that LoRA enables effective few-shot learning on DESI data, even though the original model was trained exclusively on LAMOST and Gaia XP spectra. Our analysis shows that tuning different modules of the pre-trained model yields varying performance gains, and that Gaia XP information embedded in the foundation model can enhance adaptation to DESI, despite differences in resolution and wavelength coverage, and alignment with LAMOST spectra originally. Our approach facilitates the rapid deployment of powerful pre-trained models across heterogeneous spectroscopic datasets, enabling consistent stellar parameter estimation in the era of large, diverse sky surveys.

<sup>1</sup>Department of Physics & Astronomy, The Johns Hopkins University, Baltimore, MD 21218, USA <sup>2</sup>Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA <sup>3</sup>School of Astronomy and Space Science, University of Chinese Academy of Sciences, Beijing 100049, People’s Republic of China. Correspondence to: Xiaosheng Zhao <xzhao113@jh.edu>.

*Proceedings of the 42<sup>nd</sup> International Conference on Machine Learning*, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).



**Figure 1.** Fine-tuning the SpecCLIP model with LoRA. Gray rounded blocks represent pre-trained modules: the LAMOST LRS and Gaia XP foundation models (trained separately on their respective modalities), and the projection networks trained with contrastive loss on paired projected embeddings. Two downstream MLPs, previously trained on LRS-based iron abundance labels, are also included. In the fine-tuning stage, a small number of DESI spectra are used to adapt: (1) the LRS-based MLP (**LoRA2**), (2) both the MLP and LRS foundation model (**LoRA1 + LoRA2**), (3) the projected-based MLP (**LoRA4**), or (4) the projected-based MLP, projection network, and optionally the LRS foundation model (**LoRA1 + LoRA3 + LoRA4**). The center inset illustrates the LoRA mechanism, which updates the model weights by injecting the product of two low-rank matrices as a residual into the frozen pre-trained network.

## 2. Method

### 2.1. Overview of SpecCLIP Foundation Models

SpecCLIP is a foundation model framework designed to align stellar spectra across modalities that differ in wavelength coverage, resolution, and signal-to-noise ratio.

In the based SpecCLIP model, each modality—such as LAMOST LRS (DR11<sup>1</sup>) and Gaia XP (Gaia Collaboration et al., 2023) spectra—is assigned its own transformer-based model with multi-head attention layers pre-trained using masked modeling (Devlin et al., 2018) or multilayer perceptron (MLP)-based autoencoder. This unsupervised pre-training retrain the most important information which is further aligned by the contrastive training followed. It allows the resulting embeddings to support downstream tasks such as stellar parameter estimation, similarity search, and cross-modal prediction, with moderate supervision.

**LAMOST LRS Foundation Model:** We select 966,082 (9 : 1 train-valid split, likewise, in the following cases) high-quality spectra from LAMOST DR11 for pre-training, applying quality cuts of G-band  $S/N_g > 50$  and absolute g-

band magnitude  $< 15.8$ . We restrict the wavelength range to 400–560 nm, yielding 1462 flux points sampled uniformly in logarithmic wavelength space per spectrum. These are normalized and segmented into overlapping tokens of size 20 with a stride of 10, resulting in 146 tokens per input. A transformer encoder with 6 layers processes these tokens to produce a spectral embedding. During pre-training, masked modeling is applied by masking 6 randomly selected chunks (each 20 tokens).

**Gaia XP Foundation Model:** We construct two types of models for Gaia XP spectra, which consist of 343 flux points spanning 336–1021 nm. For pre-training, we use 1,000,000 spectra. The input spectra are normalized by dividing by the V-band (550 nm) flux to provide a color-based inductive bias. Our model for Gaia XP spectra is an MLP-based autoencoder trained without masking. Both models (LRS and XP) share the same embedding dimension (768), number of parameters (42.7M), and number of training epochs.

**Contrastive Training:** We use a dataset of 820,568 paired spectra from LAMOST LRS and Gaia XP, where each spectrum is processed by its corresponding pre-trained foundation model. The resulting embeddings are then mapped into a shared latent space using modality-specific projection

<sup>1</sup><https://www.lamost.org/dr11/>

networks. Each projection head (with 7.1M parameters) includes a cross-attention block, where a learnable query vector attends to the sequence of spectral features to produce a 768-length representation for contrastive alignment.

**Iron Abundance Estimation:** We trained downstream MLP models to estimate stellar iron abundance ([Fe/H]) using 90,106 LAMOST LRS spectra, split into training and validation sets with a 9:1 ratio. The MLP takes as input either the spectral embeddings from the LAMOST LRS foundation model or the projected embeddings after contrastive alignment.

## 2.2. LoRA Fine-Tuning for Cross-Survey Adaptation

To adapt SpecCLIP to new spectroscopic domains with minimal supervision, we apply Low-Rank Adaptation (LoRA) to fine-tune the pre-trained models on stellar spectra from the DESI Early Data Release (EDR). LoRA introduces low-rank update matrices into pre-existing linear layers, allowing the original pre-trained weights to remain frozen while efficiently injecting task-specific information through a small number of additional parameters.

In essence, LoRA decomposes the weight update  $\Delta W \in \mathbb{R}^{m \times n}$  into a product of two low-rank matrices:  $\Delta W = AB$ , where  $A \in \mathbb{R}^{m \times r}$  and  $B \in \mathbb{R}^{r \times n}$  for a given rank  $r \ll \min(m, n)$ . This enables substantial parameter savings while maintaining expressive adaptation capacity. In our setup, LoRA modules are inserted into selected attention and/or feed-forward layers of the pre-trained transformer backbones. A schematic of LoRA is shown in Fig. 1 (middle panel).

**Fine-Tuning with LoRA.** Our final model architecture is shown in Fig. 1. We fine-tune four modules using Low-Rank Adaptation (LoRA): (1) the LAMOST-LRS foundation model (**LoRA1**), where LoRA is applied to all self-attention layers ( $\{\mathbf{q}, \mathbf{k}, \mathbf{v}, \mathbf{out}\}_{\text{proj}}$ ) with rank = 4 and scaling factor  $\alpha = 8$ , adding 129,024 trainable parameters (0.30% of the model); (2) the projection network used for contrastive alignment (**LoRA3**), where we apply LoRA to the  $\mathbf{out}_{\text{proj}}$  and subsequent linear layers with rank = 16,  $\alpha = 32$ , and 147,456 trainable parameters (0.29%); and (3) the downstream MLPs (**LoRA2** and **LoRA4**) trained on SpecCLIP embeddings, where LoRA is added to all linear layers with rank = 8,  $\alpha = 16$ , and 31,752 parameters (2.30% of the MLP size). We compare performance under different configurations and evaluate whether Gaia XP knowledge embedded in the pre-trained SpecCLIP model improves adaptation to DESI spectra without direct supervision. All experiments complete in around a few tens of seconds on a single NVIDIA V100 GPU.

## 3. Data

We base our experiments on the DESI Early Data Release (EDR) Milky Way Stars Value-Added Catalog, `mwsall-pix-fuji.fits`, which provides physical parameters and spectroscopic metadata for 625,588 sources. We extract stellar samples following a series of quality filters: we retain only rows where the spectral type is STAR, `RVS_WARN` is zero, the observation is marked as `PRIMARY`, signal-to-noise ratio (SNR) exceeds 20, celestial coordinates are finite, and the Gaia `SOURCE_ID` is valid (not 999999). This results in a clean sample of DESI spectra suitable for foundation model adaptation.

To provide ground-truth labels for downstream tasks, we cross-match this DESI sample with APOGEE DR17 (Abdurro’uf et al., 2022), yielding 495 stars with high-fidelity iron abundance ([Fe/H]) measurements. Among these, we randomly allocate 89 samples for training, 9 for validation, and the remaining 396 for testing. These are used to fine-tune **LoRA2** and **LoRA4**.

To fine-tune **LoRA1** and **LoRA3**, we exclude all 495 APOGEE-matched stars and randomly select an additional 200 unmatched DESI samples—100 for training and 100 for validation—with SNR further constrained to exceed 50. Although label-free, these spectra support contrastive learning and masked modeling objectives. While one could in principle use the same unlabeled DESI spectra for both foundation and downstream fine-tuning (due to the lack of label usage in the former), we deliberately use separate subsets in this study to make the adaptation task more challenging and test the generalization capacity of our method.

All DESI spectra, retrieved via the SPectra Analysis & Retrievable Catalog Lab (SPARCL, Juneau et al., 2024; Juneau et al., 2021) at NOIRLab’s Astro Data Lab (Fitzpatrick et al., 2014; Nikutta et al., 2020), are normalized using the same pipeline as for LAMOST LRS, and then interpolated onto the LAMOST wavelength grid spanning 400–560 nm. This standardization ensures compatibility between DESI and LAMOST representations.

For full SpecCLIP fine-tuning (combining LoRA1 and LoRA3), we further cross-match the selected DESI samples with the Gaia XP catalog to retrieve their fluxes. These are processed identically to the XP spectra in the original SpecCLIP paper: fluxes are divided by the value at 550 nm to yield a 343-dimensional color vector.

## 4. Results

We evaluate the effectiveness of LoRA-based fine-tuning under several configurations, comparing zero-shot performance with targeted adaptations to different parts of the SpecCLIP architecture. Table 1 summarizes the robust stan-

**Table 1.** Performance of iron abundance ([Fe/H]) prediction on the DESI EDR test set. We report robust standard deviation ( $\sigma$ , using Tukey Biweight Scale Estimator(Hoaglin et al., 1983)) of the residual and  $R^2$  scores for different fine-tuning configurations. Results are shown for the full test set and for two iron abundance subsets. Top three entries are highlighted with decreasing shades of blue. Numbers in brackets indicate the source counts.

Method	$\sigma$	$R^2$
<i>Full Test Set (396)</i>		
Zero-shot (MLP1)	0.2730	0.7358
LoRA2	0.2663	0.7156
LoRA1 + LoRA2	0.2227	0.7719
Zero-shot (MLP2)	0.2560	0.7203
LoRA4	0.2023	0.7937
LoRA1 + LoRA3 + LoRA4	0.2297	0.7801
<i>[Fe/H] &lt; -1 (60, metal-poor)</i>		
Zero-shot (MLP1)	0.4444	-0.5130
LoRA2	0.5872	-1.2881
LoRA1 + LoRA2	0.5151	-0.9143
Zero-shot (MLP2)	0.4658	-0.7579
LoRA4	0.5803	-0.8357
LoRA1 + LoRA3 + LoRA4	0.5970	-0.8159
<i>[Fe/H] ≥ -1 (336, metal-rich)</i>		
Zero-shot (MLP1)	0.2479	0.0702
LoRA2	0.2272	0.2378
LoRA1 + LoRA2	0.1924	0.4173
Zero-shot (MLP2)	0.2371	0.0725
LoRA4	0.1621	0.5106
LoRA1 + LoRA3 + LoRA4	0.1851	0.4274

dard deviation ( $\sigma$ ) and coefficient of determination ( $R^2$ ) on the DESI EDR test set, both overall and for subsets divided by iron abundance. Fig.2 shows the corresponding scatter plots for all experiments.

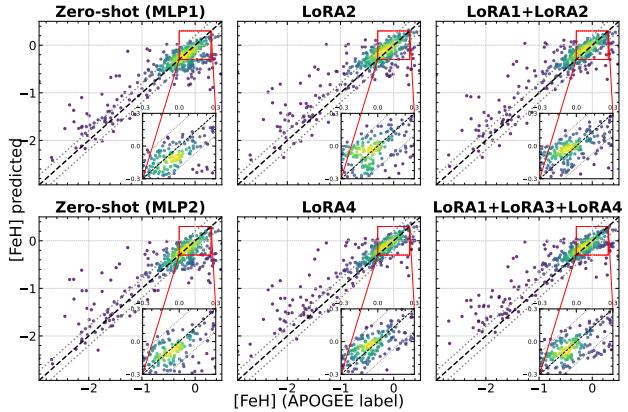
**Full test set.** The best overall performance is achieved by LoRA4, which fine-tunes only the downstream MLP for the Gaia XP-aligned embeddings, reaching the lowest robust scatter ( $\sigma = 0.2023$ ) and the highest  $R^2 = 0.7937$ . Close behind is the full multi-module adaptation (LoRA1 + LoRA3 + LoRA4), showing consistent improvement over zero-shot baselines. Notably, LoRA1 + LoRA2 also performs well, indicating that the foundation model adaptations contribute meaningfully.

**Metal-poor subset ( $[Fe/H] < -1$ ).** All methods struggle in this regime. Interestingly, the best performance comes from the zero-shot MLP1 model ( $\sigma = 0.4444$ ,  $R^2 = -0.5130$ ), suggesting that fine-tuning may overfit or underperform in this sparse label regime. All LoRA-enhanced variants yield higher  $\sigma$  and more negative  $R^2$ , particularly LoRA2 and LoRA1 + LoRA3 + LoRA4, which degrade the predictive performance ( $\sigma \sim 0.58\text{--}0.60$ ,  $R^2 < -0.8$ ). This highlights the difficulty of generalizing to metal-poor stars with limited

data.

**Metal-rich subset ( $[Fe/H] \geq -1$ ).** In contrast, all fine-tuned models significantly outperform the zero-shot baselines. LoRA4 achieves the best results ( $\sigma = 0.1621$ ,  $R^2 = 0.5106$ ), with LoRA1 + LoRA3 + LoRA4 and LoRA1 + LoRA2 close behind. This suggests that LoRA fine-tuning is highly effective for the dominant population of metal-rich stars in DESI, where the training distribution is better matched.

Overall, these results demonstrate that LoRA offers a flexible and lightweight way to adapt foundation models to new survey domains. Fine-tuning the downstream MLP alone (LoRA2 or LoRA4) is effective, but further gains can be achieved by jointly adapting the foundation model. However, the poor performance in the metal-poor regime suggests the need for better regularization or targeted augmentation to ensure robust generalization in low-density regions of label space.



**Figure 2.** Scatter plots of predicted versus APOGEE iron abundance ([Fe/H]) for all fine-tuning configurations, with point density indicated by the viridis colormap. The dashed black line indicates the ideal 1:1 relation, and the gray dotted lines mark  $\pm 0.2$  dex deviations. An inset zooms into the  $\pm 0.3$  dex region for closer inspection. LoRA fine-tuning improves overall predictive accuracy compared to zero-shot baselines, particularly in the metal-rich regime ( $[Fe/H] \geq -1$ ). However, performance in the metal-poor regime ( $[Fe/H] < -1$ ) remains limited, with larger scatter and systematic deviations across all configurations.

## 5. Summary

We presented the first application of Low-Rank Adaptation (LoRA) to fine-tune a spectral foundation model for stellar spectroscopy. Building on the SpecCLIP framework, we explored how LoRA can efficiently adapt different modules—pre-trained LRS foundation model, projection heads, and downstream MLPs—to the DESI EDR stellar spectra with only a few (100 to 200) labeled examples.

Our experiments demonstrate that LoRA fine-tuning leads to

noticeable improvements over zero-shot baselines, particularly in the metal-rich regime. We find that even lightweight adaptation of the downstream predictor achieves strong performance, while further gains can be achieved by jointly tuning the foundation model layers. However, performance on metal-poor stars remains limited, suggesting that additional strategies—such as better regularization, improved sample balancing for fine-tuning, careful hyperparameter optimization, or uncertainty-aware modeling—may be necessary to enhance generalization in this regime.

These findings suggest that LoRA provides a scalable and effective pathway for extending pre-trained spectral foundation models to new surveys, enabling flexible and data-efficient inference in heterogeneous spectroscopic environments.

## Broader Impact

This paper presents work whose goal is to advance the field of Machine Learning in the context of astrophysical data analysis. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## Acknowledgements

We acknowledge support through the generosity of Eric and Wendy Schmidt, by recommendation of the Schmidt Futures program. 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). This research uses services or data provided by the SPectra Analysis and Retrievable Catalog Lab (SPARCL) and the Astro Data Lab, which are both part of the Community Science and Data Center (CSDC) Program of NSF NOIRLab. NOIRLab is operated by the Association of Universities for Research in Astronomy (AURA), Inc. under a cooperative agreement with the U.S. National Science Foundation.

## References

- Abdurro’uf, Accetta, K., Aerts, C., Silva Aguirre, V., Ahumada, R., Ajgaonkar, N., Filiz Ak, N., Alam, S., Allende Prieto, C., Almeida, A., Anders, F., Anderson, S. F., Andrews, B. H., Anguiano, B., Aquino-Ortíz, E., Aragón-Salamanca, A., Argudo-Fernández, M., Ata, M., Aubert, M., Avila-Reese, V., Badenes, C., Barbá, R. H., Barger, K., Barrera-Ballesteros, J. K., Beaton, R. L., Beers, T. C., Belfiore, F., Bender, C. F., Bernardi, M., Bershadsky, M. A., Beutler, F., Bidin, C. M., Bird, J. C., Bizyaev, D., Blanc, G. A., Blanton, M. R., Boardman, N. F., Bolton, A. S., Bouquin, M., Borissova, J., Bovy, J., Brandt, W. N., Brown, J., Brownstein, J. R., Brusa, M., Buchner, J., Bundy, K., Burchett, J. N., Bureau, M., Burgasser, A., Cabang, T. K., Campbell, S., Cappellari, M., Carlberg, J. K., Wanderley, F. C., Carrera, R., Cash, J., Chen, Y.-P., Chen, W.-H., Cherinka, B., Chiappini, C., Choi, P. D., Chojnowski, S. D., Chung, H., Clerc, N., Cohen, R. E., Comerford, J. M., Comparat, J., da Costa, L., Covey, K., Crane, J. D., Cruz-Gonzalez, I., Culhane, C., Cunha, K., Dai, Y. S., Damke, G., Darling, J., Davidson, Jr., J. W., Davies, R., Dawson, K., De Lee, N., Diamond-Stanic, A. M., Cano-Díaz, M., Sánchez, H. D., Donor, J., Duckworth, C., Dwelly, T., Eisenstein, D. J., Elsworth, Y. P., Emsellem, E., Eracleous, M., Escoffier, S., Fan, X., Farr, E., Feng, S., Fernández-Trincado, J. G., Feuillet, D., Filipp, A., Fillingham, S. P., Frinchaboy, P. M., Fromenteau, S., Galbany, L., García, R. A., García-Hernández, D. A., Ge, J., Geisler, D., Gelfand, J., Géron, T., Gibson, B. J., Goddy, J., Godoy-Rivera, D., Grabowski, K., Green, P. J., Greener, M., Grier, C. J., Griffith, E., Guo, H., Guy, J., Hadjara, M., Harding, P., Hasselquist, S., Hayes, C. R., Hearty, F., Hernández, J., Hill, L., Hogg, D. W., Holtzman, J. A., Horta, D., Hsieh, B.-C., Hsu, C.-H., Hsu, Y.-H., Huber, D., Huertas-Company, M., Hutchinson, B., Hwang, H. S., Ibarra-Medel, H. J., Chitham, J. I., Ilha, G. S., Imig, J., Jaekle, W., Jayasinghe, T., Ji, X., Johnson, J. A., Jones, A., Jönsson, H., Katkov, I., Khalatyan, Dr., A., Kinemuchi, K., Kisku, S., Knapen, J. H., Kneib, J.-P., Kollmeier, J. A., Kong, M., Kounkel, M., Kreckel, K., Krishnarao, D., Lacerna, I., Lane, R. R., Langgin, R., Lavender, R., Law, D. R., Lazarz, D., Leung, H. W., Leung, H.-H., Lewis, H. M., Li, C., Li, R., Lian, J., Liang, F.-H., Lin, L., Lin, Y.-T., Lin, S., Lintott, C., Long, D., Longa-Peña, P., López-Cobá, C., Lu, S., Lundgren, B. F., Luo, Y., Mackereth, J. T., de la Macorra, A., Mahadevan, S., Majewski, S. R., Manchado, A., Mandeville, T., Maraston, C., Margalef-Bentabol, B., Masseron, T., Masters, K. L., Mathur, S., McDermid, R. M., McKay, M., Merloni, A., Merrifield, M., Meszaros, S., Miglio, A., Di Mille, F., Minniti, D., Minsley, R., and Monachesi, A. The Seventeenth Data Release of the Sloan Digital Sky Surveys: Complete Release of MaNGA, MaStar, and APOGEE-2 Data. *ApJS*, 259(2):35, April 2022. doi: 10.3847/1538-4365/ac4414.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv e-prints*, art. arXiv:1810.04805, October 2018. doi: 10.48550/arXiv.1810.04805.
- Fitzpatrick, M. J., Olsen, K., Economou, F., Stobie, E. B., Beers, T. C., Dickinson, M., Norris, P., Saha, A., Seaman, R., Silva, D. R., Swaters, R. A., Thomas, B., and Valdes, F. The NOAO Data Laboratory: a conceptual overview. In Peck, A. B., Benn, C. R., and Seaman,

- R. L. (eds.), *Observatory Operations: Strategies, Processes, and Systems V*, volume 9149, pp. 91491T. International Society for Optics and Photonics, SPIE, 2014. doi: 10.1117/12.2057445. URL <https://doi.org/10.1117/12.2057445>.
- Gaia Collaboration, Montegriffo, P., Bellazzini, M., De Angeli, F., Andrae, R., Barstow, M. A., Bossini, D., Bragaglia, A., Burgess, P. W., Cacciari, C., Carrasco, J. M., Chornay, N., Delchambre, L., Evans, D. W., Fouesneau, M., Frémat, Y., Garabato, D., Jordi, C., Manteiga, M., Massari, D., Palaversa, L., Pancino, E., Riello, M., Ruz Mieres, D., Sanna, N., Santoveña, R., Sordo, R., Valenari, A., Walton, N. A., Brown, A. G. A., Prusti, T., de Bruijne, J. H. J., Arenou, F., Babusiaux, C., Biermann, M., Creevey, O. L., Ducourant, C., Eyer, L., Guerra, R., Hutton, A., Klioner, S. A., Lammers, U. L., Lindegren, L., Luri, X., Mignard, F., Panem, C., Pourbaix, D., Randich, S., Sartoretti, P., Soubiran, C., Tanga, P., Bailer-Jones, C. A. L., Bastian, U., Drimmel, R., Jansen, F., Katz, D., Lattanzi, M. G., van Leeuwen, F., Bakker, J., Castañeda, J., Fabricius, C., Galluccio, L., Guerrier, A., Heiter, U., Masana, E., Messineo, R., Mowlavi, N., Nicolas, C., Nienartowicz, K., Pailler, F., Panuzzo, P., Riclet, F., Roux, W., Seabroke, G. M., Thévenin, F., Gracia-Abril, G., Portell, J., Teyssier, D., Altmann, M., Audard, M., Bellas-Velidis, I., Benson, K., Berthier, J., Blomme, R., Busonero, D., Busso, G., Cánovas, H., Carry, B., Cellino, A., Cheek, N., Clementini, G., Damerdji, Y., Davidson, M., de Teodoro, P., Nuñez Campos, M., Dell’Oro, A., Esquej, P., Fernández-Hernández, J., Fraile, E., García-Lario, P., Gosset, E., Haigron, R., Halbwachs, J. L., Hambly, N. C., Harrison, D. L., Hernández, J., Hestroffer, D., Hodgkin, S. T., Holl, B., Janßen, K., Jevardat de Fombelle, G., Jordan, S., Krone-Martins, A., Lanzafame, A. C., Löffler, W., Marchal, O., Marrese, P. M., Moitinho, A., Muinonen, K., Osborne, P., Pauwels, T., Recio-Blanco, A., Reylé, C., Rimoldini, L., Roegiers, T., Rybizki, J., Sarro, L. M., Siopis, C., Smith, M., Sozzetti, A., Utrilla, E., van Leeuwen, M., Abbas, U., Ábrahám, P., Abreu Aramburu, A., Aerts, C., Aguado, J. J., Ajaj, M., Aldea-Montero, F., Altavilla, G., Álvarez, M. A., Alves, J., Anderson, R. I., Anglada Varela, E., Antoja, T., Baines, D., Baker, S. G., Balaguer-Núñez, L., Balbinot, E., Balog, Z., Barache, C., Barbato, D., Barros, M., Bartolomé, S., Bassilana, J. L., Bauchet, N., Becciani, U., Berihuete, A., Bernet, M., Bertone, S., Bianchi, L., Binnenfeld, A., Blanco-Cuaresma, S., Boch, T., Bombrun, A., Bouquillon, S., Bramante, L., Breedt, E., Bressan, A., Brouillet, N., Brugaletta, E., Bucciarelli, B., Burlacu, A., Butkevich, A. G., Buzzi, R., Caffau, E., Cancilliere, R., Cantat-Gaudin, T., Carballo, R., Carlucci, T., Carnerero, M. I., Casamiquela, L., Castellani, M., Castro-Ginard, A., Chaoul, L., Charlot, P., Chemin, L., Chiaramida, V., Chiavassa, A., Comoretto, G., Contursi, G., Cooper, W. J., Cornez, T., Cowell, S., Crifo, F., Cropper, M., Crosta, M., Crowley, C., Dafonte, C., and Dapergolas, A. Gaia Data Release 3. The Galaxy in your preferred colours: Synthetic photometry from Gaia low-resolution spectra. *A&A*, 674:A33, June 2023. doi: 10.1051/0004-6361/202243709.
- Hoaglin, D. C., Mosteller, F., and Tukey, J. W. *Understanding Robust and Exploratory Data Analysis*. John Wiley & Sons, New York, 1983.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. *arXiv e-prints*, art. arXiv:2106.09685, June 2021. doi: 10.48550/arXiv.2106.09685.
- Juneau, S., Olsen, K., Nikutta, R., Jacques, A., and Bailey, S. Jupyter-enabled astrophysical analysis using data-approximate computing platforms. *Computing in Science & Engineering*, 23(2):15–25, 2021. doi: 10.1109/MCSE.2021.3057097.
- Juneau, S., Jacques, A., Pothier, S., Bolton, A. S., Weaver, B. A., Pucha, R., McManus, S., Nikutta, R., and Olsen, K. SPARCL: SPectra Analysis and Retrievable Catalog Lab. *arXiv e-prints*, art. arXiv:2401.05576, January 2024. doi: 10.48550/arXiv.2401.05576.
- Koleva, M., Prugniel, P., Bouchard, A., and Wu, Y. ULySS: a full spectrum fitting package. *A&A*, 501(3):1269–1279, July 2009. doi: 10.1051/0004-6361/200811467.
- Koposov, S. E., Allende Prieto, C., Cooper, A. P., Li, T. S., Bernaldo de Silva, L., Kim, B., Carrillo, A., Dey, A., Manser, C. J., Nikakhtar, F., Riley, A. H., Rockosi, C., Valluri, M., Aguilar, J., Ahlen, S., Bailey, S., Blum, R., Brooks, D., Claybaugh, T., Cole, S., de la Macorra, A., Dey, B., Forero-Romero, J. E., Gaztañaga, E., Guy, J., Kremin, A., Le Guillou, L., Levi, M. E., Manera, M., Meisner, A., Miquel, R., Moustakas, J., Nie, J., Palanque-Delabrouille, N., Percival, W. J., Rezaie, M., Rossi, G., Sanchez, E., Schlafly, E. F., Schubnell, M., Tarlé, G., Weaver, B. A., and Zhou, Z. DESI Early Data Release Milky Way Survey value-added catalogue. *MNRAS*, 533(1):1012–1031, September 2024. doi: 10.1093/mnras/stae1842.
- Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., and Zasowski, G. The Cannon: A data-driven approach to Stellar Label Determination. *ApJ*, 808(1):16, July 2015. doi: 10.1088/0004-637X/808/1/16.
- Nikutta, R., Fitzpatrick, M., Scott, A., and Weaver, B. Data lab—a community science platform. *Astronomy and Computing*, 33:100411, 2020. ISSN 2213-1337. doi: <https://doi.org/10.1016/j.ascom.2020.100411>.

URL <https://www.sciencedirect.com/science/article/pii/S2213133720300652>.

Parker, L., Lanusse, F., Golkar, S., Sarra, L., Cranmer, M., Bietti, A., Eickenberg, M., Krawezik, G., McCabe, M., Morel, R., Ohana, R., Pettee, M., Régaldo-Saint Blancard, B., Cho, K., Ho, S., and Polymathic AI Collaboration. AstroCLIP: a cross-modal foundation model for galaxies. *MNRAS*, 531(4):4990–5011, July 2024. doi: 10.1093/mnras/stae1450.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning Transferable Visual Models From Natural Language Supervision. *arXiv e-prints*, art. arXiv:2103.00020, February 2021. doi: 10.48550/arXiv.2103.00020.

Różański, T., Ting, Y.-S., and Jabłońska, M. Transformer-Payne: Enhancing Spectral Emulation Accuracy and Data Efficiency by Capturing Long-range Correlations. *ApJ*, 980(1):66, February 2025. doi: 10.3847/1538-4357/ad9b99.

Ting, Y.-S., Conroy, C., Rix, H.-W., and Cargile, P. The Payne: Self-consistent ab initio Fitting of Stellar Spectra. *ApJ*, 879(2):69, July 2019. doi: 10.3847/1538-4357/ab2331.

Xiang, M., Ting, Y.-S., Rix, H.-W., Sandford, N., Buder, S., Lind, K., Liu, X.-W., Shi, J.-R., and Zhang, H.-W. Abundance Estimates for 16 Elements in 6 Million Stars from LAMOST DR5 Low-Resolution Spectra. *ApJS*, 245(2):34, December 2019. doi: 10.3847/1538-4365/ab5364.

Xiang, M. S., Liu, X. W., Yuan, H. B., Huang, Y., Huo, Z. Y., Zhang, H. W., Chen, B. Q., Zhang, H. H., Sun, N. C., Wang, C., Zhao, Y. H., Shi, J. R., Luo, A. L., Li, G. P., Wu, Y., Bai, Z. R., Zhang, Y., Hou, Y. H., Yuan, H. L., Li, G. W., and Wei, Z. The LAMOST stellar parameter pipeline at Peking University - LSP3. *MNRAS*, 448(1):822–854, March 2015. doi: 10.1093/mnras/stu2692.

Zhang, M., Xiang, M., Ting, Y.-S., Wang, J., Li, H., Zou, H., Nie, J., Mou, L., Wu, T., Wu, Y., and Liu, J. Determining Stellar Elemental Abundances from DESI Spectra with the Data-driven Payne. *ApJS*, 273(2):19, August 2024. doi: 10.3847/1538-4365/ad51dd.

Zhao, X., Huang, Y., Xue, G., Kong, X., Liu, J., Tang, X., Beers, T. C., Ting, Y.-S., and Luo, A.-L. SpecCLIP: Aligning and Translating Spectroscopic Measurements for Stars. *arXiv e-prints*, art. arXiv:2507.01939, July 2025. doi: 10.48550/arXiv.2507.01939.