Differentiable Stellar Atmospheres with Physics-Informed Neural Networks

Jiadong Li*[1]   Mingjie Jian*[2]   Yuan-Sen Ting*[3][4]   Gregory M. Green[1]

Abstract

We present *Kurucz-a1*, a physics-informed neural network (PINN) that emulates 1D stellar atmosphere models under Local Thermodynamic Equilibrium (LTE), addressing a critical bottleneck in differentiable stellar spectroscopy. By incorporating hydrostatic equilibrium as a physical constraint during training, *Kurucz-a1* creates a differentiable atmospheric structure solver that maintains physical consistency while achieving computational efficiency. *Kurucz-a1* can achieve superior hydrostatic equilibrium and more consistent with the solar observed spectra compared to ATLAS-12 itself, demonstrating the advantages of modern optimization techniques. Combined with modern differentiable radiative transfer codes, this approach enables data-driven optimization of universal physical parameters across diverse stellar populations—a capability essential for next-generation stellar astrophysics.

1. Introduction

Understanding stellar properties requires detailed modeling of observed spectra using stellar atmosphere models. Since stellar interiors are opaque to radiation, observable absorption features originate primarily from the photosphere (Gray, 2008), the layer where light begins to escape from the atmosphere.

Stellar spectral modeling typically involves two central steps: constructing an atmospheric model and generating synthetic spectra. The first step determines the atmospheric structure—temperature, pressure, and electron density as functions of optical depth (measured in Rosseland mean opacity units)—by iteratively solving radiative transfer, radiative and hydrostatic equilibrium equations under assumptions of one-dimensional stratification and local thermodynamic equilibrium (LTE) (Hubeny & Mihalas, 2015). Traditional approaches rely on pre-computed grids such as ATLAS (Castelli & Kurucz, 2003), MARCS (Gustafsson et al., 2008), and PHOENIX (Allard, 2016). The second step uses this fixed atmospheric structure for spectral synthesis (e.g. Kurucz & Avrett, 1981; Sneden et al., 2012; Gerber et al., 2023; Piskunov & Valenti, 2017), performing radiative transfer calculations to generate synthetic spectra.

However, the output spectra are influenced by numerous poorly constrained parameters. Atomic transition parameters, such as oscillator strengths for individual lines, often require laboratory measurements or theoretical calculations that may be inaccurate (Asplund, 2005; Laverick et al., 2018). Furthermore, atmospheric modeling relies on opacity calculations (Magic et al., 2013; Colgan et al., 2016), convection treatments (Asplund et al., 2000a;b; Neilson & Lester, 2013; Pasetto et al., 2014), and assumptions about atmospheric structure (Asplund et al., 2005), which introduce systematic uncertainties into spectral modeling.

Large-scale spectroscopic surveys such as the Sloan Digital Sky Survey (York, 2000) and LAMOST (Zhao et al., 2012) provide extensive datasets that could enable optimization of these model parameters. While individual stars have different fundamental properties (effective temperature, surface gravity, metallicity), the underlying atomic physics remains universal. In principle, a differentiable end-to-end modeling framework could optimize these universal parameters by fitting large stellar samples while marginalizing over star-by-star properties—effectively using the diversity of stellar observations to constrain the physics that governs all stars.

Ideally, creating differentiable stellar spectroscopy codes would involve rewriting existing models in modern frameworks like `PyTorch` or `JAX`. However, most stellar atmosphere codes predate the widespread adoption of automatic differentiation. Many of these codes (e.g. Castelli & Kurucz, 2003) are written in legacy languages, such as `Fortran 77`, and have limited documentation. This presents a barrier to modernization, even with current automated coding tools.

Fortunately, the two modeling components—atmospheric structure model and radiative transfer—can be addressed

---

*Equal contribution [1]Max-Planck-Institut für Astronomie, Heidelberg, Germany [2]Department of Astronomy, Stockholm University, Sweden [3]Department of Astronomy, The Ohio State University, USA [4]Center for Cosmology and AstroParticle Physics, The Ohio State University, USA. Correspondence to: Jiadong Li <jdli@mpia.de>.
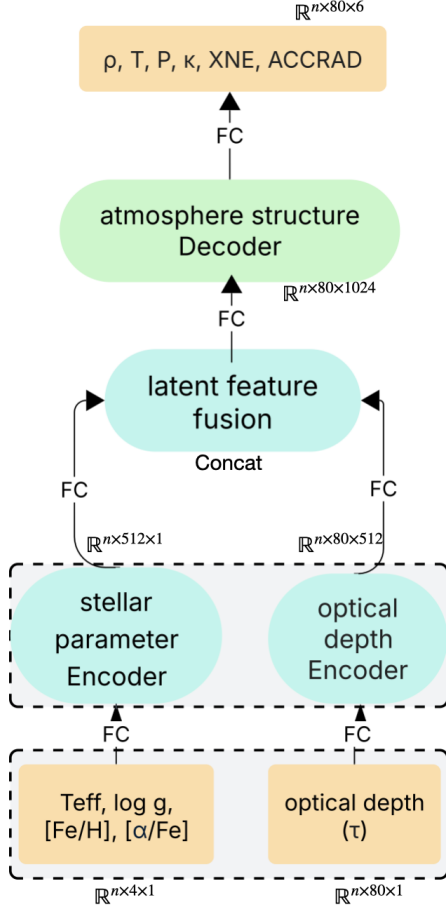
Figure 1. *Kurucz-a1* neural network architecture with dual-encoder design. Stellar parameters (T_{eff}, $\log g$, [Fe/H], [α/Fe]) and optical depth points are encoded separately, then concatenated and processed through fully connected layers (FC) to predict atmospheric parameters (ρ_x, T, P, κ, X_{NE}, ACCRAD) at each depth point.

separately. Modern frameworks such as Korg (Wheeler et al., 2024) have successfully modernized radiative transfer solvers in differentiable languages. However, the atmospheric structure solver remains the key bottleneck for end-to-end differentiable modeling.

In this study, we address this bottleneck by developing a physics-informed neural network (PINN) emulator for stellar atmospheric structure calculations. While this approach does not directly solve the differential equations in a differentiable framework, it creates a differentiable component that, when combined with modern radiative transfer solvers, enables the end-to-end optimization framework described above.

2. Motivation for PINN

Emulating stellar atmospheric structure presents challenges even in discretized form. Consider the ATLAS codes (Ku-

rucz, 1970; Kurucz & Avrett, 1981; Kurucz, 1996; 2014) —cornerstone tools in stellar astrophysics developed by Robert Kurucz[1]. The ATLAS-12 atmospheric structure solver produces high-dimensional vector fields as functions of Rosseland optical depth: six atmospheric parameters (temperature, pressure, density, electron number density, opacity, and radiative acceleration) across 80 depth points, yielding 480-dimensional outputs.

The input space is equally complex. While fundamental stellar parameters (T_{eff}, $\log g$, [Fe/H], [α/Fe]) provide the primary constraints, individual elemental abundances can affect atmospheric structure (LeBlanc et al., 2009; Ting et al., 2016). This creates a challenging high-dimensional input-output mapping problem where traditional neural network architectures lack appropriate inductive biases for the complex physical relationships governing stellar atmospheres.

However, stellar atmospheric structure is governed by well-established physical laws—primarily hydrostatic equilibrium and radiative transfer equations. These constraints provide the essential inductive bias that standard architectures cannot capture. Physics-Informed Neural Networks (PINNs) (E & Yu, 2017; Raissi et al., 2018; 2019) offer an optimal solution by integrating these physical constraints directly into the learning process through composite loss functions that enforce both data fidelity and physics compliance.

3. Method

Kurucz-a1 employs a dual-encoder architecture (Figure 1) that separates global stellar parameters from local depth information. This design reflects the underlying physics where global stellar properties establish the overall atmospheric structure, while local conditions vary systematically with atmospheric depth.

The stellar parameter encoder processes four fundamental quantities—effective temperature (T_{eff}), surface gravity ($\log g$), metallicity ([Fe/H]), and α-enhancement ([α/Fe])—through a multi-layer perceptron to produce 512-dimensional embeddings. The depth encoder transforms the 80 Rosseland optical depth points (τ_{Ross}) into 512-dimensional representations at each atmospheric layer, allowing the network to learn appropriate representations for the atmospheric depth coordinate.

The stellar parameter embedding is broadcast and concatenated with each depth embedding to form 1024-dimensional combined embeddings at all 80 depth points. These are

---

[1]Robert Kurucz, who dedicated his life to developing these transformative codes, passed away in 2025. Together with Fiorella Castelli, they built the foundation on which much of contemporary stellar spectroscopy still relies. This work aims to preserve and further modernize their shared legacy.
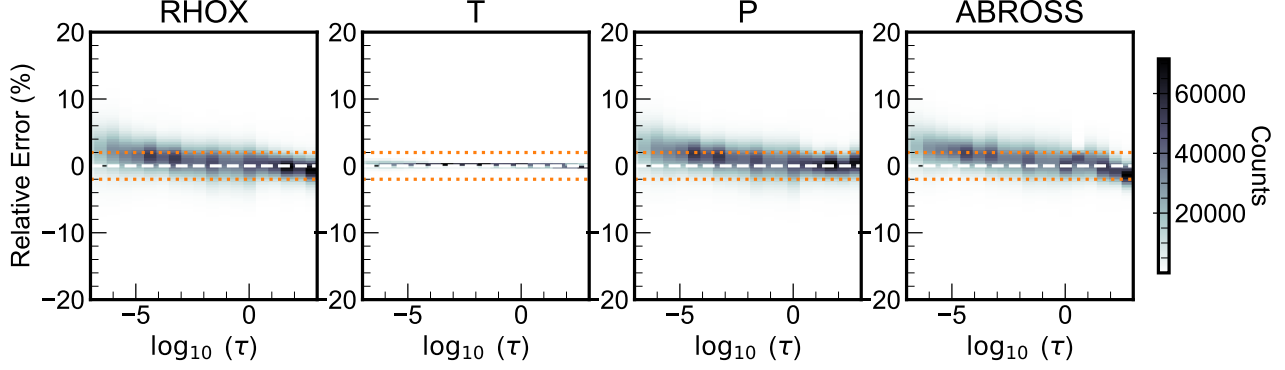
Figure 2. Relative error distributions for atmospheric parameters predicted by *Kurucz-a1* compared to ATLAS-12 models. The four panels show column mass density (RHOX), temperature (T), pressure (P), and Rosseland opacity (ABROSS) versus optical depth. The validation dataset spans Galactic stellar populations (see appendix).
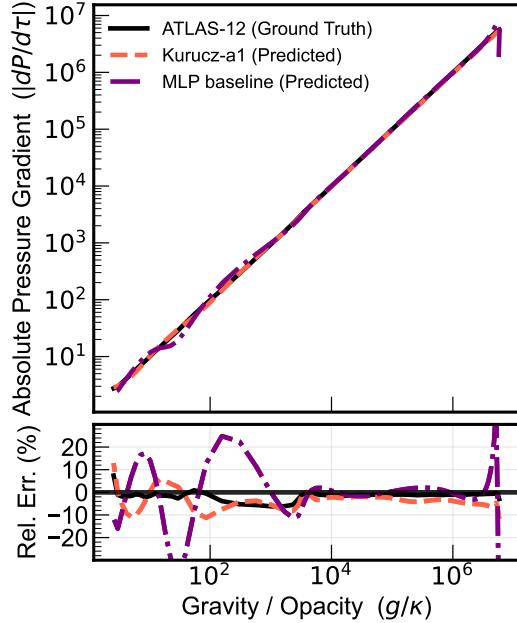

Figure 3. Hydrostatic equilibrium diagnostic for a representative red-giant branch star ($T_{\rm eff} = 4500\,{\rm K}$, $\log g = 2.5$, [Fe/H] = 0, [α/Fe] = 0). Deviations from the diagonal indicate violations of $dP/d\tau = g/\kappa$. *Kurucz-a1* (dashed) demonstrates superior agreement compared to the MLP baseline (dot-dashed), nearly matching ATLAS-12 (solid).

Figure 4. Distribution of hydrostatic equilibrium loss values across the validation dataset. *Kurucz-a1* (purple) achieves a tight distribution comparable to ATLAS-12 (gray), while the MLP baseline (red) shows higher and more scattered values. Values near zero indicate perfect adherence to the hydrostatic equilibrium constraint.

processed through a 3-layer MLP emulator (with hidden dimensions [1024, 512, 256]) to predict six atmospheric parameters at each depth: column mass density (ρ_x), temperature (T), gas pressure (P), electron number density ($X_{\rm NE}$), Rosseland mean opacity ($\kappa_{\rm Ross}$), and radiative acceleration (ACCRAD). We use GeLU activations (Hendrycks & Gimpel, 2016).
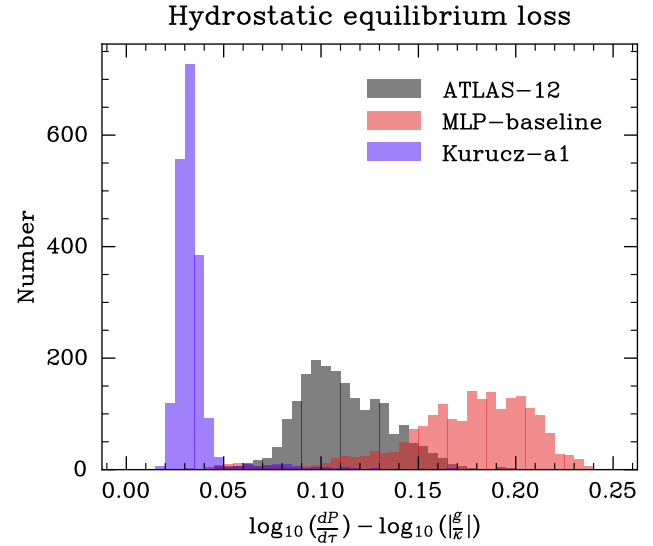
The total loss function combines data reproduction and physics enforcement:

$$\mathcal{L}_{\rm total} = (1 - \alpha) \cdot \mathcal{L}_{\rm data} + \alpha \cdot \mathcal{L}_{\rm physics} \qquad (1)$$

with $\alpha = 0.03$ by grid search, where $\mathcal{L}_{\rm data}$ is the MSE loss of the predicted outputs. The physics loss enforces hydrostatic equilibrium:

$$\mathcal{L}_{\rm physics} = \frac{1}{N_b \cdot N_d} \sum_{i,j} \left(\frac{dP}{d\tau} - \frac{g}{\kappa} \right)^2_{i,j} \qquad (2)$$

where $\frac{dP}{d\tau}$ is computed via automatic differentiation, enforcing the constraint that governs atmospheric stability:

$$\frac{dP}{d\tau} = \frac{g}{\kappa} \qquad (3)$$

4. Results

We evaluate *Kurucz-a1*'s performance to demonstrate that physics-informed constraints are essential for robust stellar atmosphere emulation. The validation encompasses 10,427 stellar models spanning $T_{\text{eff}} = 2500$–$50,000$ K, $\log g = -1$–5.5, [Fe/H] $= -4$–$+1.46$, and [α/Fe] $= -0.2$–$+0.62$ (see Appendix for training data details).

Figure 2 shows the relative error distributions across optical depth for all atmospheric parameters. *Kurucz-a1* achieves median errors below 0.12% for temperature, 1.1% for pressure and density, and 1.5% for opacity across the full range of stellar parameters tested. These errors are well within acceptable bounds for stellar atmosphere modeling. For context, even state-of-the-art grids like MARCS and ATLAS exhibit temperature differences of 10 K to 80 K in their upper layers ($\tau_{\text{Ross}} \leq 10^{-2}$), which is considered the current frontier of uncertainty stemming from different input physics (Gustafsson et al., 2008). Comparisons for the APOGEE survey show the MARCS and ATLAS9 grids agree to within $< 1\%$ for temperature (Heiter et al., 2015) and 2-3% for pressure (Mészáros et al., 2012). Since our model's median temperature error is well below these accepted inter-model discrepancies, it demonstrates a high degree of fidelity.

However, accurate stellar atmosphere modeling requires more than point-wise precision—the physical consistency of atmospheric profiles is crucial (Magic et al., 2013). Radiative transfer calculations depend critically on pressure and temperature gradients, not just their absolute values (Hayek et al., 2010).

Figure 3 demonstrates the advantage of *Kurucz-a1* in maintaining hydrostatic equilibrium. *Kurucz-a1* achieves hydrostatic equilibrium comparable to ATLAS-12 itself, while outperforming the MLP baseline that lacks physics constraints. This validates the vital role of physics-informed loss terms.

Notably, *Kurucz-a1*, on average, achieves even better hydrostatic equilibrium than ATLAS-12, likely due to ATLAS-12's discrete numerical scheme and optimization methods that predate modern gradient-based techniques. The small deviations from precise equilibrium in ATLAS-12 ($\sim$0.1% level) reflect limitations of finite-difference discretization.

Finally, we validated *Kurucz-a1*'s practical utility by synthesizing solar spectra using PySME (Wehrhahn et al., 2023; Piskunov & Valenti, 2017) and comparing with observations (Figure 5). The synthetic spectra generated using *Kurucz-a1* closely match both ATLAS-12 models and observations from the Melchior database (Royer, 2024), confirming that the emulated atmospheric output is indistinguishable from reference models. *Kurucz-a1* even outperforms ATLAS-12 in some spectral line wings, likely due to better hydrostatic equilibrium adherence.

5. Broader Impact

Kurucz-a1, available open-source at `https://github.com/jiadonglee/kurucz-a1`, provides a practical tool and methodological template for the astronomical community, achieving a mean inference time of 0.37 ± 0.06 millisecond per model on an Apple M1 Pro chip.

For stars hotter than $\sim$4500K, where molecular transitions are minimal, radiative transfer codes typically execute within seconds to minutes. Consequently, stellar spectral analysis commonly employs sparse grids of pre-computed atmospheric models, running only spectral synthesis code like `Korg` (Wheeler et al., 2024) and `PySME` (Piskunov & Valenti, 2017; Wehrhahn et al., 2023) to synthesize lines for comparison with observations.

However, as demonstrated by Ting et al. (2016), these sparse atmospheric grids introduce systematic errors in stellar line analysis. While solving stellar atmospheric structure with codes like ATLAS requires computational overhead (minutes to hours per model), current sparse grid approaches also limit analysis to few parameters, grouping elemental abundances into broad categories like [Fe/H] and [α/Fe].

Kurucz-a1 provides a low-cost replacement for stellar atmosphere synthesis models. It generates atmospheric output within seconds, making it practical to eliminate the interpolation errors inherent in sparse grids while maintaining computational efficiency. This framework also eases data storage and transfer requirements when additional abundance dimensions are included.

The current implementation uses four stellar parameters consistent with existing atmospheric grids. Future work will extend the PINN framework to higher-dimensional parameter spaces, enabling incorporation of additional abundance dimensions (alpha-elements, CNO, helium) essential for modeling diverse stellar populations, as well as incorporating additional physics such as equation of state and energy balance.

Beyond practical applications, this work addresses a fundamental bottleneck in achieving end-to-end differentiable stellar spectroscopy. Combined with modern differentiable radiative transfer codes, this approach opens pathways to data-driven calibration of atomic physics parameters that current stellar models must assume.
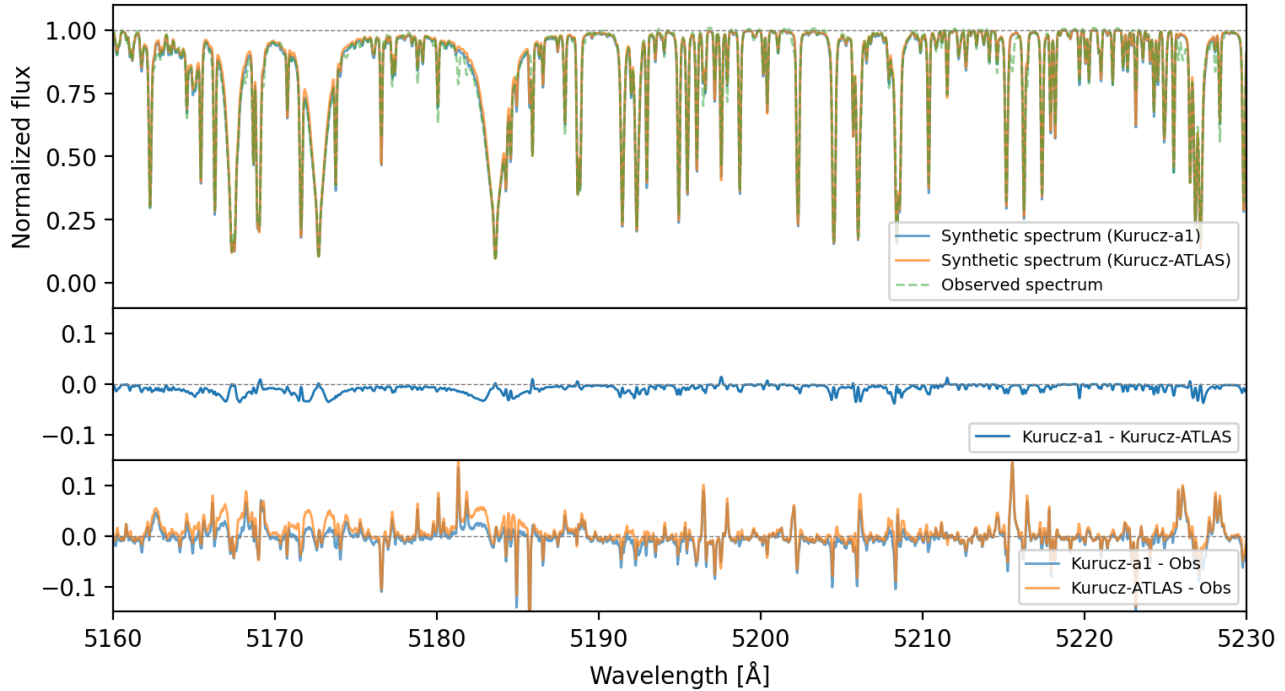
4

Figure 5. Solar spectrum synthesis validation. **Top**: Observed solar spectrum (green dashed) compared to synthetic spectra from *Kurucz-a1* (blue) and ATLAS-12 (orange). **Middle**: Direct comparison between *Kurucz-a1* and ATLAS-12 synthetic spectra, showing minimal differences. **Bottom**: Residuals relative to observed spectrum demonstrate that *Kurucz-a1* achieves better agreement than ATLAS-12 due to superior hydrostatic equilibrium adherence.

6. Acknowledgements

References

Allard, F. The PHOENIX Model Atmosphere Grid for Stars. In Reylé, C., Richard, J., Cambrésy, L., Deleuil, M., Pécontal, E., Tresse, L., and Vauglin, I. (eds.), *SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, pp. 223–227, December 2016.

Asplund, M. Uncertainties in stellar abundance analyses. *Highlights of Astronomy*, 13:542–547, 2005. doi: 10.1017/S1539299600016555.

Asplund, M., Nordlund, Å., Trampedach, R., Allende Prieto, C., and Stein, R. F. Line formation in solar granulation. I. Fe line shapes, shifts and asymmetries. *A&A*, 359:729–742, July 2000a. doi: 10.48550/arXiv.astro-ph/0005320.

Asplund, M., Nordlund, Å., Trampedach, R., and Stein, R. F. Line formation in solar granulation. II. The photospheric Fe abundance. *A&A*, 359:743–754, July 2000b. doi: 10.48550/arXiv.astro-ph/0005321.

Asplund, M., Grevesse, N., and Sauval, A. J. The Solar Chemical Composition. In Barnes, III, T. G. and Bash, F. N. (eds.), *Cosmic Abundances as Records of Stellar Evolution and Nucleosynthesis*, volume 336 of *Astronomical Society of the Pacific Conference Series*, pp. 25, September 2005.

Bressan, A., Marigo, P., Girardi, L., Salasnich, B., Dal Cero, C., Rubele, S., and Nanni, A. PARSEC: stellar tracks and isochrones with the PAdova and TRieste Stellar Evolution Code. *MNRAS*, 427(1):127–145, November 2012. doi: 10.1111/j.1365-2966.2012.21948.x.

Castelli, F. and Kurucz, R. L. New Grids of ATLAS9 Model Atmospheres. In Piskunov, N., Weiss, W. W., and Gray, D. F. (eds.), *Modelling of Stellar Atmospheres*, volume 210 of *IAU Symposium*, pp. A20, January 2003. doi: 10.48550/arXiv.astro-ph/0405087.

Colgan, J., Kilcrease, D. P., Magee, N. H., Sherrill, M. E., Abdallah, Jr., J., Hakel, P., Fontes, C. J., Guzik, J. A., and Mussack, K. A. A New Generation of Los Alamos

Opacity Tables. *ApJ*, 817(2):116, February 2016. doi: 10.3847/0004-637X/817/2/116.

E, W. and Yu, B. The Deep Ritz method: A deep learning-based numerical algorithm for solving variational problems. *arXiv e-prints*, art. arXiv:1710.00211, September 2017. doi: 10.48550/arXiv.1710.00211.

Gerber, J. M., Magg, E., Plez, B., Bergemann, M., Heiter, U., Olander, T., and Hoppe, R. Non-LTE radiative transfer with Turbospectrum. *A&A*, 669:A43, January 2023. doi: 10.1051/0004-6361/202243673.

Gray, D. F. *The Observation and Analysis of Stellar Photospheres*. 2008.

Gustafsson, B., Edvardsson, B., Eriksson, K., Jørgensen, U. G., Nordlund, Å., and Plez, B. A grid of MARCS model atmospheres for late-type stars. I. Methods and general properties. *A&A*, 486(3):951–970, August 2008. doi: 10.1051/0004-6361:200809724.

Hayek, W., Asplund, M., Carlsson, M., Trampedach, R., Collet, R., Gudiksen, B. V., Hansteen, V. H., and Leenaarts, J. Radiative transfer with scattering for domain-decomposed 3D MHD simulations of cool stellar atmospheres. Numerical methods and application to the quiet, non-magnetic, surface of a solar-type star. *A&A*, 517:A49, July 2010. doi: 10.1051/0004-6361/201014210.

Heiter, U., Jofré, P., Gustafsson, B., Korn, A. J., Soubiran, C., and Thévenin, F. Gaia FGK benchmark stars: Effective temperatures and surface gravities. *A&A*, 582:A49, October 2015. doi: 10.1051/0004-6361/201526319.

Hendrycks, D. and Gimpel, K. Gaussian Error Linear Units (GELUs). *arXiv e-prints*, art. arXiv:1606.08415, June 2016. doi: 10.48550/arXiv.1606.08415.

Hubeny, I. and Mihalas, D. *Theory of Stellar Atmospheres. An Introduction to Astrophysical Non-equilibrium Quantitative Spectroscopic Analysis*. 2015.

Kurucz, R. L. Atlas: a Computer Program for Calculating Model Stellar Atmospheres. *SAO Special Report*, 309, February 1970.

Kurucz, R. L. Status of the ATLAS 12 Opacity Sampling Program and of New Programs for Rosseland and for Distribution Function Opacity. In Adelman, S. J., Kupka, F., and Weiss, W. W. (eds.), *M.A.S.S., Model Atmospheres and Spectrum Synthesis*, volume 108 of *Astronomical Society of the Pacific Conference Series*, pp. 160, January 1996.

Kurucz, R. L. ATLAS12: Opacity sampling model atmosphere program. Astrophysics Source Code Library, record ascl:1303.024, March 2013.

Kurucz, R. L. Model Atmosphere Codes: ATLAS12 and ATLAS9. In Niemczura, E., Smalley, B., and Pych, W. (eds.), *Determination of Atmospheric Parameters of B*, pp. 39–51. 2014. doi: 10.1007/978-3-319-06956-2_4.

Kurucz, R. L. and Avrett, E. H. Solar Spectrum Synthesis. I. A Sample Atlas from 224 to 300 nm. *SAO Special Report*, 391, May 1981.

Laverick, M., Lobel, A., Merle, T., Royer, P., Martayan, C., David, M., Hensberge, H., and Thienpont, E. The Belgian repository of fundamental atomic data and stellar spectra (BRASS). I. Cross-matching atomic databases of astrophysical interest. *A&A*, 612:A60, April 2018. doi: 10.1051/0004-6361/201731933.

LeBlanc, F., Monin, D., Hui-Bon-Hoa, A., and Hauschildt, P. H. Stellar model atmospheres with abundance stratification. *A&A*, 495(3):937–944, March 2009. doi: 10.1051/0004-6361:200810848.

Magic, Z., Collet, R., Asplund, M., Trampedach, R., Hayek, W., Chiavassa, A., Stein, R. F., and Nordlund, Å. The Stagger-grid: A grid of 3D stellar atmosphere models. I. Methods and general properties. *A&A*, 557:A26, September 2013. doi: 10.1051/0004-6361/201321274.

Mészáros, S., Allende Prieto, C., Edvardsson, B., Castelli, F., García Pérez, A. E., Gustafsson, B., Majewski, S. R., Plez, B., Schiavon, R., Shetrone, M., and de Vicente, A. New ATLAS9 and MARCS Model Atmosphere Grids for the Apache Point Observatory Galactic Evolution Experiment (APOGEE). *AJ*, 144(4):120, October 2012. doi: 10.1088/0004-6256/144/4/120.

Neilson, H. R. and Lester, J. B. Spherically symmetric model stellar atmospheres and limb darkening. II. Limb-darkening laws, gravity-darkening coefficients and angular diameter corrections for FGK dwarf stars. *A&A*, 556:A86, August 2013. doi: 10.1051/0004-6361/201321888.

Pasetto, S., Chiosi, C., Cropper, M., and Grebel, E. K. Theory of stellar convection: removing the mixing-length parameter. *MNRAS*, 445(4):3592–3609, December 2014. doi: 10.1093/mnras/stu1933.

Piskunov, N. and Valenti, J. A. Spectroscopy Made Easy: Evolution. *A&A*, 597:A16, January 2017. doi: 10.1051/0004-6361/201629124.

Raissi, M., Yazdani, A., and Karniadakis, G. E. Hidden Fluid Mechanics: A Navier-Stokes Informed Deep Learning Framework for Assimilating Flow Visualization Data. *arXiv e-prints*, art. arXiv:1808.04327, August 2018. doi: 10.48550/arXiv.1808.04327.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019. doi: 10.1016/j.jcp.2018.10.045.

Royer, P., e. a. MELCHIORS. The Mercator Library of High Resolution Stellar Spectroscopy. *A&A*, 681:A107, January 2024. doi: 10.1051/0004-6361/202346847.

Sneden, C., Bean, J., Ivans, I., Lucatello, S., and Sobeck, J. MOOG: LTE line analysis and spectrum synthesis. Astrophysics Source Code Library, record ascl:1202.009, February 2012.

Ting, Y.-S., Conroy, C., and Rix, H.-W. Accelerated Fitting of Stellar Spectra. *ApJ*, 826(1):83, July 2016. doi: 10.3847/0004-637X/826/1/83.

Wehrhahn, A., Piskunov, N., and Ryabchikova, T. PySME. Spectroscopy Made Easier. *A&A*, 671:A171, March 2023. doi: 10.1051/0004-6361/202244482.

Wheeler, A. J., Casey, A. R., and Abruzzo, M. W. Korg: Fitting, Model Atmosphere Interpolation, and Brackett Lines. *AJ*, 167(2):83, February 2024. doi: 10.3847/1538-3881/ad19cc.

York, D. G., e. a. The Sloan Digital Sky Survey: Technical Summary. *AJ*, 120:1579–1587, September 2000. doi: 10.1086/301513. URL http://adsabs.harvard.edu/abs/2000AJ....120.1579Y.

Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., and Deng, L.-C. LAMOST spectral survey — An overview. *Research in Astronomy and Astrophysics*, 12(7):723–734, July 2012. doi: 10.1088/1674-4527/12/7/002.
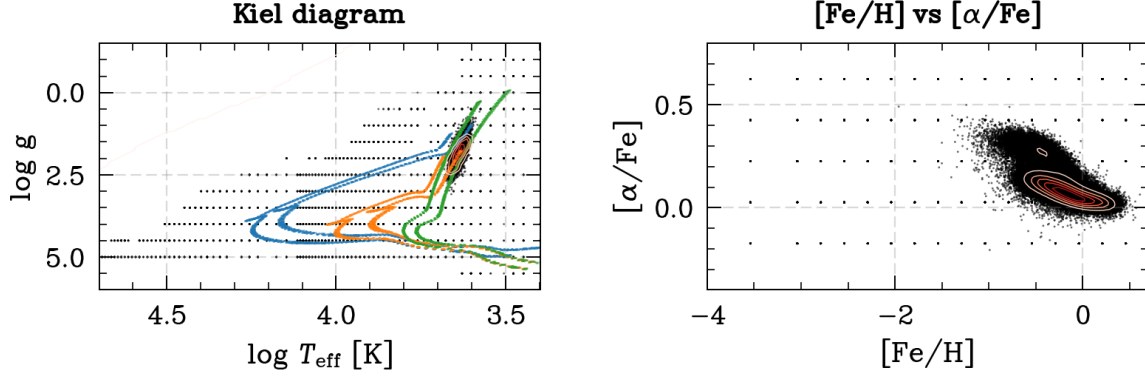
Figure 6. Distribution of stellar atmospheric models in the training dataset. *Left panel:* Kiel diagram showing coverage in effective temperature ($\log T_{\text{eff}}$) versus surface gravity ($\log g$) parameter space. The colored evolutionary tracks from PARSEC models (Bressan et al., 2012) represent different stellar ages (0.1, 1, 10 Gyr) and metallicities (-1 and 0). *Right panel:* Distribution in the [Fe/H]–[α/Fe] abundance plane, showing the density of atmospheric models with contour lines indicating regions of highest concentration.

A. Appendix: Training Data

Our model utilizes pre-computed stellar atmosphere grids from the ATLAS-12 code (Kurucz, 2013), as illustrated in Figure 6. The grid encompasses effective temperatures (T_{eff}) spanning 2500–50,000 K, surface gravities ($\log g$) from -1 to 5.5, metallicities ([Fe/H]) from -4 to $+1.46$, and α-element abundances ([α/Fe]) from -0.2 to $+0.62$.

The dataset comprises 104,269 atmospheric models, each providing six fundamental parameters as functions of 80 optical depth points: column mass density (ρ_x), temperature (T), gas pressure (P), electron number density (X_{NE}), Rosseland mean opacity (κ_{Ross}), and radiative acceleration (ACCRAD).

To optimize sampling for Galactic stellar studies, we implemented enhanced density for red giant branch (RGB) stars in both the Kiel diagram and [Fe/H]–[α/Fe] plane, given their importance for understanding Galactic dynamics and chemical evolution. The enhanced sampling around metallicity ([Fe/H] ~ -1 to $+0.5$) and α-abundance patterns ([α/Fe] ~ 0.0–0.4) reflects the focus on Galactic stellar populations.

The dataset was randomly partitioned into training (90%, 93,842 models) and validation (10%, 10,427 models) subsets for model development and performance evaluation.