









A method to derive self-consistent NLTE astrophysical parameters for four million high-resolution 4MOST stellar spectra in half a day with invertible neural networks

Victor F. Ksoll^{1,*}, Nicholas Storm^{2,3}, Maria Bergemann², Katherine Lee², Ralf S. Klessen^{1,4},
R. Albarracín^{2,3}, Guillaume Guiglion^{5,2,6}, and Gražina Tautvaišienė⁷

- ¹ Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Überle-Str. 2, 69120 Heidelberg, Germany
² Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
³ Universität Heidelberg, Grabengasse 1, 69117 Heidelberg, Germany
⁴ Universität Heidelberg, Interdisziplinäres Zentrum für Wissenschaftliches Rechnen, Im Neuenheimer Feld 225, 69120 Heidelberg, Germany
⁵ Zentrum für Astronomie der Universität Heidelberg, Landessternwarte, Königstuhl 12, 69117 Heidelberg, Germany
⁶ Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany
⁷ Vilnius University, Faculty of Physics, Institute of Theoretical Physics and Astronomy, Sauletekio av. 3, 10257 Vilnius, Lithuania

Received 16 December 2025 / Accepted 19 February 2026

ABSTRACT

Context. Modern spectroscopic surveys have the capacity to obtain the spectra of millions of stars. However, classical spectroscopic methods can often be computationally expensive, rendering them impractical for the analysis of large datasets.

Aims. We introduce a novel simulation-based, deep-learning approach for the efficient analysis of high-resolution stellar spectra that will be obtained with the upcoming high-resolution 4MOST spectrograph.

Methods. We used a suite of synthetic non-local thermodynamic equilibrium (NLTE) spectra generated with Turbospectrum to mimic 4MOST observations and trained a conditional invertible neural network (cINN) for the purpose of predicting self-consistently stellar surface parameters and chemical abundances. The cINN is a neural network architecture that estimates full posterior distributions for the target stellar properties, providing an intrinsic uncertainty estimate. We evaluated the predictive performance of the trained cINN model on both synthetic data and the observed spectra of stars.

Results. We found that our new cINN trained on NLTE synthetic spectra is capable of recovering stellar parameters with average errors (σ) of 33 K for T_{eff} , 0.16 dex for $\log(g)$, and 0.12 dex for [Fe/H], 0.11 dex for [Ca/Fe], 0.11 dex for [Mg/Fe], and 0.51 dex for [Li/Fe], respectively, at a signal-to-noise ratio (S/N) of 250 per Angstrom. From the analysis of the observed spectra of *Gaia*-ESO / 4MOST / PLATO benchmark stars, we verified that our NLTE estimates for stellar parameters and abundances are consistent with results obtained with the independent code TSFitPy. We conclude that the NLTE cINN is robust and that it can, in theory, evaluate four million high-resolution 4MOST spectra in less than a day, using GPU acceleration.

Key words. methods: statistical – techniques: spectroscopic – stars: abundances – stars: atmospheres

1. Introduction

The past decade has seen revolutionary advances in observational stellar and Galactic astronomy. Large-scale stellar spectroscopic surveys, such as *Gaia* (Gaia Collaboration 2016) and *Gaia*-ESO (Gilmore et al. 2012, 2022; Randich et al. 2013, 2022), GALAH (De Silva et al. 2015), LAMOST (Zhao et al. 2012), and APOGEE (Majewski et al. 2016, 2017) have provided millions of stellar spectra, and upcoming facilities such as WEAVE (Jin et al. 2024), 4MOST (4-metre Multi-Object Spectroscopic Telescope; de Jong et al. 2019), and SDSS-V (Kollmeier et al. 2026) will further increase the data quantity by several orders of magnitude.

To make use of this enormous wealth in observational data, all of these spectroscopic datasets require efficient analysis approaches. Classic fitting tools such as MOOG

(Snedden 1973), SIU (Reetz 1991), PySME (Wehrhahn et al. 2023), KORG (Wheeler et al. 2023), and TSFitPy (Gerber et al. 2023; Storm & Bergemann 2023) are highly flexible in terms of micro-physics. However, these codes typically require hours of CPU time in order to provide a detailed analysis of chemical abundances and stellar parameters for a single star. Although this approach is highly desirable and justified for individual stars, it becomes prohibitively expensive in terms of computation time when homogeneous, precise, and fast analyses of millions of high-resolution stellar spectra are required.

Consequently, approaches that employ machine learning have been gaining considerable traction in recent years as a more efficient alternative. In particular, neural network-based methods have proven to be promising. Ting et al. (2019) introduced a spectral fitting tool known as ‘the Payne’, which uses a simple feed-forward neural network to facilitate a precise interpolation of synthetic spectra in a very high-dimensional parameter space. It was applied, for instance, in the analysis of

* Corresponding author: v.ksoll@uni-heidelberg.de

GALAH DR4 (Buder et al. 2025). Another data-driven spectral analysis approach is known as ‘the Cannon’ (Ness et al. 2015), which uses real, well-analysed spectra to build an efficient generative model for the probability distribution of the flux per wavelength given the stellar parameters. It was adapted for APOGEE spectra in Ness et al. (2016). The application of convolutional neural networks (CNNs) has also been explored in, for instance, the analysis of RAVE spectra (Guiglion et al. 2020) and *Gaia* RVS spectra (Guiglion et al. 2024). The analysis of spectra has also seen exploration of various generative deep learning approaches. Liu et al. (2024) introduced the StellarGAN, a spectral classification approach for O to M-type stars based on generative adversarial networks (GANs; Goodfellow et al. 2014) that is trained on SDSS and APOGEE data. Another GAN-based approach is the GANDALF framework (Santoveña et al. 2024; Manteiga et al. 2025), which is trained on APOGEE and GALAH spectra and employs an autoencoder to embed spectra in a lower dimensional representation space prior to deriving stellar atmospheric properties. More recently, Cvrček et al. (2025) presented a method that uses variational autoencoders (VAEs; Kingma & Welling 2013; Rezende et al. 2014) to both predict stellar parameters and simulate spectra, which they trained on both real and synthetic spectra from the HARPS (Mayor et al. 2003) instrument. Gebran & Bentley (2025) also employed conditional VAEs (cVAEs; Sohn et al. 2015) to build a physics-aware surrogate model for spectral synthesis, trained on synthetic spectra from SYNPEC (Hubeny & Lanz 2011). However, most of these approaches rely on stellar parameters and chemical abundances, obtained using highly simplified stellar synthetic spectra computed under assumptions of local thermodynamic equilibrium (LTE). As demonstrated in numerous recent studies (e.g. Korn et al. 2003; Asplund 2005; Caffau et al. 2010; Lind et al. 2011; Bergemann et al. 2012b,a; Asplund et al. 2021; Amarsi et al. 2020; Lind & Amarsi 2024; Bergemann & Hoppe 2025), derivation of precise and accurate stellar parameters and abundance requires that stellar spectroscopic model grids take into account the effects of non-LTE (NLTE).

Here, we present a novel approach to provide accurate astrophysical characterisation for gigantic sets of data; specifically, from the upcoming 4MOST high-resolution disc and bulge survey 4MIDABLE-HR (Bensby et al. 2019). In this survey (PI: Bensby T. and Bergemann M.), around four million spectra will be obtained at high resolution ($R = 20\,000$) to study the evolution of the Galactic disc and bulge. Dictated by the need to have highly efficient, stable, and adaptable algorithms, as well as the need to include NLTE in the chemical analysis (Lind & Amarsi 2024; Bergemann & Hoppe 2025), our aim is to expand upon the code presented in Kovalev et al. (2019) and applied to the *Gaia*-ESO survey therein. In particular, we introduced a novel simulation-based inference approach that for the first time combines NLTE synthetic spectra with a conditional invertible neural network (cINN; Ardizzone et al. 2019a,b). Overall, cINNs are a deep learning architecture that belongs to the generative normalising flow approaches and are specifically tailored towards solving inverse problems. They estimate the full posterior distribution, providing both an intrinsic uncertainty measure for the predictions and highlighting potential degeneracies in the given inverse problem. Although VAEs and GANs can be employed in a similar fashion (as in e.g. Cvrček et al. 2025), INNs have been shown to return better calibrated posterior distributions (Ardizzone et al. 2019a) and to be more robust to mode collapse (Ardizzone et al. 2019b; Kobzyev et al. 2021; Bond-Taylor et al. 2022). In addition, cINNs have already seen a series of successful applications in astronomy. These include

the characterisation of stars from photometry (Ksoll et al. 2020), analysis of emission lines from HII regions (Kang et al. 2022, 2023a), inference of cosmic ray origins (Bister et al. 2022), evaluation of exoplanet observations (Haldemann et al. 2023), prediction of galaxy merger histories (Eisert et al. 2023), and 3D reconstruction of dust distributions in star-forming cores (Ksoll et al. 2024). In the role that is most relevant to the work presented in this study, cINNs have been proven to work well for the analysis of stellar spectra in several applications. This includes the evaluation of MUSE spectra based on LTE stellar atmosphere models from the PHOENIX library (Kang et al. 2023b, 2025) and the characterisation of high resolution spectra using a cINN trained on observed spectra of Milky Way field and cluster stars obtained with the GIRAFFE spectrograph in the *Gaia*-ESO survey (Candebat et al. 2024).

In the following, we outline our development of the first self-consistent cINN method for the analysis of observations of stars with the high-resolution spectrograph of 4MOST that is based fully on 25-D NLTE stellar spectral model grids. In Sect. 2, we introduce the training data for our cINN model, as well as a set of real benchmark spectra that we use to ascertain the performance of our model. Section 3 outlines the basic concepts of the cINN approach and implementation details for the specific inverse problem that we aim to solve in this study. In Sect. 4, we then describe the performance of our approach on both synthetic test data, as well as on archival spectra of the *Gaia*-ESO/4MOST/PLATO benchmark stars. Lastly, Sect. 5 provides a summary of our main results and outlook for the subsequent development of our method.

2. Data

In this section, we briefly outline the preparation of the training data (Sect. 2.1). We also describe the composition of the set of real benchmark spectra used in Sect. 4.3 to gauge the performance of our trained model on real data (Sect. 2.2).

2.1. Training data

We computed large grids of NLTE spectra (29 548 stellar spectral models) with our latest version of Turbospectrum NLTE (TSv20¹; Gerber et al. 2023; Storm & Bergemann 2023), using the TSPy Python wrapper². The generated spectra employed 1D MARCS model atmospheres (Gustafsson et al. 2008) and cover a wavelength range of 3670–9800 Å in steps of 0.007 Å with perfect resolution without any broadening, as well as varying stellar parameters (T_{eff} , $\log(g)$, [Fe/H]) and abundances for 21 elements (Li, C, N, O, Na, Mg, Al, Si, Ca, Ti, V, Cr, Mn, Co, Ni, Sr, Y, Zr, Ba, Ce, Eu). In our spectrum synthesis calculations, we adopted a random distribution of micro-turbulence parameter values v_{mic} , ranging from 0.5 to 2 km s⁻¹. We used the *Gaia*-ESO line lists (Heiter et al. 2021) with new atomic data for C, N, O, Si, Mg, as described in (Magg et al. 2022), and VALD for its gaps (Ryabchikova et al. 2015). The chemical composition of the Sun is adopted from our NLTE analysis in Magg et al. (2022). There were 16 elements simulated in NLTE: H (Mashonkina et al. 2008), O (Bergemann et al. 2021), Na (Ezzeddine et al. 2018), Mg (Bergemann et al. 2017), Al (Ezzeddine et al. 2018), Si (Bergemann et al. 2013; Magg et al. 2022), Ca (Mashonkina et al. 2017; Semenova et al. 2020), Ti (Bergemann 2011), Mn (Bergemann et al. 2019), Fe (Bergemann et al. 2012b; Semenova et al. 2020), Co (Bergemann et al. 2010;

¹ https://github.com/bertrandplez/Turbospectrum_NLTE

² <https://github.com/TSFitPy-developers/TSFitPy>

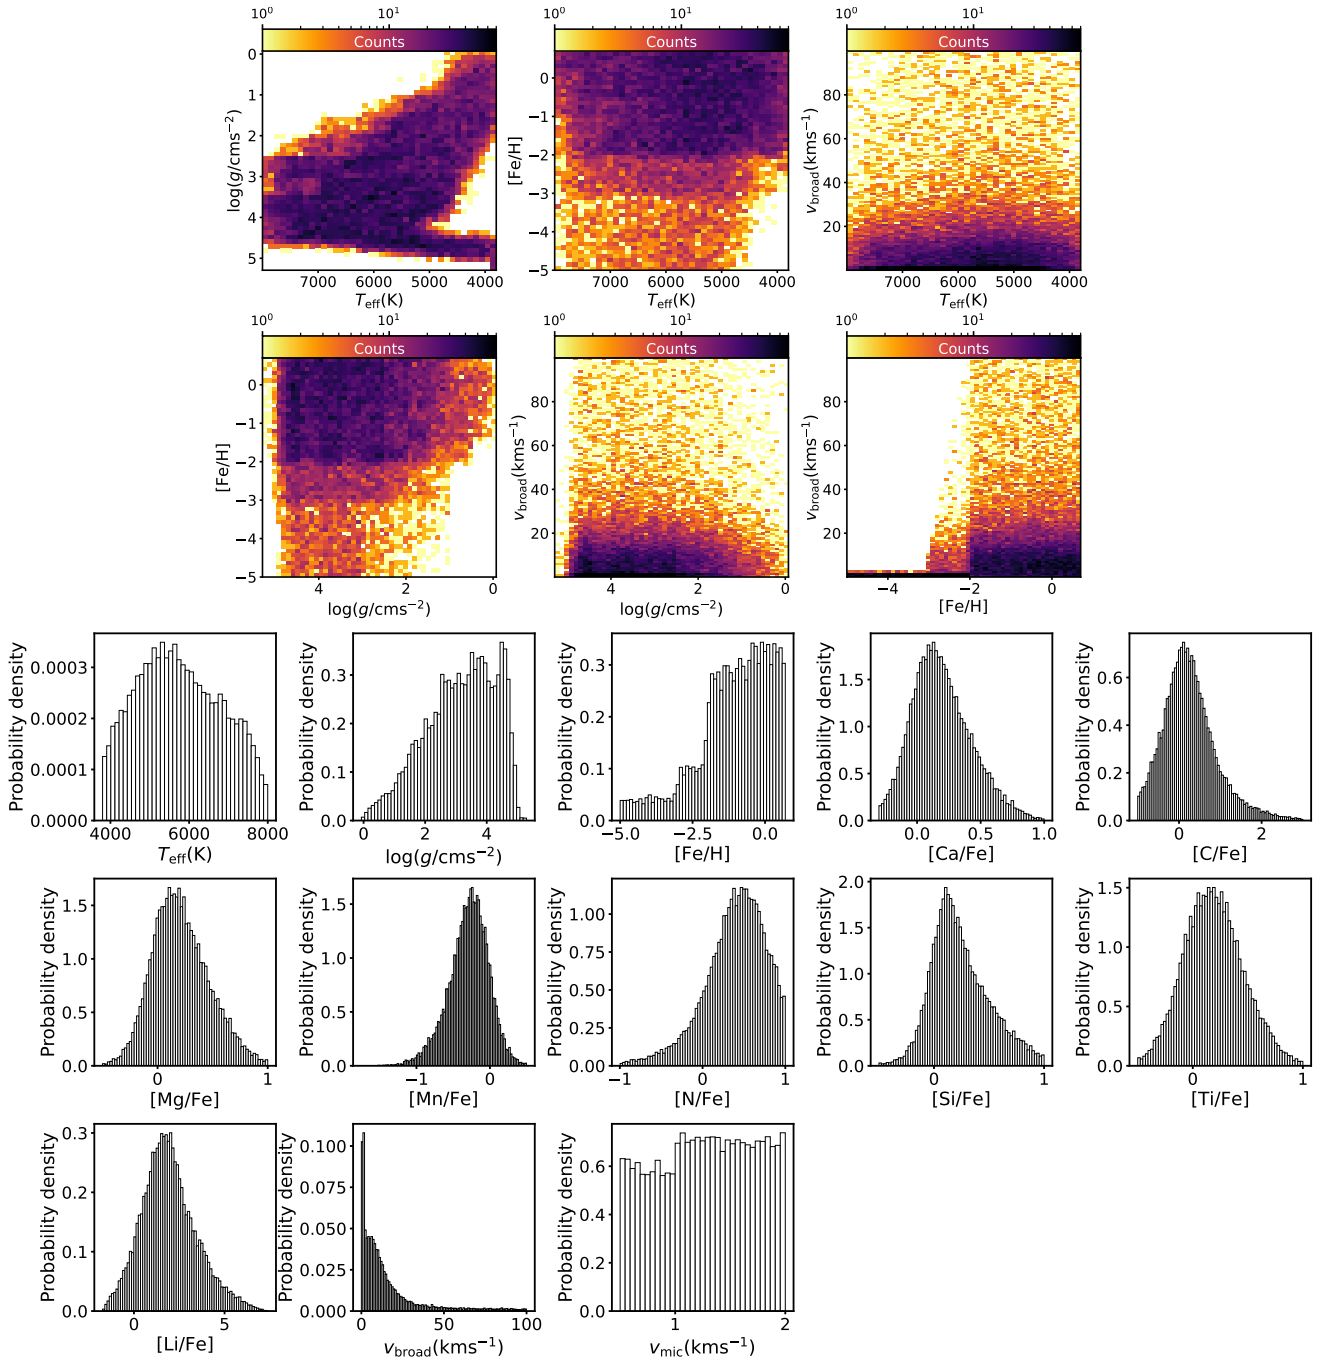


Fig. 1. Prior distributions of the target parameters in the training data. The panels in the top two rows show 2D histograms of the correlation between selected stellar properties. The panels in the bottom three rows show the 1D histograms of all target parameters.

Yakovleva et al. 2020), Ni (Bergemann et al. 2021; Voronov et al. 2022), Sr (Bergemann et al. 2012a; Gerber et al. 2023), Y (Storm & Bergemann 2023; Storm et al. 2024), Ba (Gallagher et al. 2020), and Eu (Storm et al. 2024). Subsequently, we convolved the synthetic stellar spectra with a Gaussian kernel to get a resolving power of $R = 80\,000$. We interpolated the spectra onto a grid with a pixel size of 0.025 \AA . Afterwards, rotational broadening (v_{broad}) was applied, sampled between 0 and 100 km s^{-1} . We note that we opted to not perform the immediate convolution with the 4MOST HR resolution so that we could retain the capability to test these grids on spectra with higher resolving power (also those analysed as part of 4MOST work, e.g. the *Gaia*-ESO benchmark stars, which are available at $R \gtrsim 47\,000$).

Therefore only in the final step, we ‘4MOST-ified’ the spectra by degrading them to R of the 4MOST high-resolution spectrograph ($R = 20\,000$) and to different signal-to-noise ratios (S/Ns) using the 4FS ETC package³. The synthetic spectra cover the full range of FGK-type stars in the $T_{\text{eff}}\text{-log}(g)$ diagram with $[\text{Fe}/\text{H}]$ from -5 to $+0.5$. In total, we generated a base set of 29 548 model NLTE spectra. Figure 1 shows the prior distributions of the relevant stellar surface parameters (i.e. T_{eff} , $\log(g)$, $[\text{Fe}/\text{H}]$, v_{broad} , v_{mic} , $[\text{Ca}/\text{Fe}]$, $[\text{C}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$, $[\text{Mn}/\text{Fe}]$, $[\text{N}/\text{Fe}]$, $[\text{Si}/\text{Fe}]$, $[\text{Ti}/\text{Fe}]$, and $[\text{Li}/\text{Fe}]$) considered in the training of our method

³ <https://science.aip.de/readthedocs/OpSys/etc/master/index.html>

across this base set. We note that the break in v_{broad} at around $[\text{Fe}/\text{H}] = -2$ in this model grid is based upon the assumption that old stars rotate more slowly overall. We then randomly split this database into 75, 20 and 5% sub-samples, dedicated to training, test, and validation, respectively. For each of these spectra, we then computed versions with 19 different S/N values (namely, 10, 12, 14, 16, 18, 20, 23, 26, 30, 35, 40, 45, 50, 80, 100, 130, 180, 250, and 1000 \AA^{-1}), motivated by the need to cover the entire dynamic range of spectral data quality potentially available with 4MOST (e.g. bright calibration stars versus faint distant targets in the deep bulge or disc-halo interface). After accounting for this data augmentation step, our final model spectra database across training, test, and validation sets consisted of 561 412 spectra.

2.2. Benchmark data

As 4MOST is not yet operational, we cannot benchmark our approach on real 4MOST data. To get an idea of how the cINN performs on observed spectra, we therefore used archival spectra of selected well-studied benchmark stars in addition to some other supplementary bright targets from our previous works. The spectra of the *Gaia*-ESO benchmark stars (Jofré et al. 2014; Heiter et al. 2015) were taken from Blanco-Cuaresma et al. (2014). These stars were observed with different astronomical facilities, including the Ultraviolet and Visual Echelle Spectrograph (UVES) at the 8 m Very Large Telescope (VLT) of the European Southern Observatory (ESO) in Chile and the Narval spectrograph at the 2.03 m Telescope Bernard Lyot (TBL) on the Pic du Midi de Bigorre in France. We note that some of the UVES spectra stem from the UVES Paranal Observatory Project (UVES POP, Bagnulo et al. 2003). Other stars were analysed in Fuhrmann et al. (1993); Fuhrmann (1998); Gehren et al. (2004, 2006); Bergemann & Gehren (2008); Bergemann & Cescutti (2010); Bergemann (2011), based on the spectroscopic data collected with the fibre optics Cassegrain échelle spectrograph (FOCES) at the 2.2 m Calar-Alto telescope. Table B.1 in the appendix lists the set of sources that we selected together with literature results for T_{eff} , $\log(g)$ and $[\text{Fe}/\text{H}]$ (with the corresponding references). In total, our test set consisted of 58 unique stars across 72 observed spectra, including 39 main sequence stars, 11 subgiants, and 8 giants from the K to F types. All the spectra were continuum normalised and corrected for radial velocity.

All of the benchmark spectra have a higher native resolution than 4MOST, existing on a different wavelength grid than the output of the 4MOST high resolution spectrograph and have varying S/Ns. To process them with the cINN approach built in this work, we had to first 4MOST-ify the spectra. This included a reduction of the resolution to $R = 20\,000$ (matching the target resolution of the 4MIDABLE-HR survey) via convolution with an appropriate Gaussian kernel and rebinning the spectra to the wavelength bin grid of the high resolution 4MOST spectrograph. As most of the spectra in our test set do not extend far enough into the blue to cover the blue arm (3926–4355 \AA) of the 4MOST high resolution spectrograph, the latter rebinning step was limited to the green (5160–5730 \AA) and red (6100–6790 \AA) arms only. Consequently, the network presented in this study, operates only on these two wavelength windows. While the cINN described further below was trained to handle different S/Ns in the input data, it is limited to the discrete set of S/N values currently generated in the training data (see Sect. 2.1). For the purposes of the benchmark conducted in Sect. 4.3, we decided to degrade all spectra to a fixed $(\text{S}/\text{N})_{\text{fixed}} = 250 \text{ \AA}^{-1}$. All spectra

considered have a native S/N that is higher, so that degrading the S/N (per wavelength bin) to that value was straightforward via the addition of Gaussian noise following

$$f_{i,\text{pert}} = f_i + \Delta f_i, \text{ with } \Delta f_i \propto \mathcal{N}(\mu = 0, \sigma = f_i \times (\text{S}/\text{N})_{\text{fixed}}^{-1}), \quad (1)$$

where f_i denotes the flux in wavelength bin i at the native S/N of the spectrum.

3. Method

The following provides a short overview of the cINN methodology (Sect. 3.1), the specific architecture and implementation choices for the model used in this work (Sect. 3.2), and the network training strategy (Sect. 3.3). In addition, we detail our procedure for hyperparameter optimisation (Sect. 3.4) and the post-processing procedure we applied to derive point estimates from the predicted posterior distributions (Sect. 3.5).

3.1. The conditional invertible neural network

Normalising flows (NFs, Tabak & Vanden-Eijnden 2010; Tabak & Turner 2013; Dinh et al. 2015; Rezende & Mohamed 2015; Kobyzev et al. 2021) are a family of generative deep learning models that excel at modelling complex distributions. To do so, NFs map simpler known probability distributions (that can be easily sampled from) to the complex target distribution by employing a series of invertible transformations (see e.g. Kobyzev et al. 2021, for a review). In the following, we utilise a variant of the NF framework, called the conditional invertible neural network (cINN, Ardizzone et al. 2019b). Among the NFs, the cINN architecture is particularly well suited for solving inverse problems. The latter comprise tasks, where it is the goal to infer the fundamental (physical) properties \mathbf{x} of some system (e.g. a star) given a corresponding set of observable quantities, \mathbf{y} . Inverse problems are ubiquitous in the natural sciences, but solving them (i.e. recovering the mapping $\mathbf{y} \rightarrow \mathbf{x}$) is often challenging. One of the main difficulties arises from the fact that for many problems in nature the forward process $\mathbf{x} \rightarrow \mathbf{y}$ that links the (unobservable) properties of the system to the actual observables suffers from an inherent loss of information. As a consequence, different combinations of the system properties can result in similar or even identical observations, rendering the inverse problem degenerate.

To tackle these challenges, the cINN uses a latent space approach to capture the information lost in the mapping from physical properties to actual observables. More specifically, instead of directly learning the inverse mapping $\mathbf{y} \rightarrow \mathbf{x}$, the cINN learns a mapping, f , of the physical properties, \mathbf{x} , to the latent variables, \mathbf{z} , conditioned on the observables, \mathbf{y} , that is,

$$\mathbf{z} = f(\mathbf{x}; \mathbf{c} = \mathbf{y}). \quad (2)$$

This training strategy employs the latent space to encode variance of the target parameters, \mathbf{x} , that is not captured by the corresponding observables, while ensuring that a prescribed prior distribution $P(\mathbf{z})$ for the latent variables is maintained during training. In this context, $P(\mathbf{z})$ is usually selected to be a multivariate normal distribution with zero mean and unit covariance and the dimension of the latent space matches that of the target parameter space per construction of the mapping. After a cINN is trained in this manner, given a new observation \mathbf{y}' the cINN can then query the variance encoded in the latent space via

sampling from the known prior $P(\mathbf{z})$ and, leveraging the invertible architecture of the cINN, compute samples of the posterior distribution $p(\mathbf{x}|\mathbf{y}')$ following

$$p(\mathbf{x}|\mathbf{y}') \sim g(\mathbf{z}; \mathbf{c} = \mathbf{y}') \quad \text{with } \mathbf{z} \propto P(\mathbf{z}), \quad (3)$$

where $g(\cdot; \mathbf{c}) = f^{-1}(\cdot; \mathbf{c})$ denotes the inverse of the forward mapping, f , for fixed condition, \mathbf{c} .

The fundamental architecture of the cINN as described above directly constrains the dimension of the latent space. However, it does not impose any such limitation on the dimension of the input feature space as the observations enter only as a conditioning input in the learned mappings f and g of the cINN. This not only allows $\dim(\mathbf{y})$ to be arbitrarily large, but also permits the addition of a secondary feature extraction network in a cINN. This feature extraction network, usually trained in tandem with the main cINN, can then transform the raw input observations into a more useful (learned) representation (usually lower dimensional), which can improve the predictive performance of a cINN (Ardizzone et al. 2019b). The cINN itself is comprised of a series of invertible, conditional affine coupling blocks (Dinh et al. 2017; Ardizzone et al. 2019b). Given an input vector, \mathbf{u} , each of these coupling blocks computes two complementary affine transformations with element-wise multiplication, \odot , and addition, \oplus , namely

$$\begin{aligned} \mathbf{v}_1 &= \mathbf{u}_1 \odot \exp(s_2(\mathbf{u}_2, \mathbf{c})) \oplus t_2(\mathbf{u}_2, \mathbf{c}), \\ \mathbf{v}_2 &= \mathbf{u}_2 \odot \exp(s_1(\mathbf{v}_1, \mathbf{c})) \oplus t_1(\mathbf{v}_1, \mathbf{c}), \end{aligned} \quad (4)$$

where \mathbf{u}_1 , \mathbf{u}_2 and \mathbf{v}_1 , \mathbf{v}_2 denote halves of the input \mathbf{u} and output vector \mathbf{v} , respectively. In these transformations, s_i and t_i represent arbitrarily complex mappings of the concatenation of $\mathbf{u}_i/\mathbf{v}_i$ and the conditioning input, \mathbf{c} (i.e. the input observations or output of a feature extraction network if present). From Eq. (4) we can see that the full transformation carried out by the coupling block is easily inverted given the output vector, $\mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2)$, following

$$\begin{aligned} \mathbf{u}_2 &= (\mathbf{v}_2 \ominus t_1(\mathbf{v}_1, \mathbf{c})) \odot \exp(-s_1(\mathbf{v}_1, \mathbf{c})), \\ \mathbf{u}_1 &= (\mathbf{v}_1 \ominus t_2(\mathbf{u}_2, \mathbf{c})) \odot \exp(-s_2(\mathbf{u}_2, \mathbf{c})), \end{aligned} \quad (5)$$

where \ominus denotes element-wise subtraction. The mappings s_i and t_i introduced in Eqs. (4) and (5) are never inverted themselves, regardless of whether the forward or backward passes of the coupling block are computed. This means that s_i and t_i are not required to be invertible transformations themselves. In fact, these mappings do not even need to be prescribed before the cINN is trained, but can be represented by small sub-networks instead (e.g. a simple fully connected feed forward network), such that s_i and t_i are learnt during training of the cINN (Ardizzone et al. 2019a,b).

3.2. Architecture and implementation details

Because most of the observed spectra in our set of benchmark observations (see Sect. 2.2) do not extend far enough to cover the blue arm (3926–4355 Å) of the high-resolution 4MOST spectrograph, the model we trained in the following used only the fluxes from the green (5160–5730 Å) and red arm (6100–6790 Å). An extension of our method to the blue arm once actual 4MOST spectra are available is certainly straightforward and preliminary tests on synthetic spectra indicate that the usage of the full 4MOST HR coverage will result in an even better performance of our approach. In addition to the fluxes, we also adopted the average S/N as a model input, so that the trained cINN can account

for the uncertainty in the input fluxes in the predicted posteriors, as previously demonstrated in Kang et al. (2023a,b, 2025). For simplicity, we further made sure that the input spectra to the cINN are both continuum normalised, as well as corrected for radial velocity. After some initial experimentation, we identified the following target parameters as viable to recover with the cINN given the input wavelength coverage (as per tests on synthetic data): T_{eff} , $\log(g)$, $[\text{Fe}/\text{H}]$, v_{mic} , v_{broad} , $[\text{Ca}/\text{Fe}]$, $[\text{C}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$, $[\text{Mn}/\text{Fe}]$, $[\text{N}/\text{Fe}]$, $[\text{Si}/\text{Fe}]$, $[\text{Ti}/\text{Fe}]$, and $[\text{Li}/\text{Fe}]$. As both T_{eff} and v_{broad} cover a rather large range, the network is trained to predict $\log(T_{\text{eff}}/\text{K})$ and $\log(v_{\text{broad}}/\text{km s}^{-1})$ instead. This procedure may avoid training and convergence issues that can occur when the target parameters span vastly different orders of magnitude. In addition, both the input fluxes (per wavelength bin) and all target parameters are also standardised across the training data. Standardisation refers here to the practice of subtraction of the mean and subsequent normalisation by the inverse standard deviation.

To implement the cINN for the purposes of this study, we employed the dedicated Framework for Easily Invertible Architectures (FrEIA⁴, Ardizzone et al. 2019a,b) package, based on the PYTORCH (Paszke et al. 2017) Python deep learning module. Furthermore, we adopted the Generative Flow (GLOW; Kingma & Dhariwal 2018) architecture for the affine coupling blocks, which introduces only two sub-networks per coupling block, jointly modelling (s_1, t_1) and (s_2, t_2) by one sub-network each, to improve computational efficiency. For the set-up of the sub-networks themselves, we selected a simple, fully connected network design that consists of three layers of size 512 with rectified linear unit (ReLU) activation functions. Ardizzone et al. (2019b) note that the exponential term $\exp(s)$ in Eqs. (4) and (5) could lead to instabilities during training of a cINN due to an exploding magnitude problem. To mitigate this issue, we also adapted the clamping procedure proposed in Ardizzone et al. (2019b), which modifies the argument, s , of the exponential function according to

$$s_{\text{clamp}} = \frac{2\alpha}{\pi} \arctan\left(\frac{s}{\alpha}\right), \quad (6)$$

where $\alpha = 1.9$. In addition, we introduced random permutation layers before each affine coupling block, which permute the input vector in a random but fixed (and therefore invertible) manner. The inclusion of these permutation layers serves as a way to intermix the information between the two streams \mathbf{u}_1 and \mathbf{u}_2 of the coupling blocks. As the input feature space is rather high dimensional given the high resolution of the 4MOST spectra, we also decided to add a feature extraction network to our cINN, so that the cINN can extract worthwhile features from the input data before passing them to the main body of the network. Specifically, we used a fully connected network with three hidden layers of size 2048, 1024, and 512, each employing the ReLU activation function, which returns 256 features in the output layer. We note that in preliminary experiments we also tested small convolutional neural networks for feature extraction in a cINN, but found that they did not yield significant performance improvements over the simpler fully connected architecture, while at the same time slowing down the inference. As described in Sect. 3.1, the feature extraction network is trained together with the main cINN itself on the same loss function. In our final network set-up (as determined by hyperparameter optimisation, see Sect. 3.4), the cINN itself consisted of 16 affine GLOW coupling blocks. A schematic overview of the full architecture is given in Fig. 2.

⁴ Available at <https://github.com/vislearn/FrEIA>

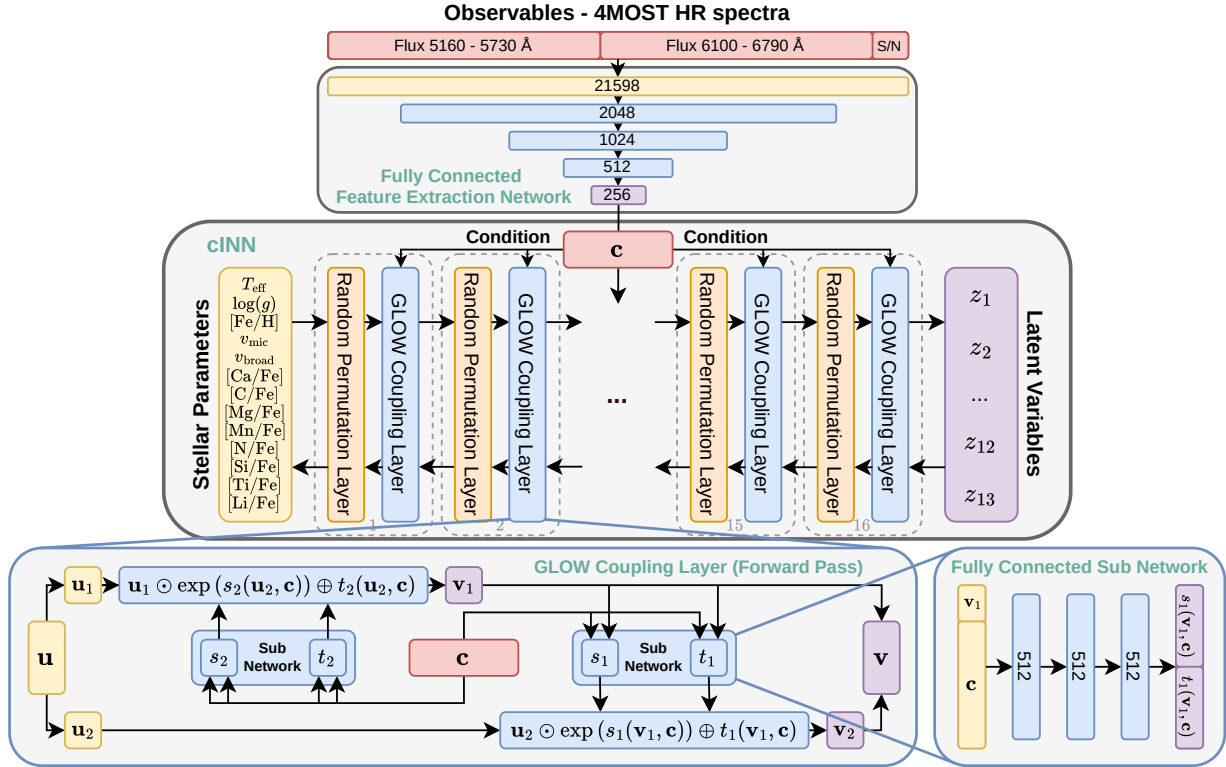


Fig. 2. Schematic overview of the cINN architecture. We note that the bottom zoom-in only highlights the forward pass of a GLOW coupling layer, following Eq. (4), but not the backward pass (Eq. (5)). Each of the 16 GLOW coupling blocks in our architecture employs two sub-networks, following the layout shown in the bottom right zoom-in.

3.3. Training set-up

To train the cINN, we followed the approach in Ardiszone et al. (2019b) and employed the maximum likelihood loss,

$$\mathcal{L} = \mathbb{E}_i \left(\frac{\|f(\mathbf{x}_i; \mathbf{c}_i, \Theta)\|_2^2}{2} - J_i \right), \quad (7)$$

where J_i is the determinant of the Jacobian matrix $J_i = \det(\partial f / \partial \mathbf{x}|_{\mathbf{x}=\mathbf{x}_i})$, evaluated at training example \mathbf{x}_i , as the primary optimisation goal for the network weights Θ . In addition, we imposed a secondary loss \mathcal{L}_{rev} for the backward pass of the cINN. Specifically, after the training example $(\mathbf{x}_i, \mathbf{c}_i)$ has passed through the cINN in the forward direction, that is $\mathbf{z}_i = f(\mathbf{x}_i; \mathbf{c}_i, \Theta)$, in the computation of \mathcal{L} , we can derive the secondary loss, which is also to be minimised as

$$\mathcal{L}_{\text{rev}} = \mathbb{E}_i [\mathbf{x}_i - g(\mathbf{z}_i + \delta \mathbf{z}_i; \mathbf{c}_i, \Theta)]^2, \quad (8)$$

where $\delta \mathbf{z}_i$ denotes a small amount of Gaussian noise, that is $\delta \mathbf{z}_i \sim \mathcal{N}(\mu = \mathbf{0}, \Sigma = 10^{-5} \times \mathbb{I})$. Here, \mathcal{L}_{rev} is supposed to promote a more robust embedding in the latent space, such that the invertibility of the network holds even when subject to small perturbations in the latent space. In practise, \mathcal{L}_{rev} drops significantly faster than \mathcal{L} , however, so that the maximum likelihood loss is overall the main driving force for the convergence of the weights during training.

To carry out the weight optimisation, we employed a standard stochastic gradient descent approach, namely the adaptive learning rate, momentum-based Adam (adaptive moment, Kingma & Dhariwal 2018) optimiser (with parameters $\beta_1 = \beta_2 = 0.8$). We initialised the Adam optimiser with a learning rate of $l_{\text{init}} = 0.001$ (in Adam this corresponds to a scaling factor applied

to the adaptive step size in the weight updates) and we made use of the common L2 weight regularisation formalism with $\lambda = 1.0 \times 10^{-5}$ in the optimisation process. The latter reduces the risk of overfitting to the training data by penalizing large weights in the network, which encourages convergence to a more generalisable solution (Goodfellow et al. 2016). In addition, we adopted a decaying learning rate schedule, reducing l by a factor of $\gamma = 0.95$ after every training epoch and we trained the cINN for 123 epochs (stopping early once the training loss converged) in total, processing 128 batches of size 1024 in every epoch (i.e. 30% of the total training data). We note that we did not loop over the entirety of our training data in every epoch here for the purpose of reducing the overall training time. Training a single cINN took about 7.5 h when using a NVIDIA RTX 6000 Ada graphics card for GPU acceleration.

3.4. Hyperparameter search

To derive the hyperparameters of our cINN architecture outlined in Sects. 3.3 and 3.2 above, we employed an optimisation approach called Hyperband (Li et al. 2018). It uses early stopping and adaptive resource allocation to provide a more efficient framework for random grid hyperparameter searches. In the Hyperband set-up, a maximum amount of training epochs per one network configuration is specified as a budget. Based on this maximum budget, Hyperband first generates differently sized groups of randomly generated hyperparameter configurations for the network. Based on the group size, Hyperband then employs a more or less aggressive pruning tactic during the hyperparameter search process (the larger the group, the more aggressive is the pruning). This pruning strategy determines for how many epochs at a time the different network configurations are allowed

to train before the least promising set-ups (based on the validation loss) are discarded. Specifically, the largest group trains only for a few epochs at a time before pruning, whereas smaller groups are allowed to train longer. In each group, this discarding procedure proceeds iteratively until only one network configuration is allowed to train for the specified maximum training epoch budget. In the final step, the ‘winning’ configurations in each group are then ranked by their validation loss. This hyperparameter search strategy tries to strike a balance between only training configurations until the end that appear promising early (i.e. converge fast) and allowing for some slower converging models to be trained as well, which could potentially converge to a better solution late in training. For more details on the method, we refer to Li et al. (2018).

3.5. Point estimates and sample rejection

As outlined in Sect. 3.1, our approach returns an estimate of the full posterior distribution of the target parameters given the input observations. However, for comparison with e.g. ground truth values in the subsequent performance analysis, it is often instructive to provide point estimates based on the predicted posterior distributions. In the following, we describe how we used the maximum a posteriori (MAP) estimate for this purpose, namely, the most likely combination of target parameters as determined by the posterior distribution. To find this most likely solution, we can use the cINN itself, as it not only generates samples of the posterior distributions, but also estimates the probability density of each generated sample. For each query observation, the MAP estimates are then identified as the posterior sample with the highest estimated probability density in the full 13 dimensional target parameter space of our model. In the following analysis presented below, we always generated 4096 posterior samples for each query observation, from which the MAP estimates are then determined. In our experiments, this number of samples has proven reliable enough to produce robust MAP estimates; thus, increasing the number of generated samples beyond this value did not markedly change the returned MAP values.

It is possible that some of the generated posterior samples fall outside of the boundaries of the training data. There are two main reasons for this. Firstly, the cINN always returns smooth posterior distributions. Consequently, it does not handle hard boundaries very well. If a prediction is close to the edges of the trained parameter space, the predicted posterior can, therefore, extend slightly beyond the formal limits of the training data. In the second scenario, the input observation is either just outside the space spanned by the training examples or vastly different altogether. In the former case, some slight extrapolation can be expected, resulting in a fraction of posterior samples to fall outside the training boundaries, whereas in the latter the cINN will likely have to extrapolate significantly. As with most machine learning approaches, there is no guarantee that the prediction is still sensible, if the model has to extrapolate far from its training boundaries (Bishop 2009; Goodfellow et al. 2016). To mitigate this potential issue, we included a sample rejection approach in our analysis of the predicted posteriors (e.g. when we determine the MAP). Specifically, we rejected all posterior samples that fell outside the target parameter boundaries of the training data (for the limits see Fig. 1). In principle, we could also use softer margins at the training boundaries to allow for some minor extrapolation, but for the analysis presented in this work, we decided to be strict. Using this rejection approach, we also introduced a ‘reliability’ metric. Specifically, for each query

spectrum, we determined the fraction of rejected posterior samples, f_{reject} (out of the 4096 initially generated ones). If f_{reject} is large (i.e. for $f_{\text{reject}} > 0.5$), the cINN extrapolated more often; thus, the prediction might have to be treated with caution.

3.6. INNs versus other generative approaches

As illustrated in the introduction, INNs are not the only generative deep learning approach that have seen usage in the context of spectra analysis. Conceptionally, models such as variational autoencoders (VAEs; Kingma & Welling 2013; Rezende et al. 2014) or generative adversarial networks (GANs; Goodfellow et al. 2014), or (more specifically) their conditional versions (cVAEs; Sohn et al. 2015, or cGANs; Mirza & Osindero 2014; Isola et al. 2016), could be applied to solve inverse problems in a similar fashion to INNs (Ardizzone et al. 2019a). While concrete performance will naturally vary depending on the individual architectures used, INNs or NFs generally offer some advantages. One central difference of NFs to VAEs or GANs is that they can not only efficiently generate samples of a target distribution, but, at the same time, also compute an exact density estimate for the generated samples (Kobyzev et al. 2021; Bond-Taylor et al. 2022). As outlined in Sect. 3.5, we used the latter capability, for instance, to derive the MAP estimates from the generated posteriors. In addition, Ardizzone et al. (2019a) demonstrated that INNs tend to return better calibrated posterior distributions and more accurate point estimates than a comparable cVAE approach. The posterior distributions predicted by INNs also default to the prior distribution of a certain target parameter if the conditioning input (i.e. input observation) does not contain enough information to constrain the parameter, whereas cVAEs may return spurious distributions in this case (Ardizzone et al. 2019a). In general, as VAEs only offer limited posterior approximations, they are more likely to underestimate variance in the posteriors, to induce biases in their MAP estimates and to suffer from mode collapse (Kobyzev et al. 2021; Bond-Taylor et al. 2022). The latter issue is also often encountered in GANs (Ardizzone et al. 2019b; Kobyzev et al. 2021; Bond-Taylor et al. 2022). In contrast, if trained via maximum likelihood, INNs are extremely robust to mode collapse by construction, as it is heavily penalised in this training strategy (Ardizzone et al. 2019b). Furthermore, maximum likelihood training is very stable, contrary to the adversarial training regime in GANs, which is notoriously hard and unstable (Ardizzone et al. 2019b; Bond-Taylor et al. 2022). Lastly, Liu & Regier (2020) also demonstrated that GANs struggle significantly with modelling low-dimensional density distributions in comparison to NFs.

4. Results

In this section, we summarise the performance of our trained cINN in terms of runtime (Sect. 4.1), on synthetic stellar spectra (Sect. 4.2) and on observed archival spectra of the *Gaia*-ESO/PLATO/4MOST benchmark stars (Sect. 4.3). We note that the following focuses on the predictive performance results for T_{eff} , $\log(g)$, v_{mic} , v_{broad} , [Fe/H], [Ca/Fe], [Mg/Fe], and [Li/Fe]. Even though our NLTE cINN also (jointly) predicts the abundances [C/Fe], [Mn/Fe], [N/Fe], [Si/Fe], and [Ti/Fe], we do not provide results for these chemical elements in this paper. A validation of the other elemental abundances will be a subject of a follow-up work by our group.

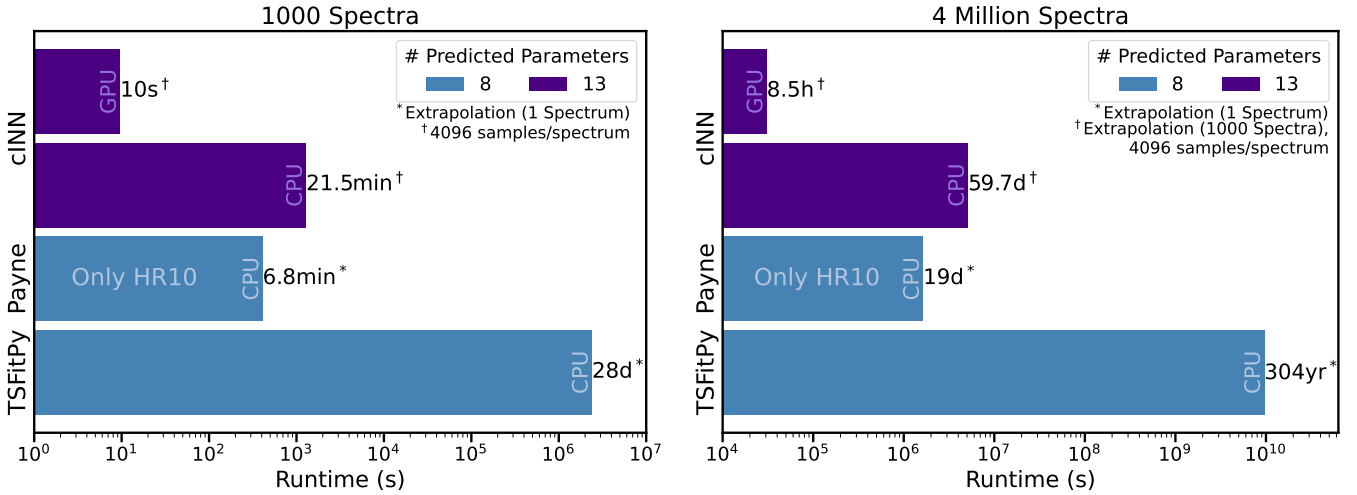


Fig. 3. Runtime comparison between the proposed cINN approach to the previously established methods TSFitPy and the Payne. The performance of the cINN was measured for the case of generating 4096 posterior samples per spectrum, using one core of an AMD EPYC 9554P (3.7 GHz) CPU and an NVIDIA RTX 6000 Ada GPU, respectively. The total cINN runtime is a combination of posterior generation and subsequent point estimation, but the latter becomes negligible quickly. We note that the runtimes in the right panel are extrapolated from the 1000 spectra (cINN) and 1 spectrum (TSFitPy and Payne) results, respectively. The Payne approach shown here operates on a more limited input parameter space, denoted as HR10, which covers the wavelength range from 534 to 562 nm and is based on the set-up of the GIRAFFE spectrograph.

4.1. Network efficiency

Figure 3 provides a summary of our cINN runtime, consisting of posterior generation plus subsequent point estimation according to Sect. 3.5, for example datasets of one thousand spectra and four million spectra, respectively. As demonstrated, the cINN approach is highly efficient and takes only 21.5 min on a single CPU core (AMD EPYC 9554P @3.7 GHz) or 10 s on a GPU (NVIDIA RTX 6000 Ada) to process 1000 query spectra without leveraging any additional parallelisation. Based on this, the cINN could process the output of, for instance, the 4MIDABLE-HR survey, expected to obtain more than four million spectra with the 4MOST facility⁵, in less than half a day when using GPU acceleration on a single GPU (when generating 4096 posterior samples per spectrum; i.e. more than 10^{10} posterior samples in total). For context, Fig. 3 also shows a comparison of the cINN runtime to the state-of-the-art classical fitting method TSFitPy, which is also our primary benchmark for the elemental abundances on the real spectra in Sect. 4.3. This highlights the great advantage of deep learning approaches, such as the cINN, as it would require 28 days (when determining only eight parameters on a single CPU) to process just 1000 spectra (or more than 300 years for four million spectra), for example. For reference, Fig. 3, also shows the runtime of a different machine learning method, namely the neural network approach known as Payne. We note that the Payne set-up used here operates on a smaller input space (denoted as HR10) that covers a wavelength range of 534–562 nm and is based on the set-up of the GIRAFFE spectrograph. This comparison demonstrates that the cINN runtime is on the same order of magnitude as the Payne, while, at the same time, accounting for a much larger input space, predicting five parameters more, and returning full posterior distributions in addition to a point estimate.

⁵ We note that this is the estimate of the number of spectra, whereas the number of stars to be observed is lower. For many stars, multi-epoch spectroscopic data will be available.

4.2. Tests on synthetic spectra

After training the cINN, we evaluated the predictive performance on the $N = 5910$ dedicated, held-out synthetic test spectra. In this, we computed two metrics based on the MAP estimates of the cINN X_{MAP} with respect to the ground truth, X_{GT} , to quantify the performance. The first are the mean, $\bar{\Delta}_{\text{cINN-GT}}$,

$$\bar{\Delta}_{\text{cINN-GT}}(X) = \frac{1}{N} \sum_{i=1}^N X_{i,\text{MAP}} - X_{i,\text{GT}}, \quad (9)$$

and the standard deviation $\sigma_{\text{cINN-GT}}$,

$$\sigma_{\text{cINN-GT}}(X) = \sqrt{\frac{\sum_{i=1}^N |X_{i,\text{MAP}} - X_{i,\text{GT}} - \bar{\Delta}_{\text{cINN-GT}}(X)|^2}{N-1}}, \quad (10)$$

of the residuals, used to gauge bias in the cINN MAP estimates and the average error, respectively. The second is the normalised root mean squared error (NRMSE), expressed as

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (X_{i,\text{MAP}} - X_{i,\text{GT}})^2}}{\hat{X}_{\text{max}} - \hat{X}_{\text{min}}}, \quad (11)$$

where \hat{X}_{max} and \hat{X}_{min} refer to the upper and lower boundaries of target parameter X in the training data, that serves to compare the predictive performance between different target parameters. In Fig. 4, we provide a summary of the cINN predictive performance in terms of $\bar{\Delta}_{\text{cINN-GT}}$ and $\sigma_{\text{cINN-GT}}$ across all S/N that the network is trained on (see also Tables A.1 and A.2 in the appendix). As we can see, the network performance depends notably on the quality of the input spectra as $\sigma_{\text{cINN-GT}}$ (and the NRMSE) increases across all target parameters the more noisy the input observations become. The bias $\bar{\Delta}_{\text{cINN-GT}}$ (although overall small) likewise tends to grow larger (in magnitude), with trends towards overestimation for T_{eff} and $[\text{Fe}/\text{H}]$, tendencies towards underestimation for $[\text{Ca}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$ and $[\text{Li}/\text{Fe}]$, and

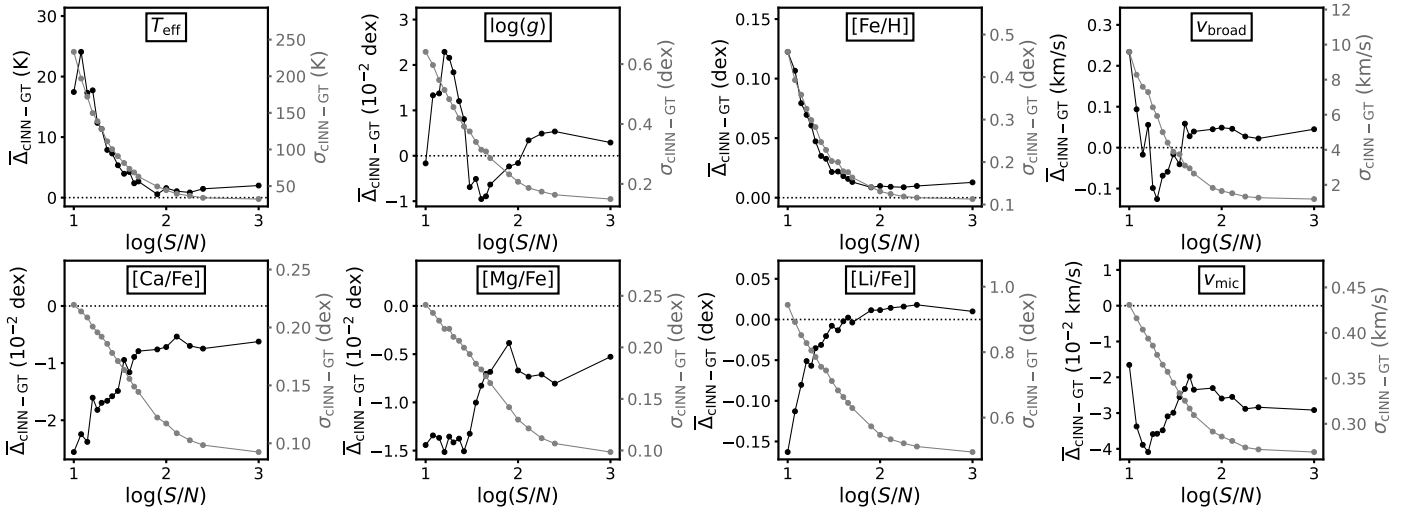


Fig. 4. Summary of cINN performance on synthetic spectra as a function of S/N. In each panel, the mean residual $\bar{\Delta}_{\text{cINN-GT}}$ is plotted in black on the left y-axis scale, while the standard deviation $\sigma_{\text{cINN-GT}}$ is indicated in grey and plotted on the right y-axis scale. The black dotted line indicates, where $\bar{\Delta}_{\text{cINN-GT}} = 0$ for reference. For more details, see Table A.1 and Fig. A.1.

mixed behaviour for $\log(g)$, v_{broad} and v_{mic} as S/N decreases (for more details see Appendix A and Fig. A.1). This is accompanied by a widening of the predicted posterior distributions with decreasing S/N, as illustrated by the median width of the 68% confidence interval $\Delta_{68\%}$ (of each 1D component of the predicted posterior) in Fig. A.2 in the appendix. Here, in particular, the abundances become quite uncertain at low S/N, exhibiting particularly wide posteriors. This result demonstrates that the cINN, on the one hand, correctly internalised that an increased level of white noise (which is assumed, but may not obviously be true in observed data) in the input fluxes warrants a greater uncertainty in the predicted parameters. On the other hand, it also indicates that the stellar parameters and abundances become harder to constrain for the cINN at low S/N, as evidenced by the increased $\sigma_{\text{cINN-GT}}$ (and NRMSE) of the MAP point estimates. Specifically, $\bar{\Delta}_{\text{cINN-GT}} \pm \sigma_{\text{cINN-GT}}$ increases from 2 ± 32 K, 0.003 ± 0.15 dex, 0.01 ± 0.11 , 0.05 ± 1.19 km s⁻¹, -0.03 ± 0.27 km s⁻¹, -0.006 ± 0.09 dex, -0.005 ± 0.1 dex and 0.01 ± 0.49 dex at S/N = 1000 Å⁻¹ for T_{eff} , $\log(g)$, [Fe/H], v_{broad} , v_{mic} , [Ca/Fe], [Mg/Fe] and [Li/Fe], respectively, to 18 ± 233 K, -0.002 ± 0.64 dex, 0.12 ± 0.46 dex, 0.23 ± 9.59 km s⁻¹, -0.017 ± 0.43 km s⁻¹, -0.026 ± 0.22 dex, -0.014 ± 0.24 dex and -0.16 ± 0.94 dex at S/N = 10 Å⁻¹. As demonstrated in Fig. 4 (see also Fig. A.2 in the appendix), our final cINN performs best in the S/N range between 80 and 1000 Å⁻¹, maintaining a comparable performance. Below S/N = 80 Å⁻¹, however, both intrinsic uncertainties (i.e. posterior widths) and average errors of the prediction start to increase significantly. We note that for the analysis of low S/N observations, a more specialised cINN, trained only on low S/N data, could potentially perform better than the general purpose model presented here.

Next, we discuss the results obtained with the NLTE cINN on synthetic spectra at S/N = 250 Å⁻¹ more in detail, as this is the expected S/N of the high-resolution spectra of stars to be observed in the 4MOST 4MIDABLE-HR program. Figure 5 shows a 2D histogram comparing the cINN MAP estimates on the synthetic test spectra to the ground truth. At this S/N, we find that our NLTE cINN recovers atmospheric parameters and abundances of stars with $\bar{\Delta}_{\text{cINN-GT}} \pm \sigma_{\text{cINN-GT}}$ values of

1.46 ± 33 K, 0.005 ± 0.16 dex, 0.01 ± 0.12 dex, -0.007 ± 0.1 dex, -0.008 ± 0.11 dex, 0.02 ± 0.51 dex, 0.02 ± 1.27 km s⁻¹ and -0.03 ± 0.27 km s⁻¹ for T_{eff} , $\log(g)$, [Fe/H], [Ca/Fe], [Mg/Fe], [Li/Fe], v_{broad} and v_{mic} , respectively. Looking at the NRMSEs, we find that T_{eff} is comparatively recovered the best with an NRMSE of 0.008, whereas v_{mic} appeared the hardest to constrain with an NRMSE of 0.182.

Analogously to the analysis of the posterior width as a function of S/N, we also investigated how the cINN performance behaves in dependence of the primary stellar parameters, T_{eff} , $\log(g)$ and [Fe/H] for the fixed S/N = 250 Å⁻¹. This analysis revealed that the chemical composition, and in particular [Fe/H], is the most critical parameter that influences the quality of the results. Specifically, for [Fe/H] ≤ -2.5, the predicted posteriors for the abundances become so wide that the cINN returns the entire training parameter range as a possible solution (see Fig. A.3 for an illustration of the full posterior width as a function of [Fe/H]). This increase in the intrinsic uncertainty of the cINN prediction is also reflected by increased $\sigma_{\text{cINN-GT}}$ values, rising, for example, by a factor of 2 for [Mg/Fe] or [Ca/Fe] when compared to synthetic spectra with [Fe/H] > -2.5. One potential explanation for this behaviour is the large dynamical range of chemical abundances spanning five orders of magnitude in all key elements, such as C, Fe, Mg, Ti. This leads to significant variations in atomic and molecular opacity over the parameter space of FGKM-type stars, being especially critical for red giants with their lower surface temperatures favouring formation of molecules (Plez 2012; Eitner et al. 2025). Apart from that, the grid of available synthetic stellar spectra is rather sparse for [Fe/H] ≤ -2. This regime is typically associated with very metal-poor (VMP) and extremely metal-poor (EMP) stars (Christlieb et al. 2005). Due to this relative scarcity of training examples, the cINN cannot achieve the same precision in this (VMP) metallicity regime, which was also demonstrated in other studies employing machine learning (see Table 1 in Kovalev et al. 2019). We note that there are efforts within the 4MOST Infrastructure Group 7 (IWG7) to compute dedicated spectral grids in this domain, and therefore we reserve the optimisation of our cINN approach in the VMP and EMP regime to a follow-up study.




Fig. 5. 2D histograms that compare the cINN MAP estimates to the corresponding ground truth for 5910 synthetic spectra with $S/N = 250$ held-out in the test set. The dotted black guide line indicates a perfect one-to-one match.

Finally, we explored the decision-making process of the trained cINN by computing saliency maps for all our target parameters. Saliency describes the average network internal gradients of each target parameter with respect to each input feature $\partial X/\partial f_\lambda$ (where f_λ is the input flux for a given wavelength bin). Looking, in particular, at the resulting saliency maps of the predicted abundances (see Figs. A.4 and A.5 in the appendix for an illustration), we found that the most significant peaks in the saliency (i.e. the features with the highest importance) almost always coincide with absorption lines of the corresponding element (see e.g. the [Ca/Fe], [Mg/Fe] and [Li/Fe] panels in Fig. A.5). This indicates that the trained cINN has learned to pay attention to parts of the spectrum, i.e. the spectral lines of different elements, that also play an important part in classical analysis techniques (albeit not necessarily in the same exact combination). This confirms that the decision making process of the trained network has picked up on some of the physics used in the computation of the synthetic stellar spectral grids. This is not unexpected, because the cINN is essentially a function – the result of an optimisation problem (here, training the cINN) – and by definition mathematically constructed such that the input data are reproduced as closely as possible given the degeneracies associated with the multi-dimensionality.

4.3. Test on benchmark spectra

After verifying the network performance on the synthetic test spectra, we applied the cINN to the set of archival spectra for the benchmark stars, as outlined in Sect. 2.2. Stellar parameters, in particular T_{eff} and $\log(g)$, for these targets were assembled from the literature (see Table B.1). For the *Gaia*-ESO benchmark stars, T_{eff} and $\log(g)$ were obtained using methods that are almost independent of spectroscopy (Heiter et al. 2015; Jofré et al. 2014). In particular the T_{eff} values were constrained using

interferometric angular diameters, whereas the $\log(g)$ values were obtained either on the basis of evolutionary phase considerations (models of stellar interior) or from global asteroseismic parameters. For the stars observed with the FOCES instrument (Sect. 2.2), however, no such constraints are available, therefore their estimates rely on spectroscopy.

Since $[X/H]$ abundance ratios are very sensitive to the physics of stellar spectral models and of the spectroscopic analysis approach (Bergemann et al. 2012b; Lind et al. 2012; Bergemann et al. 2017, 2019; Amarsi et al. 2019, 2024), we adopted our own abundances for these benchmark stars. Here we used the TSFitPy code (see Sect. 2) to determine NLTE values of [Fe/H], [Li/Fe], [Mg/Fe], and [Ca/Fe] for each star in Table B.1 directly. The diagnostic lines are provided in Table B.2 (appendix). In these abundance calculations, the values of T_{eff} , $\log(g)$, v_{broad} and v_{mic} were set to the values of the cINN MAP estimates. The results of our TSFitPy calculations are provided in Table B.3. For some spectra, it was not possible to determine some elemental abundances, such as [Li/Fe], because of the weakness of the spectral features.

When comparing our initial cINN predictions for stellar parameters and abundances with the afore-mentioned estimates, we identified a small bias in the predictions, which correlates with T_{eff} (see Table B.4 for the raw cINN predictions). This bias is also known from previous studies with the Payne approach (Gent et al. 2022), which requires the use of line masks in the form of a wavelength-dependent noise model, or alternatively more sophisticated Bayesian-supported noise models to improve the solutions (Bestenlehner et al. 2024). Here, we introduced a simple linear correction (see Appendix B.1 and Table B.5 for details), which removes the underlying T_{eff} -dependent bias in our parameter and abundance estimates. As a result of applying this correction, we were able to achieve stable and consistent results for the entire parameter space of our observed FGKM-type stars.

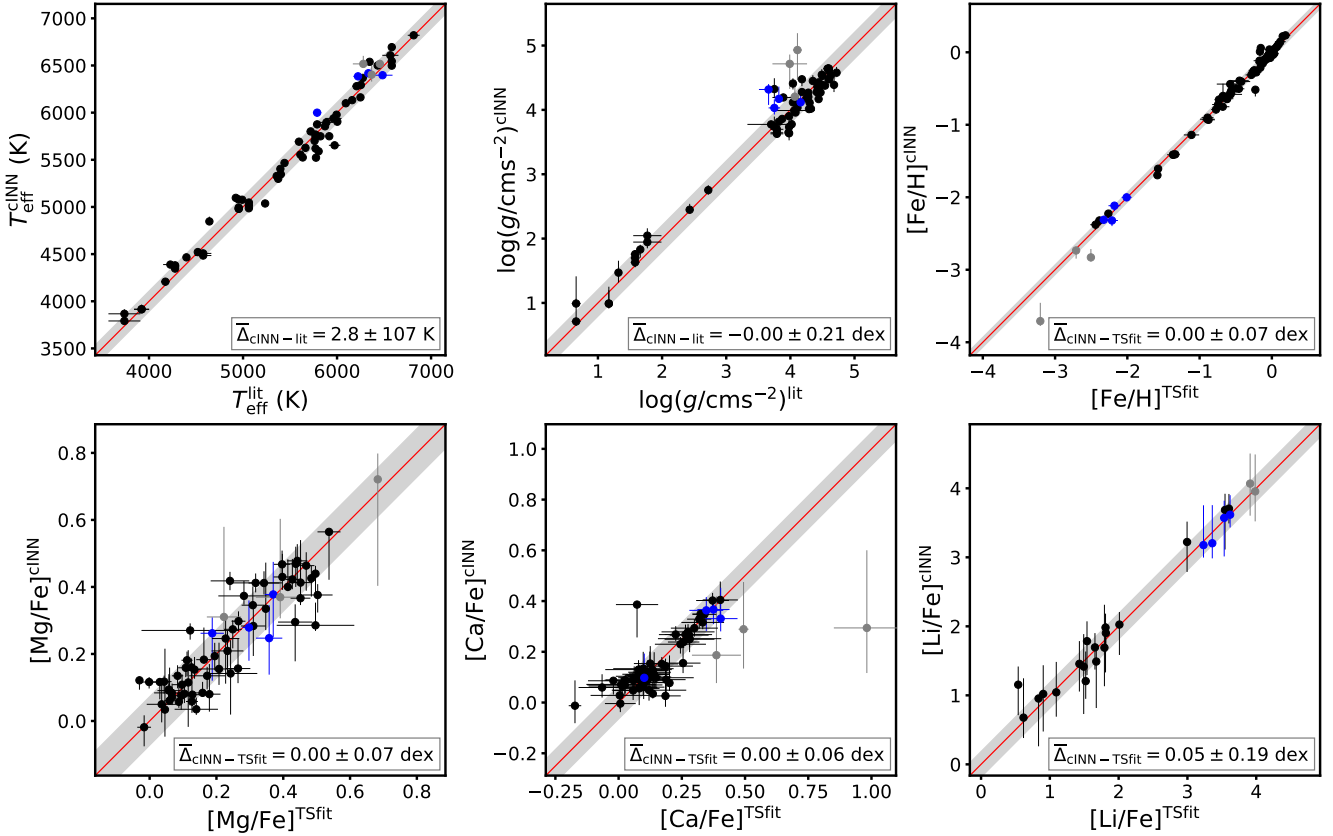


Fig. 6. Comparison of cINN predictions after correction for systematic biases to the literature stellar parameters (X^{lit}) and abundance fits with Turbospectrum (X^{TSfit}) for the real benchmark spectra listed in Table B.1. We note that some stars have multiple spectra in our benchmark set. For some spectra, there are no uncertainties listed in the literature for T_{eff} and $\log(g)$, so no x-error bars are displayed in these cases (see Table B.1). The T_{eff} uncertainties are often too small to be visible in the top left panel. The y-error bars indicate the 1D 68% confidence intervals determined from the predicted posterior distributions. Stars indicated in grey have $[\text{Fe}/\text{H}] \leq -2.5$ and were not considered in the fit for the bias correction (but are corrected here). The grey shaded area around the 1-to-1 correlation represents the 1σ standard deviation of the residuals of the MAP estimate with respect to the reference value $X^{\text{cINN,MAP}} - X^{\text{lit/TSfit}}$. The stars highlighted in blue correspond to the example stars shown in Figs. 7 and B.2.

In Fig. 6, we show our resulting stellar parameters and chemical abundances obtained with the NLTE cINN (y-axis) in comparison to those obtained with TsFitPy and from literature sources (x-axis). Our final values are also provided in Table B.3. For all elements ($[\text{Fe}/\text{H}]$, $[\text{Ca}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$ and $[\text{Li}/\text{Fe}]$), the agreement between independent estimates is very good and it is typically within the combined uncertainties of both values. The typical errors ($\sigma_{\text{cINN-lit/TSfit}}$) of our predictions are 107 K for T_{eff} , 0.21 dex for $\log(g)$, 0.07 dex for $[\text{Fe}/\text{H}]$, 0.07 dex for $[\text{Mg}/\text{Fe}]$, 0.06 dex for $[\text{Ca}/\text{Fe}]$, and 0.19 dex for $[\text{Li}/\text{Fe}]$, respectively. It is important to note here, that our correction currently extends only down to $[\text{Fe}/\text{H}] \approx -2.5$ (for the reasons outlined in the previous section) and is limited to an interval of 0.5 to 3.9 in $[\text{Li}/\text{Fe}]$. The predictions in Table B.3 for the few very metal poor stars and where $[\text{Li}/\text{Fe}]^{\text{cINN}} < 0$ may, therefore, need to be treated with caution. In a greater context, the performance of our cINN does not only allow for a robust analysis of observed spectra, as indicated by the successful recovery of the Turbospectrum NLTE abundances, but should also be sufficiently accurate to facilitate studies that aim to chemically distinguish Galactic populations, such as the thin and thick disc, as demonstrated for example by Gent et al. (2024).

As a final, qualitative evaluation of the NLTE cINN performance on the real data, we used Turbospectrum (Gerber et al. 2023; Storm & Bergemann 2023) to compute theoretical spectra that correspond to the cINN estimates of stellar parameters

and abundances. Our model spectra are then compared to the observed wavelength-dependent datasets. We note that for this analysis, we used all 13 output parameters of the cINN, but only applied the bias correction to the four abundances (i.e. $[\text{Fe}/\text{H}]$, $[\text{Mg}/\text{Fe}]$, $[\text{Ca}/\text{Fe}]$, and $[\text{Li}/\text{Fe}]$) it is available for. The reason for this is that there is no one-to-one mapping between stellar parameters and the cINN functional behaviour, and stellar parameters have a global non-linear impact on the spectra, also affecting the shape of the continuum. For abundances it is, however, appropriate to do this, because the diagnostic elements (Mg I, Fe I, Li I, Ca I) are trace (minority) species in FGKM-type stellar atmospheres (Bergemann & Nordlander 2014) and have no impact on the optical continuum. Therefore, by applying abundance corrections we can be confident that only the fluxes in the lines of these specific elements are affected. Furthermore, to visualise the effect of the uncertainties in the predicted abundances of $[\text{Fe}/\text{H}]$, $[\text{Mg}/\text{Fe}]$, $[\text{Ca}/\text{Fe}]$, and $[\text{Li}/\text{Fe}]$, we kept all nine other predicted parameters fixed to the MAP estimate and computed four additional models using the 2.5, 16, 84, and 97.5% quantiles of the abundances to approximate the 68 and 95% confidence intervals of the re-simulated flux.

In Fig. 7, we show the observed spectra of four stars and overplot the simulated spectra computed for the cINN parameter predictions. We in particular highlight the diagnostic absorption lines of Mg I, Ca I, Li I, and Fe I. We refer to Fig. B.2 in the appendix, where the spectra are presented in full. Generally, we




Fig. 7. Comparison of the 4MOST-ified input spectra to model spectra corresponding to the cINN MAP estimates for four low-metallicity stars, focusing on prominent absorption lines of Mg, Ca, Li and Fe. We note that the re-simulation employed all 13 stellar parameters predicted by the cINN, but used the bias correction from Sect. 4.3 only for the abundances of Ca, Mg, Li, and Fe (see also Table B.5).

find that the model spectra computed using the cINN estimates fit the observed data well. Averaged over all ($72 \times 25\,197$) wavelength bins of our benchmark set, the fluxes of the re-simulated spectra (corresponding to the MAP estimates) and the input observations match to within $0.69^{+2.12}_{-0.52}$ %.

5. Summary and outlook

In this paper, we introduce a new, simulation-based, machine-learning method that employs a conditional invertible neural network (cINN) to determine astrophysical parameters of stars from observed stellar spectra. The parameters include effective temperature, T_{eff} , surface gravity, $\log(g)$, metallicity, $[\text{Fe}/\text{H}]$, broadening, v_{broad} , microturbulent velocity, v_{mic} and the elemental abundance ratios, $[\text{Ca}/\text{Fe}]$, $[\text{Mg}/\text{Fe}]$, and $[\text{Li}/\text{Fe}]$. The framework presented can be applied to data with the spectral

characteristics (resolving power, noise, wavelength coverage) of the 4MOST high-resolution ($R = 20\,000$, wavelength coverage: $5160\text{--}5730\text{ \AA}$, $6100\text{--}6790\text{ \AA}$) spectrograph (de Jong et al. 2019).

Our cINN was trained on a grid of 29 548 synthetic stellar spectra, computed in NLTE and using MARCS model atmospheres (Gustafsson et al. 2008) with the Turbospectrum code (Gerber et al. 2023; Storm & Bergemann 2023). We further augmented this base model grid by computing a version of each spectrum at 19 different S/N values, totalling 561 412 synthetic spectra. The latter were then split 75, 20, and 5% into the dedicated training, test and validation sub-sets, respectively. The cINN is part of the normalising flow family and with that returns predictions of the full posterior distribution of the target stellar parameters and elemental abundances. This allows the cINN to both provide an internally consistent uncertainty estimate and to highlight potential degeneracies in the inverse problem, while

remaining highly efficient at inference time. The cINN presented in this study can evaluate 1000 high-resolution 4MOST spectra within a few 10s of seconds when making use of GPU acceleration (on an NVIDIA RTX 6000 Ada), which means it has the capability to process the entire output of a large 4MOST survey, such as e.g. 4MIDABLE-HR (over four million spectra), in less than half a day.

We first tested the performance of our trained cINN on held-out synthetic test spectra, where we found that our NLTE cINN has the capacity to recover stellar parameters and abundances with average errors σ (i.e. standard deviation of the residuals) of 33 K for T_{eff} , 0.16 dex for $\log(g)$, 0.12 dex for [Fe/H], 1.27 km s⁻¹ for v_{broad} , 0.27 km s⁻¹ for v_{mic} , 0.1 dex for [Ca/Fe], 0.11 dex for [Mg/Fe], and 0.51 dex for [Li/Fe], respectively, at an average S/N = 250 Å⁻¹ of the input spectra, for example. A more detailed analysis revealed that our NLTE cINN performs the best for stars with [Fe/H] > -2.5 and spectra with S/N > 80 Å⁻¹. Outside of this regime, we observed a significant increase in the intrinsic uncertainties of the cINN predictions (i.e. width of the posteriors) that is accompanied by an increase in average errors σ . Here, we found that the [Fe/H] dependency of the predictive performance can likely be attributed to a relative lack of training examples at very low metallicity in the available training data. The decrease in performance with S/N, on the other hand, could be explained by the general increase in complexity that comes with more noisy input data and the general purpose approach of the presented cINN that aimed to make the method applicable to a wide range of S/Ns from 10 to 1000 Å⁻¹, rather than specialising in low-quality spectra alone, for instance.

Next, we applied the NLTE cINN to a set of 72 real benchmark and bright star spectra, degraded to the resolution and wavelength coverage of 4MOST. We compared the predicted maximum a posteriori (MAP) estimates of the cINN to literature results for T_{eff} and $\log(g)$, and to [Fe/H] and elemental abundances derived with classical evaluation techniques, using the TSPy code. For this introductory paper, we focused here on the Ca, Mg, and Li abundances, reserving an analysis of the other elements to a follow-up study. In this comparison, we found a generally good agreement of the cINN estimates with the literature/TSPy results, but we identified a slight systematic offset as a function of the predicted T_{eff} . To account for this apparent bias, we computed a linear correction (see Eq. (B.1) and Table B.5). When including this correction, we found an agreement of the cINN MAP estimates with the independent measures with σ values of 107 K for T_{eff} , 0.21 dex for $\log(g)$, 0.07 dex for [Fe/H], 0.07 dex for [Mg/Fe], 0.06 dex for [Ca/Fe], and 0.19 dex for Li. This shows that a cINN-based approach is a promising avenue for the comprehensive and efficient NLTE analysis of large datasets to be obtained with the 4MOST facility in, for instance, the upcoming 4MIDABLE HR survey (Bensby et al. 2019).

In a planned follow-up to this proof-of-concept study, we will extend our analysis to the remaining elemental abundances that are predicted by the network (i.e. C, Mn, N, Si, and Ti) to verify their proper recovery. In addition, we plan to expand our suite of synthetic NLTE spectra with additional examples at very low metallicity to improve the cINN predictive performance in this regime. Further improvements can also be made in the S/N coverage by moving beyond the discrete set of options shown in this work to a continuous treatment of the S/N by using, for instance, a training strategy with dynamic data augmentation. Once we have realised these complementary analyses and improvements, we plan to release our model as a free tool to the astronomical community.

Acknowledgements. The team at Heidelberg University acknowledges financial support from the European Research Council via the ERC Synergy Grant “ECOGAL” (project ID 855130), from the German Excellence Strategy via the Heidelberg Cluster “STRUCTURES” (EXC 2181 - 390900948), and from the German Ministry for Economic Affairs and Climate Action in project “MAINN” (funding ID 50002206). They are also grateful for computing resources provided by the Ministry of Science, Research and the Arts (MWK) of the State of Baden-Württemberg through bwHPC and the German Science Foundation (DFG) through grants INST 35/1134-1 FUGG and 35/1597-1 FUGG, and also for data storage at SDS@hd funded through grants INST 35/1314-1 FUGG and INST 35/1503-1 FUGG. V.K. also acknowledges financial support by the Carl-Zeiss-Stiftung. R.A. and N.S. acknowledge support from the Lise Meitner grant from the Max Planck Society (grant PI: M. Bergemann) and from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 949173, PI: M. Bergemann). G.G. acknowledges support by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project-IDs: eBer-22-59652 (GU 2240/1-1 “Galactic Archaeology with Convolutional Neural-Networks: Realising the potential of *Gaia* and 4MOST”).

References

- Abdurro’uf, Accetta, K., Aerts, C., et al. 2022, *ApJS*, 259, 35
 Aguilera-Gómez, C., Ramírez, I., & Chanamé, J. 2018, *A&A*, 614, A55
 Amarsi, A. M., Nissen, P. E., & Skúladóttir, Á. 2019, *A&A*, 630, A104
 Amarsi, A. M., Lind, K., Osorio, Y., et al. 2020, *A&A*, 642, A62
 Amarsi, A. M., Ogneva, D., Buldgen, G., et al. 2024, *A&A*, 690, A128
 Arduzzone, L., Kruse, J., Rother, C., & Köthe, U. 2019a, in *International Conference on Learning Representations*
 Arduzzone, L., Lüth, C., Kruse, J., Rother, C., & Köthe, U. 2019b, *CoRR*, abs/1907.02392 [arXiv:1907.02392]
 Asplund, M. 2005, *ARA&A*, 43, 481
 Asplund, M., Amarsi, A. M., & Grevesse, N. 2021, *A&A*, 653, A141
 Bagnulo, S., Jehin, E., Ledoux, C., et al. 2003, *The Messenger*, 114, 10
 Bensby, T., Bergemann, M., Rybizki, J., et al. 2019, *The Messenger*, 175, 35
 Bergemann, M. 2011, *MNRAS*, 413, 2184
 Bergemann, M., & Cescutti, G. 2010, *A&A*, 522, A9
 Bergemann, M., & Gehren, T. 2008, *A&A*, 492, 823
 Bergemann, M., & Hoppe, R. 2025, arXiv e-prints [arXiv:2511.04254]
 Bergemann, M., & Nordlander, T. 2014, in *Determination of Atmospheric Parameters of B*, eds. E. Niemczura, B. Smalley, & W. Pych (Berlin: Springer), 169
 Bergemann, M., Pickering, J. C., & Gehren, T. 2010, *MNRAS*, 401, 1334
 Bergemann, M., Hansen, C. J., Bautista, M., & Ruchti, G. 2012a, *A&A*, 546, A90
 Bergemann, M., Lind, K., Collet, R., Magic, Z., & Asplund, M. 2012b, *MNRAS*, 427, 27
 Bergemann, M., Kudritzki, R.-P., Würl, M., et al. 2013, *ApJ*, 764, 115
 Bergemann, M., Collet, R., Amarsi, A. M., et al. 2017, *ApJ*, 847, 15
 Bergemann, M., Gallagher, A. J., Eitner, P., et al. 2019, *A&A*, 631, A80
 Bergemann, M., Hoppe, R., Semenova, E., et al. 2021, *MNRAS*, 508, 2236
 Bestenlehner, J. M., Enßlin, T., Bergemann, M., et al. 2024, *MNRAS*, 528, 6735
 Bishop, C. M. 2009, *Pattern Recognition and Machine Learning*, 8th edn., Information science and statistics (New York: Springer), XX, 738
 Bister, T., Erdmann, M., Köthe, U., & Schulte, J. 2022, *Euro. Phys. J. C*, 82, 171
 Blanco-Cuaresma, S., Soubiran, C., Jofré, P., & Heiter, U. 2014, *A&A*, 566, A98
 Bond-Taylor, S., Leach, A., Long, Y., & Willcocks, C. G. 2022, *IEEE Trans. Pattern Anal. Mach. Intell.*, 44, 7327
 Buder, S., Kos, J., Wang, X. E., et al. 2025, *PASA*, 42, e051
 Caffau, E., Ludwig, H.-G., Steffen, M., & Bonifacio, P. 2010, *IAU Symp.*, 268, 329
 Candebat, N., Sacco, G. G., Magrini, L., et al. 2024, *A&A*, 692, A228
 Casamiquela, L., Soubiran, C., Jofré, P., et al. 2026, *A&A*, 705, A167
 Christlieb, N., Beers, T. C., Thom, C., et al. 2005, *A&A*, 431, 143
 Cvrček, V., Romaniello, M., Šára, R., Freudling, W., & Ballester, P. 2025, *A&A*, 693, A256
 de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *The Messenger*, 175, 3
 De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, 449, 2604
 Dinh, L., Krueger, D., & Bengio, Y. 2015, in *3rd International Conference on Learning Representations, ICLR 2015*, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings, eds. Y. Bengio, & Y. LeCun
 Dinh, L., Sohl-Dickstein, J., & Bengio, S. 2017, in *International Conference on Learning Representations*
 Eisert, L., Pillepich, A., Nelson, D., et al. 2023, *MNRAS*, 519, 2199
 Eitner, P., Bergemann, M., Hoppe, R., et al. 2025, *A&A*, 703, A199
 Ezzeddine, R., Merle, T., Plez, B., et al. 2018, *A&A*, 618, A141
 Fuhrmann, K. 1998, *A&A*, 338, 161

- Fuhrmann, K., Axer, M., & Gehren, T. 1993, *A&A*, 271, 451
- Gaia Collaboration (Prusti, T., et al.) 2016, *A&A*, 595, A1
- Gallagher, A. J., Bergemann, M., Collet, R., et al. 2020, *A&A*, 634, A55
- García Pérez, A. E., Sánchez-Blázquez, P., Vazdekis, A., et al. 2021, *MNRAS*, 505, 4496
- Gebran, M., & Bentley, I. 2025, *Astronomy*, 4, 13
- Gebran, M., Farah, W., Paletou, F., Monier, R., & Watson, V. 2016, *A&A*, 589, A83
- Gehren, T., Liang, Y. C., Shi, J. R., Zhang, H. W., & Zhao, G. 2004, *A&A*, 413, 1045
- Gehren, T., Shi, J. R., Zhang, H. W., Zhao, G., & Korn, A. J. 2006, *A&A*, 451, 1065
- Gent, M. R., Bergemann, M., Serenelli, A., et al. 2022, *A&A*, 658, A147
- Gent, M. R., Eitner, P., Serenelli, A., et al. 2024, *A&A*, 683, A74
- Gerber, J. M., Magg, E., Plez, B., et al. 2023, *A&A*, 669, A43
- Gilmore, G., Randich, S., Asplund, M., et al. 2012, *The Messenger*, 147, 25
- Gilmore, G., Randich, S., Worley, C. C., et al. 2022, *A&A*, 666, A120
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., et al. 2014, in *Advances in Neural Information Processing Systems*, eds. Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (USA: Curran Associates, Inc.), 27
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge: MIT Press)
- Guiglion, G., Matijević, G., Queiroz, A. B. A., et al. 2020, *A&A*, 644, A168
- Guiglion, G., Nepal, S., Chiappini, C., et al. 2024, *A&A*, 682, A9
- Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, *A&A*, 486, 951
- Haldemann, J., Ksoll, V., Walter, D., et al. 2023, *A&A*, 672, A180
- Heiter, U., Jofré, P., Gustafsson, B., et al. 2015, *A&A*, 582, A49
- Heiter, U., Lind, K., Bergemann, M., et al. 2021, *A&A*, 645, A106
- Hubeny, I., & Lanz, T. 2011, *Astrophysics Source Code Library* [record ascl:1109.022]
- Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. 2016, Doi: <https://ui.adsabs.harvard.edu/abs/2016arXiv161107004I>
- Jin, S., Trager, S. C., Dalton, G. B., et al. 2024, *MNRAS*, 530, 2688
- Jofré, P., Heiter, U., Soubiran, C., et al. 2014, *A&A*, 564, A133
- Kang, D. E., Pellegrini, E. W., Ardizzone, L., et al. 2022, *MNRAS*, 512, 617
- Kang, D. E., Klessen, R. S., Ksoll, V. F., et al. 2023a, *MNRAS*, 520, 4981
- Kang, D. E., Ksoll, V. F., Itrich, D., et al. 2023b, *A&A*, 674, A175
- Kang, D. E., Itrich, D., Ksoll, V. F., et al. 2025, *A&A*, 697, A39
- Kingma, D. P., & Dhariwal, P. 2018, in *Advances in Neural Information Processing Systems*, eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (USA: Curran Associates, Inc.), 31
- Kingma, D. P., & Welling, M. 2013, arXiv e-prints [arXiv:1312.6114]
- Kobyzev, I., Prince, S. J., & Brubaker, M. A. 2021, *IEEE Trans. Pattern Anal. Mach. Intell.*, 43, 3964
- Kollmeier, J. A., Rix, H.-W., Aerts, C., et al. 2026, *AJ*, 171, 52
- Korn, A. J., Shi, J., & Gehren, T. 2003, *A&A*, 407, 691
- Kovalev, M., Bergemann, M., Ting, Y.-S., & Rix, H.-W. 2019, *A&A*, 628, A54
- Ksoll, V. F., Ardizzone, L., Klessen, R., et al. 2020, *MNRAS*, 499, 5447
- Ksoll, V. F., Reissl, S., Klessen, R. S., et al. 2024, *A&A*, 683, A246
- Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., & Talwalkar, A. 2018, *J. Mach. Learn. Res.*, 18, 1
- Lind, K., & Amarsi, A. M. 2024, *ARA&A*, 62, 475
- Lind, K., Asplund, M., Barklem, P. S., & Belyaev, A. K. 2011, *A&A*, 528, A103
- Lind, K., Bergemann, M., & Asplund, M. 2012, *MNRAS*, 427, 50
- Liu, T., & Regier, J. 2020, arXiv e-prints [arXiv:2006.10175]
- Liu, W., Cao, S., Yu, X.-C., et al. 2024, *ApJS*, 271, 53
- Luck, R. E. 2017, *AJ*, 153, 21
- Magg, E., Bergemann, M., Serenelli, A., et al. 2022, *A&A*, 661, A140
- Majewski, S. R., APOGEE Team, & APOGEE-2 Team. 2016, *Astron. Nachr.*, 337, 863
- Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, 154, 94
- Manteiga, M., Santoveña, R., Álvarez, M. A., et al. 2025, *A&A*, 694, A326
- Mashonkina, L., Zhao, G., Gehren, T., et al. 2008, *A&A*, 478, 529
- Mashonkina, L., Jablonka, P., Sitnova, T., Pakhomov, Y., & North, P. 2017, *A&A*, 608, A89
- Mayor, M., Pepe, F., Queloz, D., et al. 2003, *The Messenger*, 114, 20
- Mirza, M., & Osindero, S. 2014, arXiv e-prints [arXiv:1411.1784]
- Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015, *ApJ*, 808, 16
- Ness, M., Zasowski, G., Johnson, J. A., et al. 2016, *ApJ*, 819, 2
- Pal, T., Khan, I., Worthey, G., Gregg, M. D., & Silva, D. R. 2023, *ApJS*, 266, 41
- Paszke, A., Gross, S., Chintala, S., et al. 2017, in *NIPS Autodiff Workshop*
- Plez, B. 2012, *Astrophysics Source Code Library* [record ascl:1205.004]
- Randich, S., Gilmore, G., & Gaia-ESO Consortium 2013, *The Messenger*, 154, 47
- Randich, S., Gilmore, G., Magrini, L., et al. 2022, *A&A*, 666, A121
- Reetz, J. K. 1991, Diploma Thesis, Universität München, Germany
- Rezende, D., & Mohamed, S. 2015, *Proc. Mach. Learn. Res.*, 37, 1530
- Rezende, D. J., Mohamed, S., & Wierstra, D. 2014, *Proc. Mach. Learn. Res.*, 32, 1278
- Ryabchikova, T., Piskunov, N., Kurucz, R. L., et al. 2015, *Phys. Scr.*, 90, 054005
- Santoveña, R., Dafonte, C., & Manteiga, M. 2024, arXiv e-prints [arXiv:2411.05960]
- Semenova, E., Bergemann, M., Deal, M., et al. 2020, *A&A*, 643, A164
- Snedden, C. A. 1973, PhD thesis, University of Texas, Austin
- Sohn, K., Lee, H., & Yan, X. 2015, in *Advances in Neural Information Processing Systems*, eds. C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, & R. Garnett (USA: Curran Associates, Inc.), 28
- Soubiran, C., Brouillet, N., & Casamiuela, L. 2022, *A&A*, 663, A4
- Soubiran, C., Creevey, O. L., Lagarde, N., et al. 2024, *A&A*, 682, A145
- Steinmetz, M., Guiglion, G., McMillan, P. J., et al. 2020, *AJ*, 160, 83
- Stonkutė, E., Chorniy, Y., Tautvaišienė, G., et al. 2020, *AJ*, 159, 90
- Storm, N., & Bergemann, M. 2023, *MNRAS*, 525, 3718
- Storm, N., Barklem, P. S., Yakovleva, S. A., et al. 2024, *A&A*, 683, A200
- Tabak, E., & Turner, C. 2013, *Commun. Pure Appl. Math.*, 66, 145
- Tabak, E., & Vanden-Eijnden, E. 2010, *Commun. Math. Sci.*, 8, 217
- Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2019, *ApJ*, 879, 69
- Voronov, Y. V., Yakovleva, S. A., & Belyaev, A. K. 2022, *ApJ*, 926, 173
- Wehrhahn, A., Piskunov, N., & Ryabchikova, T. 2023, *A&A*, 671, A171
- Wheeler, A. J., Abruzzo, M. W., Casey, A. R., & Ness, M. K. 2023, *AJ*, 165, 68
- Yakovleva, S. A., Belyaev, A. K., & Bergemann, M. 2020, *Atoms*, 8, 34
- Zhao, G., Zhao, Y.-H., Chu, Y.-Q., Jing, Y.-P., & Deng, L.-C. 2012, *Res. Astron. Astrophys.*, 12, 723

Appendix A: Tests on synthetic data

This appendix provides additional material related to the tests of the trained cINN on the held-out synthetic test spectra in Sect. 4.2. Supplemental to Fig. 4, Tables A.1 and A.2 and provide summaries of the $\bar{\Delta}_{\text{cINN-GT}} \pm \sigma_{\text{cINN-GT}}$ and NRMSE performance metrics of the cINN MAP estimates on the synthetic test spectra across all S/N values that the cINN is trained for. Figure A.1 further expands upon Fig. 4 by providing a breakdown of the mean residual $\bar{\Delta}_{\text{cINN-GT}}$ and standard deviation $\sigma_{\text{cINN-GT}}$ as a function of the S/N for different intervals of the respective target parameters. This figure shows that for most target parameters the trend of the mean residual is inverted between the lower third and upper third of the range. Specifically, there is a trend towards overestimation as the S/N decreases in the lower third, whereas in the upper third we observe a trend towards underestimation instead. Figure A.2 shows the median width $\Delta_{68\%}$ of the 68% confidence interval of the predicted posteriors for all of the target parameters of the cINN as a function of the S/N of the input spectrum. This figure demonstrates that the cINN learned to associate a larger uncertainty in the input fluxes with larger uncertainties in the predicted physical parameters. We also observe a notable increase in the posterior widths below $S/N = 80 \text{ \AA}^{-1}$, rendering, in particular, the elemental abundances a lot harder to constrain. We note that a network specifically trained for the lower S/N regime might improve upon the performance of the general purpose model that is presented here. Analogously to Fig. A.2, Fig. A.3 presents the median full width of the predicted posterior distributions as a function of [Fe/H] for a fixed $S/N = 250 \text{ \AA}^{-1}$, highlighting a notable increase in the cINN intrinsic uncertainties in particular in the $[\text{Fe}/\text{H}] < -2.5$ regime. Figures A.4 and A.5 shows a feature importance analysis on the trained network in the form of saliency maps for the target parameters. In particular, we computed the median network internal gradients $(\partial X/\partial f_{\lambda})_i$ of each target parameter, X , with respect to the flux, f_{λ} , in each wavelength bin over all synthetic test spectra with $S/N = 250 \text{ \AA}^{-1}$. To highlight the relative importance of each bin, the gradients $(\partial X/\partial f_{\lambda})_i$ are normalised in Figs. A.4 and A.5 to the magnitude of the total gradient $|\nabla_{\mathbf{f}} X|_i$ of X with respect to the entire flux vector \mathbf{f} for each test spectrum i prior to taking the average. Bins with a median gradient that is larger by more than 3σ than the mean median gradient are highlighted in black. This analysis, in particular when looking at the elemental abundances, demonstrates that the cINN learned to pay attention to sections of the spectra that are also part of classical evaluation techniques for characterising spectra, as the highest peaks in saliency almost always coincide with absorption lines of the corresponding elements.

Appendix B: Application to benchmark spectra

This appendix contains supplementary material regarding the application of the trained cINN to real spectra in Sect. 4.3. Table B.1 lists the set of sources, for which we analysed spectra in this application of our cINN, along with literature results of T_{eff} , $\log(g)$, and [Fe/H] (with the corresponding references). Table B.2 summarises the set of absorption lines used in Sect. 4.3 to determine the [Fe/H], [Ca/Fe], [Mg/Fe] and [Li/Fe] abundance ratios with TSPy for our set of real benchmark spectra. Table B.3 summarises the TSPy abundances determined for our set of real benchmark spectra as outlined in Section 4.3. It also lists the final cINN parameter estimates for all benchmark spectra after the bias correction that is detailed in the following Sect. B.1. Table B.4 provides an overview of the raw MAP estimates of the cINN for all 13 target parameters and the sample rejection fraction f_{reject} across the entire set of real spectra analysed in Sect. 4.3.

B.1. Bias correction

In our initial application to the set of observed spectra in Section 4.3, we identified a small bias in the predicted stellar parameters and abundances in dependence of the T_{eff} estimates. This is illustrated in Fig. B.1, which shows a 2D histogram of the residuals of the cINN predictions with respect to the literature values for T_{eff} and $\log(g)$, and to the TSPy outcome for [Fe/H], [Mg/Fe], [Ca/Fe] and [Li/Fe] as a function of the cINN estimates of T_{eff} for all benchmark spectra. For each star we plot all valid posterior samples in this comparison (i.e. all samples remaining after the sample rejection procedure outlined in Sect. 3.5). To account for this systematic effect, we performed a linear correction, following

$$\tilde{X}_{\text{cINN}} = X_{\text{cINN}} - (a_{X,\text{cal}} T_{\text{eff,cINN}} + b_{X,\text{cal}}), \quad (\text{B.1})$$

where $a_{X,\text{cal}}$ and $b_{X,\text{cal}}$ refer to the slope and intercept of the linear fit (indicated in red in Fig. B.1) for target parameter X , and $T_{\text{eff,cINN}}$ is the cINN estimate of the effective temperature. We note that, following our assessment of the optimal performance range in our evaluation on the synthetic data in Sect. 4.2, we excluded all stars from these fits that have a TSPy metallicity of $[\text{Fe}/\text{H}] < -2.5$. Table B.5 summarises the respective coefficients for all six target parameters considered in this analysis. We also note that we did not correct v_{broad} or v_{mic} , as reliable literature/ground truth values were not available for these parameters to compute the respective residuals.

B.2. Re-simulation

Figure B.2 provides additional results from our re-simulation experiment in Sect. 4.3, showing the entire re-simulated spectra of the example stars shown in Fig. 7 in comparison to the input observations. This diagram also indicates the position of the four example absorption lines that are the focus in Fig. 7, the H_{α} absorption line and the telluric absorption features that have not been corrected for in the input observations. Not surprisingly, some of the more apparent discrepancies visible in these plots correspond to these telluric features at around 6300 \AA (see Fig. B.2), as they are not being simulated in Turbospectrum.

Table A.1. Overview of the cINN predictive performance in terms of the mean and standard deviation, $\bar{\Delta}_{\text{cINN-GT}} \pm \sigma_{\text{cINN-GT}}$, of the residuals of the target parameters with respect to the ground truth on synthetic test data for different S/N.

S/N (\AA^{-1})	T_{eff} (K)	$\log\left(\frac{g}{\text{cm s}^{-2}}\right)$	[Fe/H]	v_{broad} (kms^{-1})	v_{mic} (kms^{-1})	Ca	[x/Fe] Mg	Li
1000	2.00 ± 32	0.003 ± 0.15	0.01 ± 0.11	0.05 ± 1.19	-0.029 ± 0.27	-0.006 ± 0.09	-0.005 ± 0.10	0.010 ± 0.49
250	1.46 ± 33	0.005 ± 0.16	0.01 ± 0.12	0.02 ± 1.27	-0.028 ± 0.27	-0.007 ± 0.10	-0.008 ± 0.11	0.018 ± 0.51
180	0.86 ± 36	0.005 ± 0.17	0.01 ± 0.12	0.03 ± 1.33	-0.029 ± 0.27	-0.007 ± 0.10	-0.007 ± 0.11	0.016 ± 0.52
130	1.07 ± 39	0.003 ± 0.19	0.01 ± 0.12	0.05 ± 1.52	-0.025 ± 0.28	-0.005 ± 0.11	-0.007 ± 0.12	0.014 ± 0.53
100	1.58 ± 44	-0.002 ± 0.21	0.01 ± 0.13	0.05 ± 1.65	-0.026 ± 0.29	-0.007 ± 0.12	-0.007 ± 0.13	0.012 ± 0.55
80	0.56 ± 49	-0.002 ± 0.23	0.01 ± 0.14	0.04 ± 1.83	-0.023 ± 0.29	-0.008 ± 0.12	-0.004 ± 0.14	0.011 ± 0.57
50	2.68 ± 62	-0.006 ± 0.29	0.01 ± 0.16	0.04 ± 2.64	-0.023 ± 0.31	-0.008 ± 0.14	-0.007 ± 0.17	-0.004 ± 0.63
45	2.37 ± 68	-0.009 ± 0.31	0.02 ± 0.18	0.03 ± 2.95	-0.020 ± 0.32	-0.009 ± 0.15	-0.007 ± 0.17	0.002 ± 0.64
40	4.23 ± 73	-0.010 ± 0.31	0.02 ± 0.18	0.06 ± 3.10	-0.023 ± 0.33	-0.012 ± 0.16	-0.008 ± 0.18	-0.002 ± 0.66
35	3.95 ± 81	-0.005 ± 0.34	0.02 ± 0.20	-0.04 ± 3.74	-0.026 ± 0.33	-0.009 ± 0.16	-0.010 ± 0.18	-0.013 ± 0.68
30	5.35 ± 90	-0.007 ± 0.38	0.02 ± 0.20	-0.02 ± 3.88	-0.030 ± 0.35	-0.015 ± 0.17	-0.013 ± 0.19	-0.008 ± 0.71
26	7.29 ± 100	0.008 ± 0.39	0.03 ± 0.23	-0.06 ± 4.41	-0.031 ± 0.36	-0.016 ± 0.18	-0.015 ± 0.20	-0.020 ± 0.74
23	7.88 ± 111	0.012 ± 0.42	0.04 ± 0.25	-0.07 ± 5.01	-0.035 ± 0.37	-0.017 ± 0.19	-0.014 ± 0.21	-0.031 ± 0.76
20	11.4 ± 127	0.018 ± 0.46	0.05 ± 0.28	-0.13 ± 5.93	-0.036 ± 0.38	-0.017 ± 0.19	-0.014 ± 0.21	-0.035 ± 0.79
18	12.4 ± 138	0.022 ± 0.48	0.06 ± 0.30	-0.10 ± 6.42	-0.036 ± 0.39	-0.018 ± 0.20	-0.014 ± 0.22	-0.057 ± 0.80
16	17.7 ± 149	0.023 ± 0.51	0.07 ± 0.33	0.06 ± 7.29	-0.041 ± 0.39	-0.016 ± 0.20	-0.015 ± 0.22	-0.051 ± 0.83
14	17.3 ± 172	0.014 ± 0.55	0.08 ± 0.36	-0.02 ± 7.59	-0.039 ± 0.40	-0.024 ± 0.21	-0.014 ± 0.23	-0.080 ± 0.85
12	24.1 ± 196	0.013 ± 0.60	0.11 ± 0.39	0.09 ± 8.28	-0.034 ± 0.42	-0.022 ± 0.21	-0.013 ± 0.23	-0.113 ± 0.89
10	17.5 ± 233	-0.002 ± 0.64	0.12 ± 0.46	0.23 ± 9.59	-0.017 ± 0.43	-0.026 ± 0.22	-0.014 ± 0.24	-0.163 ± 0.94

Table A.2. Overview of the cINN predictive performance in terms of the NRMSE of the target parameters with respect to the ground truth on synthetic test data for different S/N.

S/N (\AA^{-1})	T_{eff} (K)	$\log\left(\frac{g}{\text{cm s}^{-2}}\right)$	[Fe/H]	v_{broad} (kms^{-1})	v_{mic} (kms^{-1})	Ca	[x/Fe] Mg	Li
1000	0.008	0.028	0.02	0.012	0.18	0.071	0.066	0.054
250	0.008	0.031	0.02	0.013	0.182	0.076	0.071	0.056
180	0.009	0.033	0.021	0.013	0.184	0.079	0.075	0.057
130	0.009	0.035	0.022	0.015	0.188	0.084	0.081	0.059
100	0.011	0.039	0.023	0.016	0.191	0.09	0.087	0.06
80	0.012	0.043	0.025	0.018	0.195	0.094	0.095	0.063
50	0.015	0.054	0.029	0.026	0.207	0.111	0.11	0.069
45	0.016	0.058	0.031	0.029	0.212	0.115	0.115	0.071
40	0.018	0.059	0.031	0.031	0.217	0.12	0.119	0.073
35	0.019	0.064	0.035	0.037	0.223	0.126	0.123	0.075
30	0.022	0.07	0.036	0.039	0.231	0.132	0.129	0.078
26	0.024	0.073	0.04	0.044	0.239	0.138	0.134	0.082
23	0.027	0.079	0.044	0.05	0.244	0.144	0.138	0.083
20	0.031	0.086	0.05	0.059	0.252	0.148	0.141	0.086
18	0.033	0.091	0.053	0.064	0.258	0.151	0.146	0.089
16	0.036	0.096	0.058	0.073	0.264	0.155	0.146	0.091
14	0.041	0.102	0.064	0.076	0.27	0.162	0.151	0.094
12	0.047	0.112	0.071	0.083	0.278	0.165	0.156	0.099
10	0.056	0.12	0.083	0.096	0.287	0.17	0.161	0.105

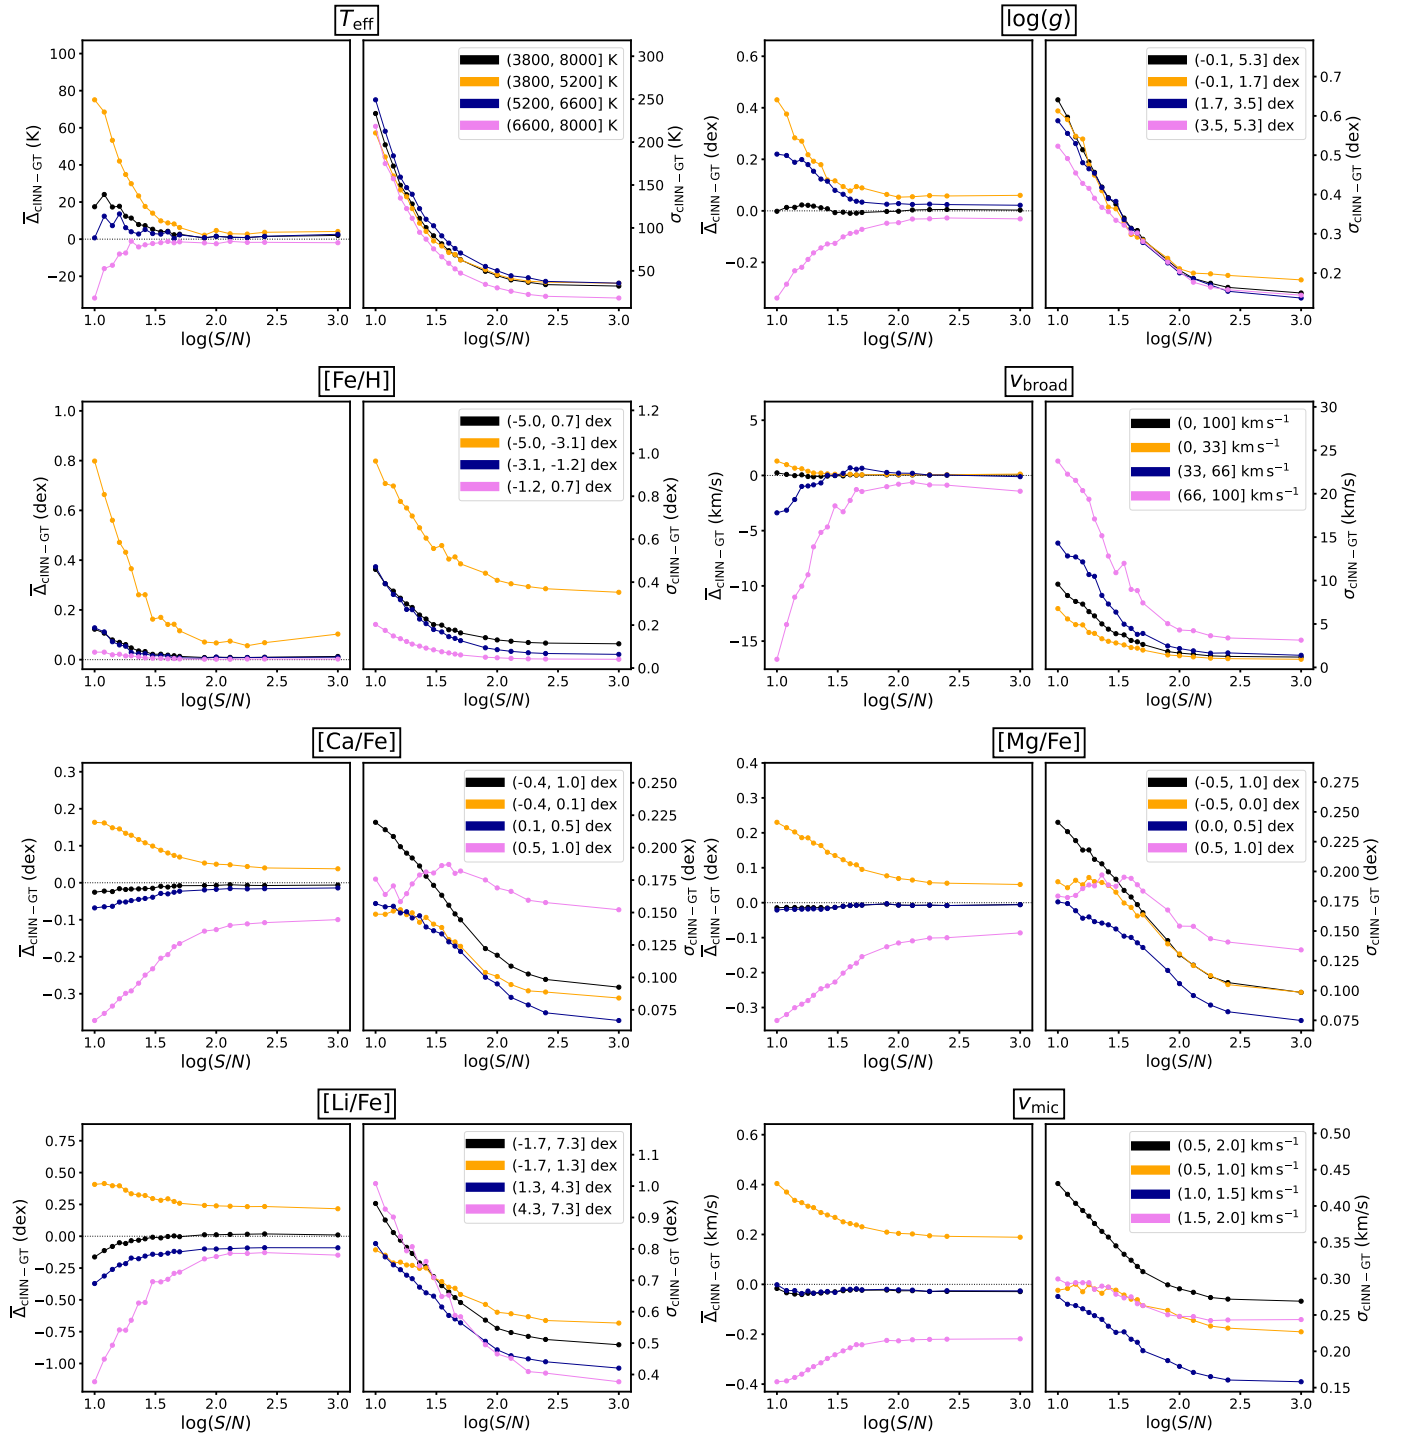


Fig. A.1. Summary of cINN performance of synthetic spectra as a function of S/N. Expanding upon Fig. 4, shown is a detailed breakdown of the mean residual $\bar{\Delta}_{\text{cINN-GT}}$ and standard deviation $\sigma_{\text{cINN-GT}}$ of each target parameter, distinguishing the lower third (orange), middle third (dark blue), and upper third (violet) of the respective range in the test data. For comparison, the black line shows the respective results over the full trained parameter range. The black dotted line indicates $\bar{\Delta}_{\text{cINN-GT}} = 0$ for reference.

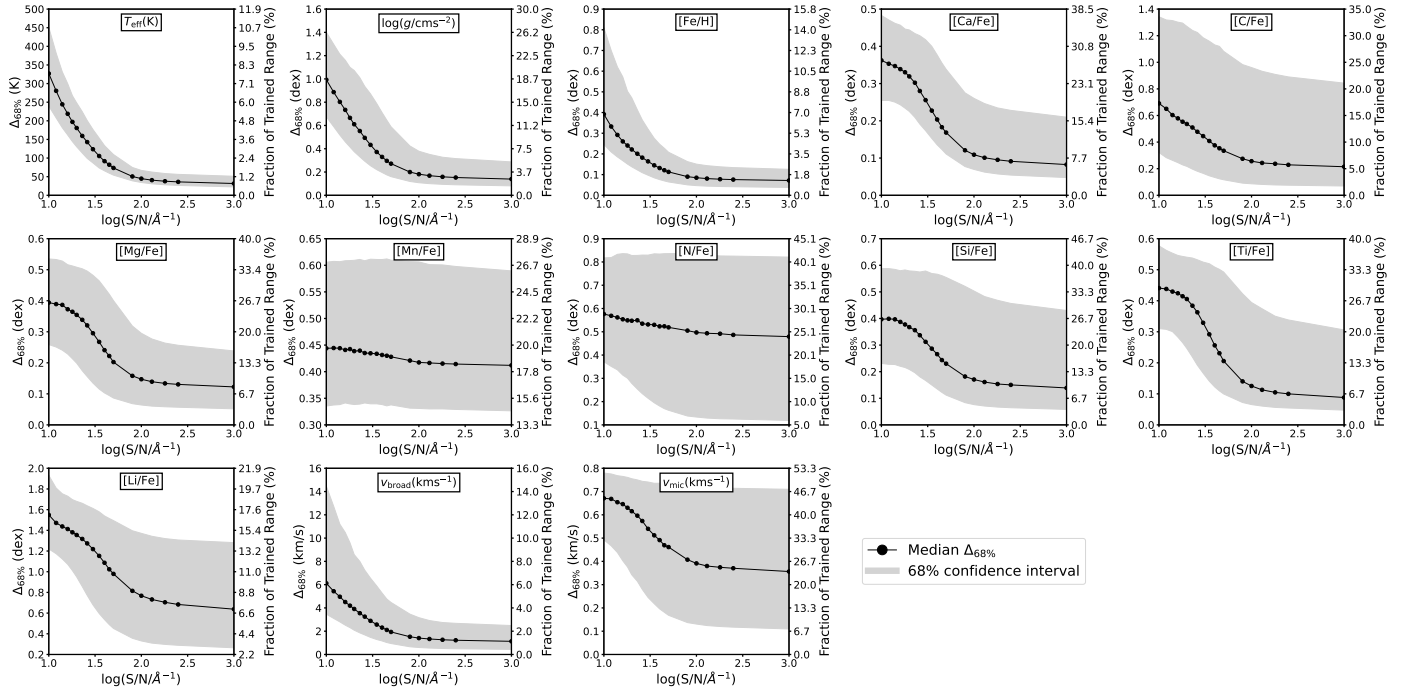


Fig. A.2. Overview of the median width of the 68% confidence interval for the individual components of the predicted posterior distributions as a function of the S/N. The grey-shaded region indicates the 68% confidence interval.

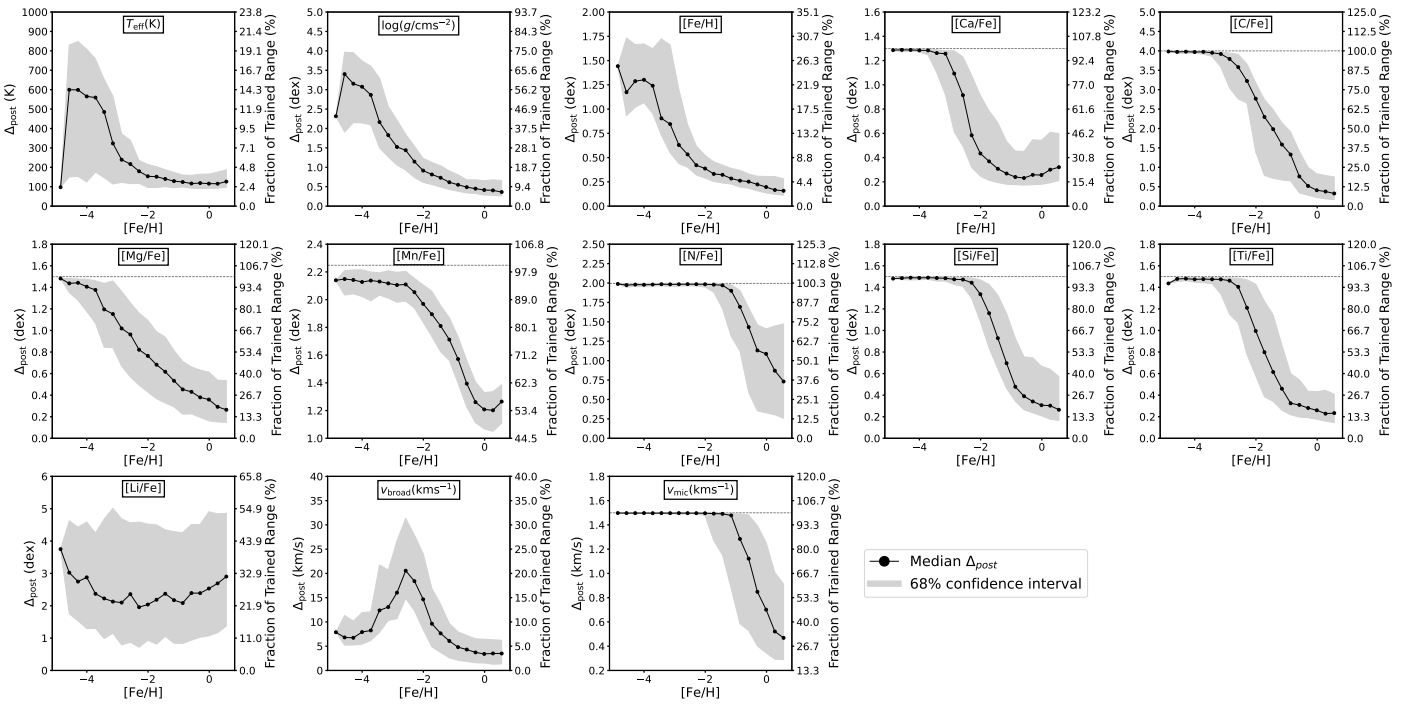


Fig. A.3. Overview of the median total width of the predicted posteriors for the individual components as a function of $[\text{Fe}/\text{H}]$ for $S/N = 250 \text{ \AA}^{-1}$ on the synthetic test spectra. The grey shaded region indicates the 68% confidence interval. The horizontal dashed line, indicates a full posterior width that corresponds to 100% of the parameter range in the training set.

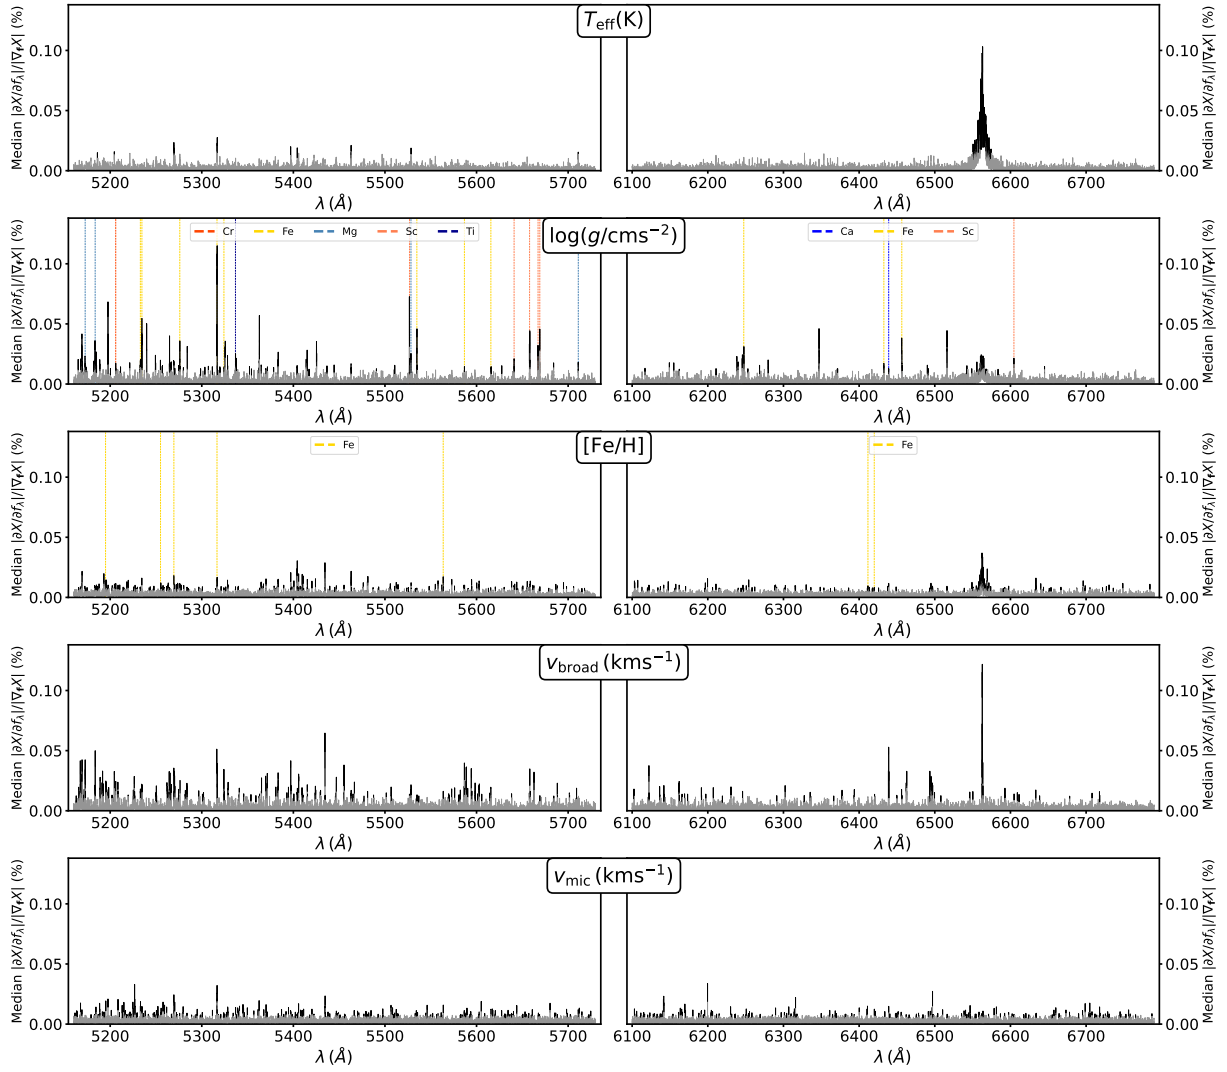


Fig. A.4. Average feature importance in the cINN for the primary stellar properties. Shown is the median relative contribution of each flux bin to the total gradient of each target parameter with respect to the total flux vector averaged over a synthetic test set with $S/N = 250 \text{ \AA}^{-1}$. Black lines indicate flux bins whose median contribution is larger than the mean by more than 3σ . In each row, the left panel shows the green 4MOST window, while the right panel presents the red window. The dashed lines in the $\log(g)$ and $[\text{Fe}/\text{H}]$ panels indicate common diagnostic absorption lines of various elements that coincide with a peak in the feature importance.

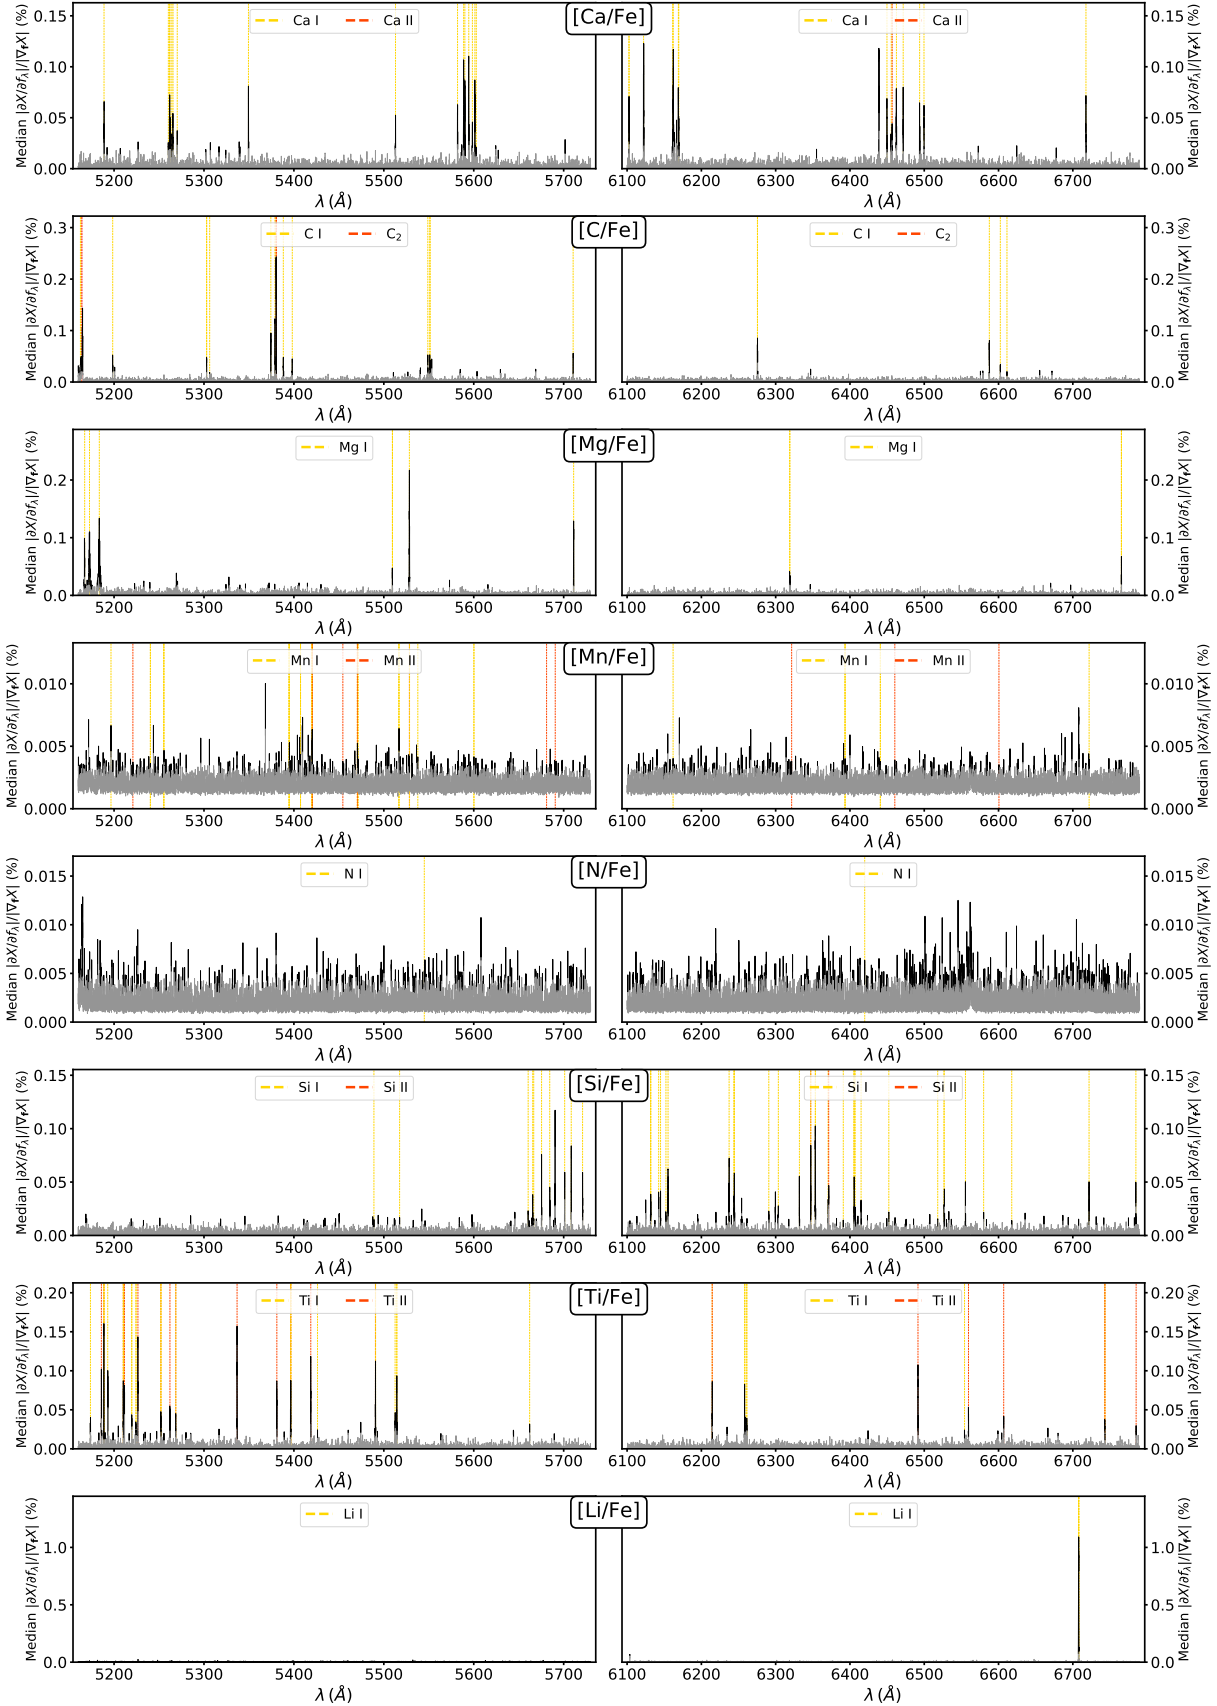


Fig. A.5. Average feature importance in the cINN for metal abundances. Shown is the median relative contribution of each flux bin to the total gradient of each target parameter with respect to the total flux vector averaged over a synthetic test set with $S/N = 250 \text{ \AA}^{-1}$. Black lines indicate flux bins whose median contribution is larger than the mean by more than 3σ . In each row, the left panel shows the green 4MOST window, while the right panel presents the red window. The dashed lines indicate absorption lines of the respective elements that are included in the Turbospectrum models and coincide with peaks in the feature importance.

Table B.1. Literature stellar properties of the set of sources used for validation.

Star	SpType	T_{eff} (K)	$\log(g/\text{cm s}^{-2})$	[Fe/H]	Reference
15PEG	F Dwarf (V)*	5728 ± 22	4.28 ± 0.03	0.13 ± 0.01	García Pérez et al. (2021)
18Sco	G Dwarf (V) [‡]	5824 ± 30	4.42 ± 0.01	0.06 ± 0.01	Soubiran et al. (2024); Casamiquela et al. (2026)
31AQL	G Subgiant (IV)*	5776 ± 3	4.28 ± 0.01	0.59 ± 0.00	García Pérez et al. (2021)
61CygA	K Dwarf (V) [‡]	4398 ± 34	4.63 ± 0.01	-0.32 ± 0.14	Soubiran et al. (2024); Casamiquela et al. (2026)
61CygB	K Dwarf (V) [‡]	4174 ± 47	4.68 ± 0.02	-0.59 ± 0.21	Soubiran et al. (2024); Casamiquela et al. (2026)
Arcturus	K Giant (III) [‡]	4277 ± 23	1.58 ± 0.07	-0.55 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
BD+26_3578	B Subgiant (IV) [†]	6432 ± 2	3.75 ± 0.01	-2.02 ± 0.02	García Pérez et al. (2021)
BD-04_3208	A Subgiant (IV) [†]	6223 ± 50	3.66 ± 0.15	-2.34 ± 0.05	Soubiran et al. (2022)
G113-22	G Subgiant (IV) [†]	5595 ± 46	3.70 ± 0.37	-1.07 ± 0.09	Soubiran et al. (2022)
G165-39	F Subdwarf (VI)*	6333 ± 2	3.82 ± 0.01	-1.90 ± 0.02	García Pérez et al. (2021)
G207-5	F Dwarf (V) [‡]	6003 ± 2	4.18 ± 0.01	-0.30 ± 0.01	García Pérez et al. (2021)
G28-43	G Subdwarf (VI)*	4923 ± 34	4.58 ± 0.07	-1.56 ± 0.02	Abdurro'uf et al. (2022)
G41-41	A Subdwarf (VI)*	6455 ± 6	4.07 ± 0.03	-3.09 ± 0.66	García Pérez et al. (2021)
G48-29	A Subdwarf (VI)*	6279 ± 111	3.99 ± 0.27	-2.52 ± 0.24	Soubiran et al. (2022)
G63-46	F Dwarf (V)*	5719 ± 42	3.89 ± 0.17	-0.89 ± 0.07	Soubiran et al. (2022)
G64-12	F Subdwarf (VI)*	6368 ± 40	4.11 ± 0.07	-3.26 ± 0.05	Soubiran et al. (2022)
G78-1	F Dwarf (V) [‡]	5995 ± 3	4.11 ± 0.02	-0.66 ± 0.02	García Pérez et al. (2021)
G9-16	F Dwarf (V)*	6858 ± -	4.28 ± -	-1.20 ± -	Pal et al. (2023)
Gmb1830	G Dwarf (V) [‡]	5235 ± 18	4.72 ± 0.01	-1.20 ± 0.06	Soubiran et al. (2024); Casamiquela et al. (2026)
HD101177	G Dwarf (V)*	5887 ± 70	4.43 ± 0.07	-0.26 ± 0.05	Soubiran et al. (2022)
HD102158	G Dwarf (V) [†]	5761 ± 20	4.32 ± 0.04	-0.43 ± 0.02	Soubiran et al. (2022)
HD103095	K Dwarf (V)*	4991 ± 30	4.60 ± 0.05	-1.26 ± 0.02	Abdurro'uf et al. (2022)
HD104304	G Subgiant (IV)*	5441 ± 6	4.18 ± 0.01	0.32 ± 0.01	García Pérez et al. (2021)
HD105631	G Dwarf (V)*	5402 ± 19	4.49 ± 0.04	0.18 ± 0.02	Soubiran et al. (2022)
HD106210	G Dwarf (V)*	5664 ± 20	4.41 ± 0.05	-0.10 ± 0.04	Soubiran et al. (2022)
HD106516	F Dwarf (V)*	6274 ± 4	4.47 ± 0.03	-0.82 ± 0.02	García Pérez et al. (2021)
HD10697	G Dwarf (V)*	5610 ± 39	3.87 ± 0.03	0.16 ± 0.01	Abdurro'uf et al. (2022)
HD10700	G Dwarf (V)*	5374 ± 2	4.35 ± 0.01	-0.19 ± 0.01	García Pérez et al. (2021)
HD107328	K Giant (III) [‡]	4576 ± 87	1.77 ± 0.22	-0.36 ± 0.06	Soubiran et al. (2024); Casamiquela et al. (2026)
HD110897	F Dwarf (V)*	5872 ± 3	4.25 ± 0.02	-0.62 ± 0.01	García Pérez et al. (2021)
HD111395	G Dwarf (V)*	6279 ± 15	3.56 ± 0.03	-0.09 ± 0.02	Steinmetz et al. (2020)
HD113139	F Dwarf (V)*	6800 ± -	4.10 ± -	-0.20 ± -	Gebran et al. (2016)
HD114174	G Subgiant (IV)*	5854 ± -	4.53 ± -	-0.02 ± -	García Pérez et al. (2021)
HD122563	K Bright Giant (II) [‡]	4642 ± 35	1.32 ± 0.03	-2.66 ± 0.05	Soubiran et al. (2024); Casamiquela et al. (2026)
HD140283	G Subgiant (IV) [‡]	5788 ± 45	3.75 ± 0.09	-2.41 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
HD144284	F Dwarf (V)*	6208 ± 48	3.79 ± -	0.19 ± -	Luck (2017)
HD162003	F Subgiant/Dwarf (IV/V)*	6567 ± 86	4.06 ± 0.32	-0.02 ± 0.11	Stonkutė et al. (2020)
HD164259	F Dwarf (V)*	6815 ± 63	4.11 ± 0.02	-0.10 ± 0.05	Aguilera-Gómez et al. (2018)
HD220009	K Giant (III) [‡]	4227 ± 77	1.66 ± 0.04	-0.71 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
HD22879	G Dwarf (V) [‡]	5962 ± 86	4.28 ± 0.04	-0.78 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)
HD284248	G Subdwarf (VI)*	6253 ± 3	4.04 ± 0.02	-1.52 ± 0.03	García Pérez et al. (2021)
HD84937	F Dwarf (V) [‡]	6484 ± 106	4.16 ± 0.05	-1.97 ± 0.07	Soubiran et al. (2024); Casamiquela et al. (2026)
OMI_AQL	F Dwarf (V)*	6249 ± 3	4.04 ± 0.01	0.28 ± 0.01	García Pérez et al. (2021)
Procyon	F Subgiant (IV) [‡]	6582 ± 5	3.98 ± 0.02	-0.03 ± 0.02	Soubiran et al. (2024); Casamiquela et al. (2026)
Sun	G Dwarf (V) [‡]	5771 ± 10	4.44 ± 0.01	0.00 ± 0.00	Soubiran et al. (2024); Casamiquela et al. (2026)
α CenA	G Dwarf (V) [‡]	5804 ± 13	4.29 ± 0.01	0.22 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)
α Cet	K Bright Giant (II) [‡]	3738 ± 170	0.66 ± 0.07	-0.57 ± 0.17	Soubiran et al. (2024); Casamiquela et al. (2026)
α Tau	K Bright Giant (II) [‡]	3921 ± 80	1.17 ± 0.05	-0.24 ± 0.11	Soubiran et al. (2024); Casamiquela et al. (2026)
β Hyi	G Subgiant (IV) [‡]	5917 ± 25	3.97 ± 0.04	-0.06 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)
β Vir	F Dwarf (V) [‡]	6093 ± 13	4.08 ± 0.02	0.13 ± 0.02	Soubiran et al. (2024); Casamiquela et al. (2026)
δ Eri	K Subgiant (IV) [‡]	4954 ± 38	3.79 ± 0.10	0.04 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
ϵ Eri	K Dwarf (V) [‡]	5062 ± 29	4.60 ± 0.01	-0.10 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)
ϵ Vir	K Giant (III) [‡]	4950 ± 10	2.72 ± 0.02	-0.03 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)
η Boo	F Subgiant (IV) [‡]	6161 ± 18	3.82 ± 0.05	0.32 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
μ Ara	G Dwarf (V) [‡]	5974 ± 61	4.30 ± 0.03	0.41 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
μ Cas	G Dwarf (V) [‡]	5358 ± 31	4.49 ± 0.01	-0.81 ± 0.04	Soubiran et al. (2024); Casamiquela et al. (2026)
μ Leo	K Giant (III) [‡]	4519 ± 23	2.43 ± 0.06	0.19 ± 0.06	Soubiran et al. (2024); Casamiquela et al. (2026)
τ Cet	G Dwarf (V) [‡]	5397 ± 15	4.50 ± 0.01	-0.47 ± 0.03	Soubiran et al. (2024); Casamiquela et al. (2026)

* From <https://simbad.cds.unistra.fr/simbad/>.† Spectral type from <https://simbad.cds.unistra.fr/simbad/>, luminosity class estimated based on $\log(g)$.‡ From <https://www.blancocuaresma.com/s/benchmarkstars>, i.e. Soubiran et al. (2024); Casamiquela et al. (2026).

Table B.2. List of absorption lines used to fit abundances with TFitPy.

Element	λ [Å]	E_{low} [eV]	$\log(gf)$	g_{up}	Isotope	Element	λ [Å]	E_{low} [eV]	$\log(gf)$	g_{up}	Isotope
Ca I	5188.844	2.933	-0.075	5		Fe I	5560.211	4.435	-1.090	7	
Ca I	5349.465	2.709	-0.310	7		Fe I	5638.262	4.220	-0.720	7	
Ca I	5581.965	2.523	-0.555	7		Fe I	5661.345	4.284	-1.756	3	
Ca I	5588.749	2.526	0.358	7		Fe I	5679.023	4.652	-0.820	7	
Ca I	5590.114	2.521	-0.571	5		Fe I	6165.360	4.143	-1.473	9	
Ca I	5594.462	2.523	0.097	5		Fe I	6187.989	3.943	-1.620	9	
Ca I	5601.277	2.526	-0.523	5		Fe I	6270.224	2.858	-2.470	3	
Ca I	6102.723	1.879	-0.850	3		Fe II	5234.623	3.221	-2.180	6	
Ca I	6122.217	1.886	-0.380	3		Fe II	5325.552	3.221	-3.160	8	
Ca I	6162.173	1.899	-0.170	3		Fe II	5425.249	3.199	-3.220	10	
Ca I	6169.563	2.526	-0.478	5		Fe II	6456.379	3.903	-2.185	6	
Ca I	6439.075	2.526	0.390	9		Li I	6707.764	0.000	-0.002	4	7
Ca I	6449.808	2.521	-0.502	5		Li I	6707.915	0.000	-0.303	2	7
Ca I	6462.567	2.523	0.262	7		Li I	6707.921	0.000	-0.002	4	6
Ca I	6493.781	2.521	-0.109	5		Li I	6708.072	0.000	-0.303	2	6
Ca I	6717.681	2.709	-0.524	3		Mg I	5167.322	2.709	-0.931	3	
Fe I	5242.491	3.634	-0.967	11		Mg I	5172.684	2.712	-0.363	3	
Fe I	5365.399	3.573	-1.020	9		Mg I	5183.604	2.717	-0.168	3	
Fe I	5379.574	3.695	-1.514	11		Mg I	5528.405	4.346	-0.547	5	
Fe I	5398.279	4.446	-0.630	5		Mg I	5711.088	4.346	-1.742	1	
Fe I	5543.936	4.218	-1.040	5							

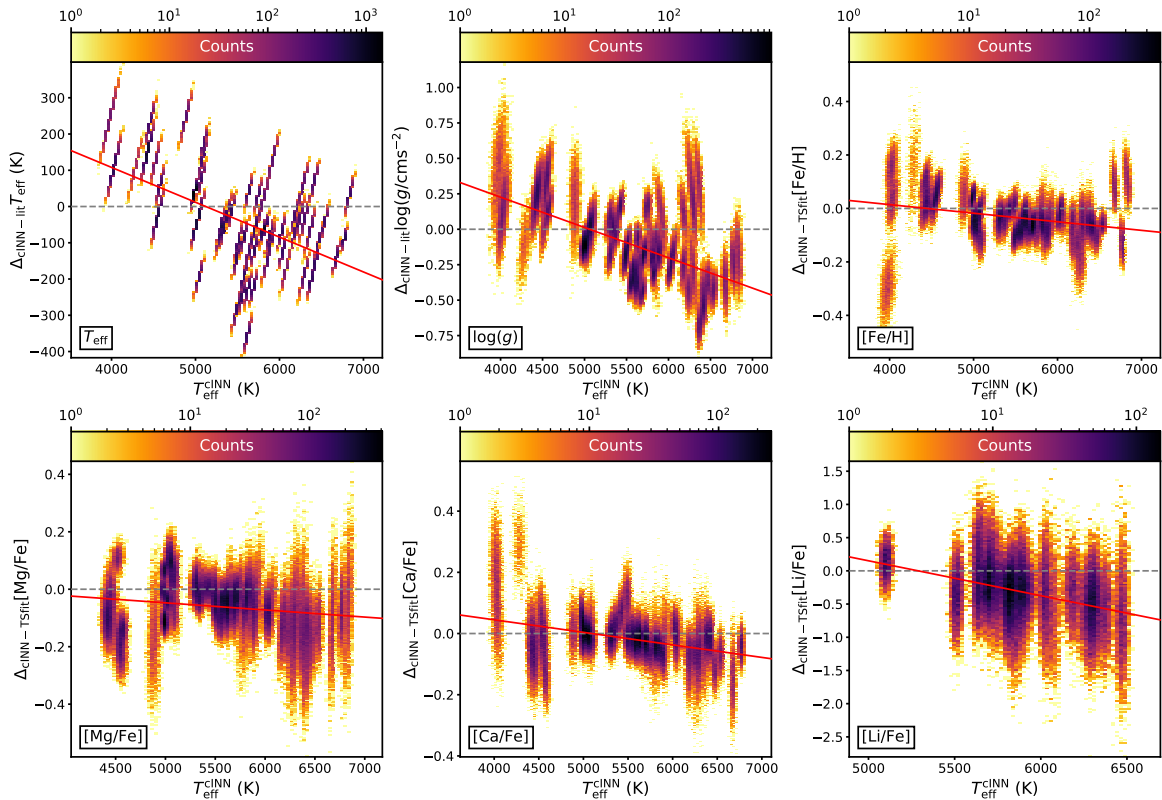
**Fig. B.1.** 2D histogram of the residuals of the cINN posterior predictions with respect to the literature/TSFitPy results as a function of the estimates for T_{eff} . The red lines indicate linear fits to the residuals used to correct for the apparent systematic bias in the cINN prediction.

Table B.3. Overview of cINN MAP predictions and TSFitPy abundances for the benchmark spectra.

Star (Spectrum)	Final cINN prediction							Turbospectrum fit*			
	T_{eff} (K)	$\log\left(\frac{g}{\text{cm}^2}\right)$	[Fe/H]	[Ca/Fe]	[Mg/Fe]	[Li/Fe]	f_{reject}	[Fe/H]	[Ca/Fe]	[Mg/Fe]	[Li/Fe]
15PEG (FOCES)	6539	3.78	-0.58	0.13	0.3	1.61	0.13	-0.57 ± 0.16	0.08 ± 0.13	0.44 ± 0.05	-
18Sco (NARVAL)	5749	4.36	-0.1	0.12	0.09	0.68	0	-0.09 ± 0.07	0.10 ± 0.04	0.06 ± 0.01	0.62 ± 0.00
31AQL (FOCES)	5523	4.2	0.23	0.06	0.12	-0.01	0	0.19 ± 0.07	-0.07 ± 0.09	0.03 ± 0.03	-
61CygA (NARVAL)	4466	4.54	-0.39	0.15	0.27	-0.2	0.5	-0.45 ± 0.13	0.13 ± 0.06	0.12 ± 0.15	-
61CygB (NARVAL)	4208	4.39	-0.6	0.39	0.13	-0.51	0.72	-0.68 ± 0.01	0.07 ± 0.08	-	-
Arcturus (UVES)	4369	1.7	-0.62	0.09	0.44	-0.39	0.01	-0.68 ± 0.07	0.19 ± 0.08	0.50 ± 0.00	-
Arcturus (NARVAL)	4347	1.63	-0.65	0.14	0.43	-0.29	0	-0.73 ± 0.06	0.19 ± 0.11	0.48 ± 0.00	-
Arcturus (UVES)	4381	1.76	-0.61	0.05	0.37	-0.17	0.01	-0.63 ± 0.07	0.12 ± 0.15	0.45 ± 0.03	-
BD+26_3578 (FOCES)	6500	4.32	-2.23	0.4	0.33	3.69	0.22	-2.26 ± 0.00	0.40 ± 0.07	0.35 ± 0.01	3.55 ± 0.00
BD-04_3208 (FOCES)	6385	4.32	-2.32	0.33	0.38	3.57	0.16	-2.21 ± 0.08	0.40 ± 0.07	0.37 ± 0.03	3.53 ± 0.00
G113-22 (FOCES)	5692	3.77	-1.14	0.27	0.47	2.03	0.1	-1.11 ± 0.11	0.26 ± 0.11	0.44 ± 0.04	2.01 ± 0.00
G165-39 (FOCES)	6418	4.17	-2	0.37	0.28	3.18	0.17	-2.01 ± 0.07	0.37 ± 0.06	0.30 ± 0.02	3.23 ± 0.00
G207-5 (FOCES)	5904	4.28	-0.5	0.06	0.18	1.49	0.36	-0.46 ± 0.08	0.10 ± 0.08	0.16 ± 0.02	1.67 ± 0.00
G28-43 (FOCES)	5096	4.64	-1.7	0.32	0.37	1.9	0.35	-1.58 ± 0.05	0.33 ± 0.08	0.28 ± 0.08	1.81 ± 0.00
G41-41 (FOCES)	6517	4.21	-2.73	0.19	0.31	4.07	0.22	-2.71 ± 0.00	0.39 ± 0.10	0.22 ± 0.05	3.91 ± 0.00
G48-29 (FOCES)	6517	4.72	-2.83	0.29	0.37	3.96	0.3	-2.50 ± 0.00	0.49 ± 0.02	0.39 ± 0.07	3.98 ± 0.00
G63-46 (FOCES)	5801	4.19	-0.79	0.25	0.43	1.04	0.23	-0.77 ± 0.07	0.28 ± 0.10	0.40 ± 0.03	1.09 ± 0.00
G64-12 (FOCES)	6403	4.93	-3.71	0.29	0.72	5.14	0.3	-3.20 ± 0.00	0.98 ± 0.13	0.68 ± 0.01	-
G78-1 (FOCES)	5978	4.12	-0.92	0.25	0.46	1.98	0.29	-0.89 ± 0.10	0.28 ± 0.13	0.47 ± 0.02	1.81 ± 0.00
G9-16 (FOCES)	6922	4.47	-1.41	0.4	0.41	2.96	0.13	-1.37 ± 0.05	0.37 ± 0.07	0.45 ± 0.04	-
Gmb1830 (NARVAL)	5037	4.58	-1.41	0.35	0.42	1.4	0.18	-1.37 ± 0.08	0.32 ± 0.03	0.24 ± 0.06	-
HD101177 (FOCES)	5897	4.27	-0.31	0.08	0.18	1.46	0.02	-0.28 ± 0.08	0.09 ± 0.08	0.11 ± 0.02	1.43 ± 0.00
HD102158 (FOCES)	5701	4.02	-0.5	0.28	0.42	1.15	0.28	-0.50 ± 0.06	0.27 ± 0.10	0.43 ± 0.02	0.54 ± 0.00
HD103095 (FOCES)	5076	4.64	-1.41	0.33	0.41	1.37	0.23	-1.34 ± 0.06	0.32 ± 0.05	0.32 ± 0.08	-
HD104304 (FOCES)	5467	4.48	0.07	0.09	0.12	0.11	0	0.06 ± 0.06	-0.02 ± 0.10	0.04 ± 0.02	-
HD105631 (FOCES)	5346	4.55	-0.05	0.1	0.09	0.38	0	-0.04 ± 0.06	0.08 ± 0.09	0.06 ± 0.03	-
HD106210 (FOCES)	5627	4.28	-0.26	0.15	0.3	0.96	0	-0.26 ± 0.07	0.17 ± 0.06	0.27 ± 0.02	-
HD106516 (FOCES)	6368	4.27	-0.75	0.23	0.56	1.99	0.24	-0.68 ± 0.09	0.24 ± 0.10	0.54 ± 0.03	-
HD10697 (FOCES)	5555	3.86	-0.03	0.06	0.06	0.95	0	0.02 ± 0.07	0.08 ± 0.07	0.09 ± 0.01	0.84 ± 0.00
HD10700 (FOCES)	5298	4.45	-0.64	0.24	0.41	0.86	0.05	-0.62 ± 0.06	0.26 ± 0.06	0.34 ± 0.05	-
HD107328 (NARVAL)	4485	1.95	-0.55	0.08	0.38	-0.27	0	-0.57 ± 0.06	0.10 ± 0.10	0.50 ± 0.04	-
HD107328 (UVES)	4508	2.05	-0.53	0.06	0.29	-0.14	0	-0.56 ± 0.05	0.08 ± 0.08	0.50 ± 0.12	-
HD110897 (FOCES)	5855	4.18	-0.64	0.11	0.27	1.21	0.04	-0.62 ± 0.09	0.11 ± 0.09	0.25 ± 0.03	1.52 ± 0.00
HD111395 (FOCES)	5525	4.38	-0.08	0.12	0.07	0.46	0	-0.02 ± 0.05	0.12 ± 0.08	0.07 ± 0.01	-
HD113139 (FOCES)	6980	4.28	0.06	0.14	0.03	1.71	0.08	-0.15 ± 0.00	-	0.05 ± 0.00	-
HD114174 (FOCES)	5676	4.27	-0.1	0.11	0.16	0.29	0	-0.10 ± 0.06	0.09 ± 0.06	0.13 ± 0.01	-
HD122563 (NARVAL)	4847	1.47	-2.32	0.09	0.14	1.91	0.22	-2.39 ± 0.06	0.11 ± 0.08	0.24 ± 0.08	-
HD140283 (UVES)	6000	4.03	-2.31	0.1	0.26	3.62	0.21	-2.33 ± 0.06	0.10 ± 0.05	0.19 ± 0.03	3.62 ± 0.00
HD140283 (NARVAL)	5876	3.74	-2.38	0.11	0.25	3.7	0.2	-2.44 ± 0.07	0.14 ± 0.04	0.23 ± 0.04	3.60 ± 0.00
HD144284 (FOCES)	6280	3.76	0.11	0.03	0.06	1.19	0.16	0.10 ± 0.07	0.14 ± 0.00	0.13 ± 0.02	-
HD162003 (FOCES)	6608	3.99	-0.08	0.1	0.08	1.69	0.17	-0.04 ± 0.05	0.15 ± 0.15	0.18 ± 0.03	1.79 ± 0.00
HD164259 (FOCES)	6820	4.03	0.01	0.08	0.11	1.24	0.1	-0.16 ± 0.00	0.20 ± 0.00	0.12 ± 0.00	-
HD220009 (NARVAL)	4389	1.83	-0.72	0.16	0.4	-0.21	0.01	-0.74 ± 0.08	0.25 ± 0.07	0.41 ± 0.00	-
HD22879 (NARVAL)	5939	4.27	-0.93	0.27	0.48	1.79	0.15	-0.88 ± 0.09	0.28 ± 0.08	0.44 ± 0.03	1.54 ± 0.00
HD284248 (FOCES)	6293	4.41	-1.61	0.29	0.28	3.22	0.28	-1.57 ± 0.05	0.30 ± 0.09	0.31 ± 0.04	2.99 ± 0.00
HD84937 (UVES)	6397	4.12	-2.12	0.36	0.25	3.2	0.24	-2.18 ± 0.08	0.35 ± 0.07	0.36 ± 0.04	3.36 ± 0.00
OMI_AQL (FOCES)	6162	4.12	0.15	0.09	0.06	1.42	0.6	0.12 ± 0.07	0.06 ± 0.11	0.06 ± 0.02	1.49 ± 0.00
Procyon (NARVAL)	6695	3.9	-0.06	0.05	0.16	1.22	0.23	-0.06 ± 0.08	0.06 ± 0.08	0.21 ± 0.06	-
Procyon (UVES)	6498	3.64	-0.14	0.07	0.16	1.37	0.06	-0.16 ± 0.07	0.01 ± 0.13	0.26 ± 0.04	-
Procyon (UVES)	6545	3.73	-0.13	0.07	0.21	1.27	0.12	-0.14 ± 0.06	0.02 ± 0.15	0.23 ± 0.05	-
Sun (NARVAL)	5746	4.33	-0.03	0.11	0.05	0.32	0.66	-0.00 ± 0.08	0.12 ± 0.04	0.04 ± 0.04	-
Sun (UVES)	5620	4.17	-0.21	0.13	0.15	0.71	0	-0.16 ± 0.04	0.13 ± 0.05	0.13 ± 0.02	-
Sun (NARVAL)	5768	4.43	-0.15	0.1	0.13	0.79	0	-0.13 ± 0.06	0.09 ± 0.05	0.08 ± 0.02	-
α CenA (UVES)	5590	4.01	-0.02	0.13	0.13	0.14	0	0.03 ± 0.06	0.10 ± 0.12	0.17 ± 0.05	-
α Cen (NARVAL)	3867	0.71	-0.52	0.23	0.05	-1.19	0.11	-0.23 ± 0.00	-	-	-
α Cen (UVES)	3791	0.99	-0.5	0.13	0.35	-1.05	0.05	-	-	-	-
α Tau (NARVAL)	3915	0.99	-0.4	0.1	0.25	-1.2	0	-0.42 ± 0.12	0.10 ± 0.02	-	-
α Tau (UVES)	3915	0.99	-0.44	-0.01	0.13	-1.16	0	-0.58 ± 0.19	-0.17 ± 0.02	-	-
β Hyi (UVES)	5750	3.65	-0.26	0.1	0.19	1.7	0.03	-0.23 ± 0.08	0.11 ± 0.10	0.19 ± 0.01	1.66 ± 0.00
β Vir (NARVAL)	6100	3.96	0.05	0.07	0.11	1.02	0.03	0.03 ± 0.08	0.01 ± 0.14	0.10 ± 0.03	0.90 ± 0.00
δ Eri (UVES)	4981	3.63	-0.12	0.09	0.08	0.16	0	-0.17 ± 0.08	0.06 ± 0.04	0.13 ± 0.02	-
δ Eri (NARVAL)	4994	3.64	-0.1	0.11	0.08	0.15	0	-0.10 ± 0.06	0.08 ± 0.07	0.16 ± 0.03	-
δ Eri (UVES)	4981	3.7	-0.12	0.09	0.16	0.22	0	-0.16 ± 0.07	0.03 ± 0.08	0.11 ± 0.02	-
ϵ Eri (UVES)	5017	4.48	-0.28	0.13	0.08	0.08	0	-0.21 ± 0.07	0.14 ± 0.06	0.07 ± 0.05	-
ϵ Eri (UVES)	5047	4.61	-0.25	0.12	0.12	0.23	0	-0.24 ± 0.08	0.12 ± 0.08	-0.03 ± 0.00	-
ϵ Eri (UVES)	4986	4.48	-0.28	0.1	0.12	0.1	0	-0.25 ± 0.07	0.14 ± 0.07	-0.00 ± 0.03	-
ϵ Vir (NARVAL)	5081	2.75	-0	-0	-0.02	-0.06	0.04	-0.03 ± 0.06	0.01 ± 0.06	-0.02 ± 0.02	-
η Boo (NARVAL)	6133	3.82	0.22	0.03	0.08	0.74	0.07	0.15 ± 0.05	0.00 ± 0.12	0.09 ± 0.02	-
μ Ara (UVES)	5654	4.11	0.09	0.11	0.08	0.38	0	0.08 ± 0.06	0.09 ± 0.06	0.10 ± 0.02	-
μ Cas (NARVAL)	5328	4.39	-0.9	0.35	0.47	1.16	0.1	-0.89 ± 0.08	0.34 ± 0.05	0.40 ± 0.02	-
μ Leo (NARVAL)	4522	2.45	0.04	0.03	0.03	-0.65	0	-0.03 ± 0.07	0.19 ± 0.06	0.14 ± 0.06	-
τ Cet (NARVAL)	5404	4.53	-0.57	0.27	0.35	0.92	0.07	-0.57 ± 0.07	0.23 ± 0.04	0.31 ± 0.03	-

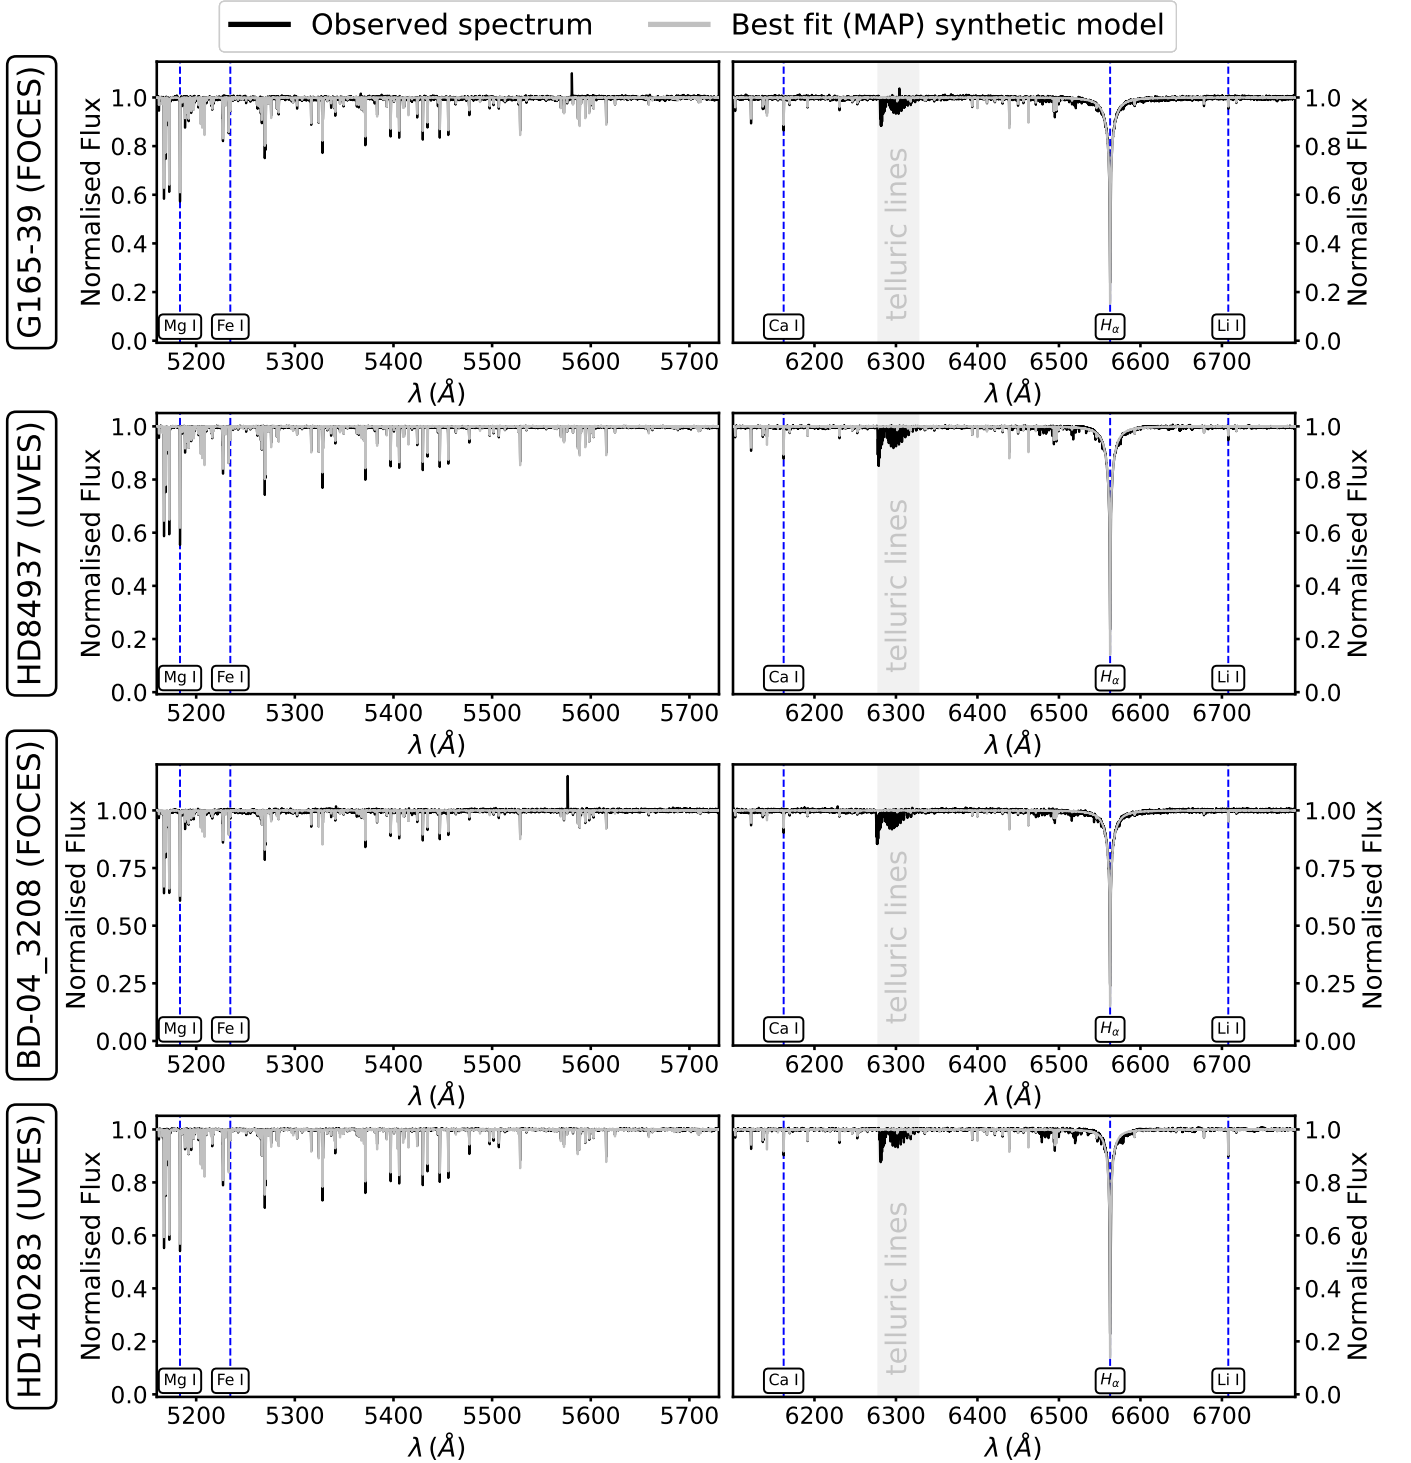
* Missing entries in the TSFitPy abundances are a result of too weak absorption line features in the considered wavelength range to perform a proper fit.

Table B.4. Overview of raw cINN MAP predictions for benchmark spectra.

Star (Spectrum)	T_{eff} (K)	$\log\left(\frac{g}{\text{cm}^2}\right)$	[Fe/H]	v_{broad} kms^{-1}	v_{mic} kms^{-1}	Ca	C	Li	[x/Fe]				f_{reject}	
									Mg	Mn	N	Si	Ti	
15PEG (FOCES)	6415	3.49	-0.64	11.14	1.66	0.08	-0.51	1.02	0.21	-0.26	0.61	-0.12	-0.04	0.13
18Sco (NARVAL)	5694	4.22	-0.14	5.61	1.2	0.09	-0.33	0.46	0.03	-0.26	0.55	-0.12	0.07	0
31AQL (FOCES)	5488	4.11	0.2	6.93	1.27	0.04	0.06	-0.12	0.06	-0.25	0.64	-0.09	0.18	0
61CygA (NARVAL)	4523	4.65	-0.39	5.02	0.54	0.18	0.14	0.2	0.23	-0.19	0.35	-0	0.4	0.5
61CygB (NARVAL)	4288	4.55	-0.6	5.4	0.61	0.42	0.25	0.02	0.1	-0.2	0.6	-0.04	0.43	0.72
Arcturus (NARVAL)	4415	1.77	-0.65	6.94	1.74	0.17	-0.1	0.17	0.39	-0.16	0.61	-0.05	0.59	0
Arcturus (UVES)	4446	1.89	-0.61	6.12	1.73	0.08	-0.15	0.27	0.33	-0.24	0.49	-0.07	0.57	0.01
Arcturus (UVES)	4435	1.83	-0.62	6.39	1.75	0.12	-0.2	0.06	0.41	-0.19	0.75	-0.03	0.57	0.01
BD+26_3578 (FOCES)	6379	4.04	-2.29	7.75	1.45	0.35	0.02	3.11	0.25	-0.34	0.2	0.32	0.22	0.22
BD-04_3208 (FOCES)	6274	4.06	-2.38	6.79	1.03	0.28	0.33	3.05	0.3	-0.31	0.33	-0.02	0.25	0.16
G113-22 (FOCES)	5642	3.65	-1.18	5.75	1.41	0.25	-0.47	1.84	0.41	-0.38	0.51	0.07	0.29	0.1
G165-39 (FOCES)	6304	3.9	-2.06	8.69	1.29	0.32	-0.23	2.64	0.2	-0.19	0.66	0.16	0.16	0.17
G207-5 (FOCES)	5835	4.11	-0.55	7.96	1.12	0.03	-0.51	1.2	0.12	-0.41	0.76	-0.1	0.06	0.36
G28-43 (FOCES)	5099	4.63	-1.72	9.1	1.19	0.32	-0.01	2	0.32	-0.56	0.81	0.24	0.32	0.35
G41-41 (FOCES)	6395	3.92	-2.8	2.76	1.46	0.13	0.29	3.49	0.23	-0.36	0.39	0.35	0.05	0.22
G48-29 (FOCES)	6395	4.43	-2.89	3.78	1.48	0.24	0.17	3.37	0.29	-0.5	0.46	0.35	0.28	0.3
G63-46 (FOCES)	5741	4.05	-0.83	5.65	1.01	0.22	-0.45	0.81	0.36	-0.25	0.86	0.02	0.2	0.23
G64-12 (FOCES)	6290	4.67	-3.77	1.22	1.12	0.25	0.32	4.62	0.64	-0.47	0.22	0.43	0.48	0.3
G78-1 (FOCES)	5903	3.94	-0.97	7.82	1.32	0.22	-0.2	1.66	0.39	-0.22	0.93	0.01	0.25	0.29
G9-16 (FOCES)	6764	4.1	-1.49	12.96	1.3	0.33	-0.27	2.19	0.32	-0.55	0.18	0.22	0.39	0.13
Gmb1830 (NARVAL)	5044	4.58	-1.43	2.67	1.05	0.35	0.08	1.53	0.37	-0.45	0.75	0.12	0.36	0.18
HD101177 (FOCES)	5829	4.1	-0.36	7.15	1.25	0.05	-0.46	1.17	0.11	-0.23	0.55	-0.15	0.08	0.02
HD102158 (FOCES)	5650	3.89	-0.54	5.84	1.04	0.26	-0.28	0.96	0.36	-0.37	0.88	-0.05	0.21	0.28
HD103095 (FOCES)	5080	4.64	-1.43	3.18	1.25	0.33	-0.1	1.47	0.36	-0.41	0.59	0.15	0.36	0.23
HD104304 (FOCES)	5436	4.4	0.04	7.12	1.23	0.07	0.11	0.03	0.06	-0.12	0.38	-0.07	0.16	0
HD105631 (FOCES)	5327	4.49	-0.07	6.17	1.22	0.09	-0.1	0.36	0.04	-0.15	0.57	-0.14	0.14	0
HD106210 (FOCES)	5583	4.17	-0.3	6.88	1.16	0.13	-0.08	0.81	0.24	-0.16	0.23	-0.03	0.21	0
HD106516 (FOCES)	6259	4.01	-0.81	9.99	1.53	0.18	-0.55	1.48	0.49	-0.27	0.67	0.06	0.14	0.24
HD10697 (FOCES)	5517	3.76	-0.07	7.86	1.24	0.04	-0.11	0.83	-0	-0.08	0.44	-0.17	0.11	0
HD10700 (FOCES)	5283	4.4	-0.67	5.64	1.07	0.23	-0.15	0.86	0.36	-0.13	0.67	0	0.3	0.05
HD107328 (UVES)	4561	2.15	-0.53	7.03	1.84	0.08	-0.31	0.24	0.25	-0.13	0.44	-0.09	0.53	0
HD107328 (NARVAL)	4541	2.06	-0.55	6.79	1.81	0.1	-0.08	0.12	0.34	-0.03	0.51	-0.07	0.52	0
HD110897 (FOCES)	5791	4.02	-0.68	6.97	1.36	0.08	-0.28	0.94	0.21	-0.31	0.65	-0.1	0.04	0.04
HD111395 (FOCES)	5490	4.29	-0.11	7.42	1.14	0.1	-0.34	0.35	0.01	-0.13	0.47	-0.1	0.13	0
HD113139 (FOCES)	6817	3.91	-0.02	91.06	1.29	0.06	-0.29	0.9	-0.06	-0.2	0.49	-0.01	0.01	0.08
HD114174 (FOCES)	5628	4.15	-0.14	6.83	1.22	0.08	-0.16	0.11	0.1	-0.22	0.36	-0.13	0.08	0
HD122563 (NARVAL)	4871	1.51	-2.33	4.56	1.53	0.1	0.17	2.13	0.1	-0.58	0.04	0.2	0.25	0.22
HD140283 (NARVAL)	5810	3.58	-2.42	4.01	1.12	0.08	0.2	3.43	0.18	-0.21	0.31	0.37	0.2	0.2
HD140283 (UVES)	5923	3.85	-2.36	2.7	1.45	0.06	0.05	3.28	0.19	-0.46	0.27	0.48	0.22	0.21
HD144284 (FOCES)	6179	3.52	0.06	31.15	1.67	-0.01	-0.45	0.72	-0.02	-0.26	0.93	-0.14	-0.06	0.16
HD162003 (FOCES)	6478	3.69	-0.15	15.32	1.77	0.04	-0.68	1.07	-0	-0.17	0.75	-0.23	-0.07	0.17
HD164259 (FOCES)	6671	3.69	-0.06	71.31	1.49	0.01	-0.51	0.51	0.03	-0.33	0.52	-0.02	0.01	0.1
HD220009 (NARVAL)	4453	1.96	-0.72	6.21	1.48	0.18	-0.18	0.23	0.37	-0.28	0.44	-0.05	0.53	0.01
HD22879 (NARVAL)	5867	4.09	-0.98	4.43	1.35	0.23	0.01	1.48	0.41	-0.27	0.81	0.05	0.27	0.15
HD284248 (FOCES)	6190	4.17	-1.66	5.59	1.32	0.25	-0.51	2.75	0.21	-0.51	0.63	0.04	0.02	0.28
HD84937 (UVES)	6285	3.85	-2.18	3.47	1.02	0.31	-0.36	2.68	0.17	-0.61	0.38	0.23	0.24	0.24
OMI_AQL (FOCES)	6071	3.9	0.09	7.23	1.31	0.05	-0.47	1.01	-0.01	-0.15	0.93	-0.19	0.02	0.6
Procyon (NARVAL)	6557	3.58	-0.12	8.93	1.92	-0.01	-0.59	0.55	0.07	-0.19	0.49	-0.17	-0.07	0.23
Procyon (UVES)	6378	3.35	-0.21	8.95	1.79	0.01	-0.6	0.8	0.07	-0.26	0.66	-0.16	-0.08	0.06
Procyon (UVES)	6421	3.44	-0.2	8.51	1.82	0.02	-0.58	0.67	0.13	-0.26	0.57	-0.15	-0.08	0.12
Sun (NARVAL)	5692	4.2	-0.07	4.24	0.92	0.09	-0.16	0.11	-0.01	-0.09	0.99	-0.16	0.07	0.66
Sun (NARVAL)	5711	4.29	-0.19	5.54	1.23	0.08	-0.18	0.57	0.07	-0.28	0.64	-0.11	0.1	0
Sun (UVES)	5576	4.06	-0.25	6.11	1.24	0.11	-0.15	0.56	0.09	-0.22	0.5	-0.11	0.09	0
α CenA (UVES)	5549	3.91	-0.05	6.74	1.32	0.11	-0.01	-0	0.07	-0.17	0.4	-0.04	0.07	0
α Cet (UVES)	3908	1.24	-0.48	6.5	1.65	0.17	-0.01	-0.32	0.33	-0.18	0.58	-0.08	0.48	0.05
α Cet (NARVAL)	3977	0.94	-0.5	6.67	1.31	0.28	-0.14	-0.5	0.02	-0.08	0.56	-0.09	0.6	0.11
α Tau (NARVAL)	4021	1.21	-0.39	6.96	1.59	0.14	-0.09	-0.53	0.22	-0.13	0.37	-0.16	0.43	0
α Tau (UVES)	4021	1.21	-0.42	5.15	1.53	0.03	-0.34	-0.5	0.11	-0.1	0.32	-0.2	0.52	0
β Hyi (UVES)	5695	3.51	-0.3	6.72	1.34	0.07	-0.4	1.48	0.13	-0.15	0.54	-0.16	0.03	0.03
β Vir (NARVAL)	6014	3.76	-0.01	7.38	1.48	0.03	-0.36	0.64	0.04	-0.21	0.62	-0.19	-0.03	0.03
δ Eri (NARVAL)	5005	3.66	-0.12	5.85	1.19	0.12	-0.06	0.3	0.04	-0.11	0.33	-0.12	0.26	0
δ Eri (UVES)	4993	3.64	-0.13	5.22	1.26	0.09	0.09	0.32	0.03	-0.2	0.32	-0.15	0.25	0
δ Eri (UVES)	4993	3.71	-0.14	5.28	1.27	0.09	0.04	0.37	0.11	-0.12	0.32	-0.13	0.25	0
ϵ Eri (UVES)	4998	4.5	-0.3	4.46	1.14	0.1	-0.07	0.25	0.07	-0.21	0.49	-0.08	0.18	0
ϵ Eri (UVES)	5053	4.61	-0.27	3.91	1.13	0.12	-0.23	0.36	0.07	-0.23	0.31	-0.09	0.22	0
ϵ Eri (UVES)	5026	4.48	-0.3	5.27	1.14	0.13	-0.14	0.22	0.04	-0.21	0.46	-0.07	0.21	0
ϵ Vir (NARVAL)	5084	2.75	-0.02	7.58	1.73	-0	-0.38	0.05	-0.07	-0.15	0.85	-0.17	0.21	0.04
η Boo (NARVAL)	6044	3.61	0.17	14.88	1.62	-0.01	-0.36	0.34	0	-0.14	0.65	-0.18	-0.01	0.07
μ Ara (UVES)	5607	3.99	0.06	6.4	1.33	0.09	-0.17	0.21	0.02	-0.19	0.73	-0.12	0.1	0
μ Cas (NARVAL)	5310	4.34	-0.93	4.27	0.97	0.34	-0.11	1.15	0.41	-0.35	0.74	0.11	0.33	0.1
μ Leo (NARVAL)	4574	2.55	0.03	7.34	1.68	0.05	-0.03	-0.28	-0	-0.19	0.59	-0.01	0.41	0
τ Cet (NARVAL)	5379	4.46	-0.6	2.62	0.99	0.26	-0.16	0.87	0.29	-0.24	0.6	-0.04	0.33	0.07

Table B.5. Coefficients for linear correction of systematic offset of cINN predictions.

X	$a_{X,\text{cal}}$	$b_{X,\text{cal}}$	X	$a_{X,\text{cal}}$	$b_{X,\text{cal}}$
T_{eff}	-0.096	492.389	[Mg/Fe]	-2.479×10^{-5}	0.077
$\log(g)$	-2.139×10^{-4}	1.082	[Ca/Fe]	-4.095×10^{-5}	0.209
[Fe/H]	-3.209×10^{-5}	0.143	[Li/Fe]	-5.254×10^{-4}	2.779

**Fig. B.2.** Comparison of the 4MOST-ified input spectra to model spectra corresponding to the cINN MAP estimates for four low-metallicity stars. The highlighted Mg, Fe, Li and Ca absorption lines correspond to the zoom-ins shown in Fig. 7.