








# cecilia: A Machine Learning-Based Pipeline for Measuring Metal Abundances of Helium-rich Polluted White Dwarfs

Mariona Badenas-Agusti,<sup>1,2,3\*</sup>  Javier Viaña,<sup>2</sup>  Andrew Vanderburg,<sup>2</sup>  Simon Blouin,<sup>4</sup> 

Patrick Dufour,<sup>5</sup>  Siyi Xu (许偲艺),<sup>6</sup>  Lizhou Sha<sup>7</sup> 

<sup>1</sup> Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

<sup>2</sup> Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

<sup>3</sup> MIT William Asbjornsen Albert Memorial Fellow

<sup>4</sup> Department of Physics and Astronomy, University of Victoria, Victoria, BC V8W 2Y2, Canada

<sup>5</sup> Département de Physique, Université de Montréal, Montréal, Québec H3C 3J7, Canada

<sup>6</sup> Gemini Observatory/NSF's NOIRLab, 670 N. A'ohoku Place, Hilo, HI, 96720, USA

<sup>7</sup> Department of Astronomy, University of Wisconsin-Madison, 475 N Charter St, Madison, WI 53706, USA

Accepted 2024 February 06. Received 2024 February 06; in original form 2023 December 19

## ABSTRACT

Over the past several decades, conventional spectral analysis techniques of polluted white dwarfs have become powerful tools to learn about the geology and chemistry of extrasolar bodies. Despite their proven capabilities and extensive legacy of scientific discoveries, these techniques are however still limited by their manual, time-intensive, and iterative nature. As a result, they are susceptible to human errors and are difficult to scale up to population-wide studies of metal pollution. This paper seeks to address this problem by presenting *cecilia*, the first Machine Learning (ML)-powered spectral modeling code designed to measure the metal abundances of intermediate-temperature ( $10,000 \leq T_{\text{eff}} \leq 20,000$  K), Helium-rich polluted white dwarfs. Trained with more than 22,000 randomly drawn atmosphere models and stellar parameters, our pipeline aims to overcome the limitations of classical methods by replacing the generation of synthetic spectra from computationally expensive codes and uniformly spaced model grids, with a fast, automated, and efficient neural-network-based interpolator. More specifically, *cecilia* combines state-of-the-art atmosphere models, powerful artificial intelligence tools, and robust statistical techniques to rapidly generate synthetic spectra of polluted white dwarfs in high-dimensional space, and enable accurate ( $\leq 0.1$  dex) and simultaneous measurements of 14 stellar parameters—including 11 elemental abundances—from real spectroscopic observations. As massively multiplexed astronomical surveys begin scientific operations, *cecilia*'s performance has the potential to unlock large-scale studies of extrasolar geochemistry and propel the field of white dwarf science into the era of Big Data. In doing so, we aspire to uncover new statistical insights that were previously impractical with traditional white dwarf characterisation techniques.

**Key words:** stars: white dwarfs - stars: atmospheres - stars: abundances - techniques: spectroscopic - methods: data analysis - methods: observational

## 1 INTRODUCTION

In recent years, the field of exoplanets has undergone remarkable growth, revolutionising our understanding of the Universe with the discovery of more than 5500 distant worlds.<sup>1</sup> Despite this notable progress, the detection of a potentially habitable Earth-sized planet with an Earth-like interior remains elusive.

In our quest to find exoplanets similar to our own, knowledge of a planet's bulk composition will be crucial to understand its formation and evolution history, geology, climate, and habitability conditions. Traditionally, transit photometry (Charbonneau et al. 2000)

and Doppler spectroscopy (Campbell et al. 1988) have been used to measure the radii and masses of exoplanets, respectively. These two parameters, when combined, provide the planetary bulk density, which can then be compared to interior models to infer the planet's most likely composition (Zeng et al. 2016). Nevertheless, this approach is incomplete because it does not account for the specific elemental abundances of the planet. In addition, for planets smaller than  $R_p \sim 4 R_{\oplus}$ , large uncertainties in radius and mass measurements ( $> 10\%$ ) pose challenges in distinguishing between different interior compositions (Zeng & Seager 2008). These problems are further complicated by degeneracies in theoretical mass-radius diagrams, which can yield various combinations of metal, silicate, water/ice, and gas for the same density measurement (Dorn et al. 2015; Rogers & Seager 2010; Seager et al. 2007).

\*E-mail: mbadenas@mit.edu

<sup>1</sup> NASA Exoplanet Database (<https://exoplanets.nasa.gov/>). Last Accessed: January 2024.

To date, the most direct —and only— way of overcoming these ambiguities and probing the composition of small, terrestrial exoplanets is through spectroscopy of externally polluted White Dwarfs (WDs). These objects are the degenerate stellar cores of low- and intermediate-mass ( $\leq 10 M_{\odot}$ ) main-sequence stars. With a typical radius similar to the size of the Earth ( $R_{WD} \approx 1 R_{\oplus}$ ), and a mass half of that of the Sun ( $M_{WD} \approx 0.6 M_{\odot}$ ), they are extremely dense and compact, which causes their atmospheres to have a stratified chemical structure. In other words, light elements (i.e. H and He) are expected to float in the thin upper layers of their atmospheres, while elements heavier than helium (or “metals,” e.g. Ca, Mg, Si, Fe) should sink rapidly towards the core with diffusion timescales much shorter than the cooling age of the white dwarf (Paquette et al. 1986; Koester 2009). In reality, however, observations suggest that between 25% and 50% of single white dwarfs are contaminated with heavy elements in their atmospheres (Koester et al. 2014; Zuckerman et al. 2010, 2003), likely from the recent or ongoing accretion of disrupted exoplanetary debris from their surviving outer planetary systems (Jura 2003; Zuckerman et al. 2003, 2010; Koester et al. 2014; Jura & Young 2014).

During the past two decades, high-resolution spectroscopy of these “polluted” white dwarfs has offered an excellent opportunity to determine the bulk composition of their accreted material. This technique has enabled the detection of more than 20 heavy elements (e.g. see references in Table 1 of Klein et al. 2021), revealing that most exoplanetary bodies are dry and rocky (Veras 2021), with some water-rich exceptions (e.g. Farihi et al. 2011, 2013; Raddi et al. 2015; Xu et al. 2017; Hoskin et al. 2020). Despite the significance of these discoveries, conventional analysis techniques of polluted white dwarfs entail time-consuming, manual, and iterative work, which makes them prone to human errors. As a result, these techniques may lead to wrong or sub-optimal results, potentially biasing our understanding of the geology and chemistry of white dwarf pollutants. With these challenges in mind, it has been difficult to transition from individual to population-wide studies of metal pollution. As we move into an era of massive multi-object spectroscopic surveys, when tens of thousands of white dwarfs will be observed and many of them are expected to show multiple heavy elements, conventional analysis techniques will thus struggle to keep up.

In this paper, we address the limitations of “classical” white dwarf characterisation methods by leveraging the power of Machine Learning (ML). Recently, multiple ML-based tools have been developed to interpolate and efficiently model spectra with the goal of estimating stellar abundances. Notable examples include *The Cannon* (aimed at deriving stellar properties based on a large training set of reference stars with well-known abundances; Ness et al. 2015), *The Payne* (designed to estimate 25 elemental abundances of red giants from low-resolution APOGEE spectra; Ting et al. 2019), *wdtools* (dedicated to the analysis of H-rich white dwarf spectra with no metal pollution; Chandra et al. 2020) and *TheHotPayne*, which is an adaptation of *The Payne* to OBA stars (Xiang et al. 2022). These ML pipelines have undoubtedly shown promising results, but they have not been designed to fit the complex and diverse spectra of polluted white dwarfs. This paper seeks to bridge this gap by presenting the first ML-based code capable of achieving this goal.

Inspired by the groundbreaking contributions of astrophysicist Cecilia Payne and building upon the ML framework of *The Payne*, we have chosen to name our code “*cecilia*” —or *Clues of Exoplanet Compositions Inferred from poLluted WhItE dWArfs*. With *cecilia*, we combine state-of-the-art atmosphere models with innovative ML methods to improve the performance of conventional white dwarf analysis techniques and ultimately enable large-scale

studies of metal pollution. From a design standpoint, *cecilia*’s ML architecture contributes to this goal by adhering to three fundamental principles: automation, efficiency, and flexibility. As an automated pipeline, *cecilia* solves the dependence of current techniques on human-based, time-intensive, and iterative procedures. In terms of efficiency, our code harnesses the speed of ML tools to rapidly estimate elemental abundances of polluted white dwarfs with a retrieval accuracy similar to that of traditional methods. Lastly, *cecilia* has a flexible ML architecture that can be re-trained at any time with new, better, and/or larger models. This also makes our pipeline a particularly useful feature extraction tool beyond the fields of white dwarf and exoplanetary science.

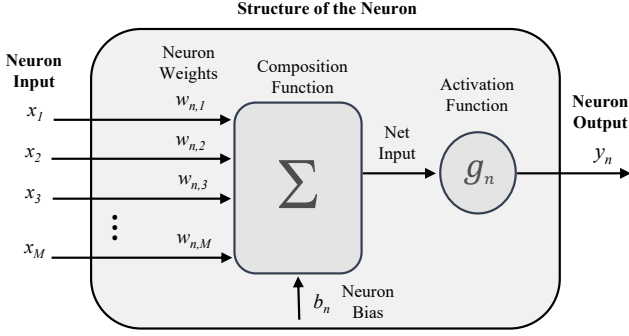
In this work, we describe the building blocks of *cecilia*, its current performance and limitations, and potential opportunities for improvement. The organisation of the paper is as follows. In Section 2, we provide a brief introduction to the concept of deep neural networks. In Section 3, we present the stellar parameters and model atmospheres used to train, validate, and test our code. Section 4 explains *cecilia*’s ML architecture, our training procedure, and our framework for estimating the elemental abundances of polluted white dwarfs from their spectra. In Section 5, we test and explore *cecilia*’s functionality with synthetic and real spectroscopic observations. Section 6 discusses the strengths and weaknesses of *cecilia*, possible future work to enhance its retrieval accuracy, and its relevance in the context of massive datasets from wide-field astronomical surveys. Finally, we offer our conclusions in Section 7.

## 2 DEEP NEURAL NETWORKS

Since the 1950s (Turing 1950; McCarthy et al. 2006), the field of modern Artificial Intelligence (AI) has sought to create intelligent systems capable of emulating human skills, both in terms of human reasoning and behavior. In recent years, the rapid development of technology has transformed this theoretical aspiration into a tangible reality, causing an important shift from model- to data-driven scientific analysis. The creation of faster and more powerful computers, coupled with the availability of extremely large and complex datasets (or “Big Data”), has only fueled the growth of AI, making neural networks increasingly relevant across a variety of disciplines, including astrophysics (Smith & Geach 2023).

To date, the task of developing computer programs that can learn, extract, and process meaningful insights from observations is typically referred to as Machine Learning (Samuel 1959). At the core of many ML algorithms are multilayered, fully connected, feed-forward neural networks (NN). These networks are learning systems where information flows *forward* through several layers of artificial nodes called “neurons.” A basic NN consists of an input layer, one or more hidden layers, and an output layer. The higher the number of hidden layers (or “depth”), the more complex the ML architecture becomes (LeCun et al. 2015). This is the case of *cecilia*, which consists of three neural networks with 5-6 layers each. In general, deeper architectures are more capable of recognising more intricate relationships in the data (Goodfellow et al. 2016).

In addition to the choice of ML architecture, the learning capabilities of a NN are also shaped by the way its neurons are interconnected with one another. This connection is regulated by three parameters known as “weights,” “biases,” and “activation functions.” The weights and biases determine the importance of the neuronal connections, while the activation functions employ non-linear operations to transform the input data of a set of neurons into useful output param-



**Figure 1.** Parameters and operations associated to a single “neuron,” which represents the smallest unit of a neural network.

eters. In particular, for a generic neuron  $n$ , the activation function  $g_n$  performs the transformation

$$y_n = g_n(b_n + \sum_{m=1}^M w_{n,m} \cdot x_m), \quad (1)$$

where  $M$  denotes the total number of inputs fed into the neuron,  $m$  is a generic input index,  $x$  and  $y$  are the input and the output data, and  $w_{n,m}$  and  $b_n$  represent the optimised weight and bias of the neuron, respectively (see Figure 1). Typically, the activation function  $g_n$  is chosen by the user experimentally (e.g. see Section 4.1).

In the field of ML, the *training* of a NN is an iterative procedure designed to fine-tune the weights and biases of the neurons with the goal of minimising the error (also known as “loss” or “cost” function) between the network’s final predictions and an assumed ground truth. During this process, the NN is exposed to labeled observations in order to identify hidden patterns in the data and learn to generate predictions that closely resemble the expected outcome for a given set of input parameters. As the training proceeds, the predictive accuracy of the NN is gradually improved with a technique known as “backpropagation.” This technique calculates the loss function at each iteration to adjust the model parameters of the network. A standard formulation of the loss function is the quadratic expression

$$\mathcal{J}_q = \frac{1}{2} \sum_{p=1}^P (y_{q,p} - \hat{y}_{q,p})^2 = \frac{1}{2} \sum_{p=1}^P e_{q,p}^2, \quad (2)$$

where  $q$  is the index of a generic instance,<sup>2</sup>  $P$  is the total number of output neurons,  $p$  is the index of a generic output neuron,  $\hat{y}_{q,p}$  is the prediction of the  $p$ -th output neuron,  $y_{q,p}$  is the known true output, and  $e$  denotes the residual error difference between  $\hat{y}$  and  $y$ . For a given training dataset, the aggregated loss function of all the instances  $Q$  can then be expressed as

$$\mathcal{J} = \sum_{q=1}^Q \mathcal{J}_q. \quad (3)$$

During training, the aggregated loss function is computed at every step in order to optimise the weights and biases of the network. These parameters are updated iteratively through the gradient of the loss function,  $\nabla \mathcal{J}(w_{n,m})$ . More specifically, the incremental change of a generic weight in the network, or  $\Delta w_{n,m}$ , is proportional to the

derivative of the total loss function with respect to the current value of the weight,

$$\Delta w_{n,m} = w_{n,m}^{\text{new}} - w_{n,m}^{\text{old}} = -\eta \frac{\partial \mathcal{J}}{\partial w_{n,m}}. \quad (4)$$

In this expression, the partial derivative  $\partial \mathcal{J} / \partial w_{n,m}$  dictates how much the network parameters should be changed in order to reach a minimum of the loss function, with the negative sign indicating that the direction of the “descent” should be against the gradient and towards the minimum. The term  $\eta$  is the “learning rate” of the network, i.e. a user-defined tuning parameter which provides stability to the learning process and controls the step size of the gradient descent.

Broadly speaking, the partial derivative  $\partial \mathcal{J} / \partial w_{n,m}$  in Eq. 4 is proportional to the inputs of the neurons, the derivatives of their activation functions with respect to their weights, and the errors of the network. For the output layer, these factors are decoupled, which makes the computation of the derivative relatively straightforward. However, for the earlier (and deeper) layers, the term  $\partial \mathcal{J} / \partial w_{n,m}$  depends on the derivative of the loss function for the outer layers, which in turn is proportional to the *product* of the derivatives of all the activation functions with respect to their weights. This product operation is implemented with the chain rule from the output to the initial layer, i.e. *backwards*, hence giving rise to the concept of *backpropagation*. We note that in many ML applications, including *cecilia*, the exact value of the gradient at each training iteration is not computed with the full training sample, but approximated with subsets (or “batches”) of the latter (e.g. [Shallue & Vanderburg 2018](#)).

### 3 DATA

In this section, we motivate the use of a neural-network spectral interpolator, describe the atmosphere models and synthetic stellar properties used to train, validate, and test *cecilia*, and discuss the normalisation techniques used to improve the learning and predictive abilities of our code.

#### 3.1 Generation of a Training Set: Motivation

Current white dwarf spectral modelling techniques rely on grids of stellar properties (or “labels”) with uniform step sizes (e.g. [Coutu et al. 2019](#)). Although this approach has worked well over the years, it can be computationally expensive for a large number of free model parameters. For example, using 13 stellar properties —as in this work (see Section 3.1.1 below)— and a relatively coarse grid of 10 points per label, conventional interpolators would require at least  $10^{13}$  models to explore the parameter space of their study. Assuming, on average, that each model takes about 1 hour to generate on a single CPU core, it would take on the order of  $10^9$  CPU-years to compute the whole grid. Therefore, classical methods are prohibitively expensive.

To address this challenge, we have developed *cecilia*, a ML generative spectral interpolator capable of sampling from a high-dimensional label space based on *randomly drawn* training parameters, rather than a uniformly spaced grid of model points. With this ML approach, our pipeline only needs about 20,000 synthetic models to predict white dwarf spectra, i.e. about  $2 \times 10^{-9}\%$  of the models that classical methods would typically require for the same task.

<sup>2</sup> In this work, we define an “instance” as a spectrum with its corresponding 13 stellar labels.

### 3.1.1 Synthetic Stellar Labels

In this work, we have chosen to characterise He-rich polluted white dwarfs with 13 stellar properties, namely: their effective temperature ( $T_{\text{eff}}$ ) and surface gravity ( $\log g$ ), their hydrogen abundance, and 10 metal abundances for Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, Ni (see Table 1).<sup>3</sup> In our models, we have also considered 14 additional heavy elements (C, N, Li, Na, Al, P, S, Cl, Ar, K, Sc, V, Co, and Cu). However, none of these metals are fitted during *cecilia*'s optimisation procedure.

To prepare *cecilia*'s training set of synthetic spectra, we generated 25,000 unique combinations of the 13 aforementioned stellar labels (i.e.  $T_{\text{eff}}$ ,  $\log g$ , and 11 elemental abundances). In particular, for  $T_{\text{eff}}$ , we generated our labels from a random uniform distribution spanning the range [10,000, 20,000] K. At higher effective temperatures, the presence of photospheric metal pollution may not be due to the accretion of exoplanetary material, but to the effects of radiative levitation pressure (Chayer et al. 1995). For cooler white dwarfs, there is less thermal energy to produce atomic transitions in the optical; as a result, there is a limited number of polluted systems exhibiting multiple heavy elements at once. With respect to  $\log g$ , we followed Chandra et al. (2020) and sampled our values from a random uniform distribution between [7,9] cgs (with  $g$  in cgs units of  $\text{cm}\cdot\text{s}^{-2}$ ). For  $\log_{10}(\text{H}/\text{He})$ , we also drew our labels from a random uniform distribution between [-7,-3] dex, as in Coutu et al. (2019).

To generate our labels for the other elements, we used the following procedure. First, we drew our calcium abundances from a random uniform distribution between [-12,-7] dex:

$$\mathcal{U}_{\log(\text{Ca}/\text{He})} = \mathcal{U}(-12, -7). \quad (5)$$

For the remaining 10 heavy elements (i.e. Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni), we generated a random Gaussian distribution centered around  $\mu_Z$  with a standard deviation  $\sigma_Z$ ,

$$\mathcal{N}_{\log(Z/\text{Ca})} = \mathcal{N}(\mu_Z, \sigma_Z), \quad (6)$$

where  $\sigma_Z = 2$  dex and

$$\mu_Z = \log_{10}(Z/\text{He})_{\odot} - \log_{10}(\text{Ca}/\text{He})_{\odot} = \log_{10}(Z/\text{Ca})_{\odot}. \quad (7)$$

In the above expressions,  $Z$  is the metal under consideration,  $\log_{10}(Z/\text{He})_{\odot}$  is the abundance of metal  $Z$  in the solar photosphere (Asplund et al. 2009), and  $\log_{10}(\text{Ca}/\text{He})_{\odot}$  is the solar calcium abundance relative to helium (Asplund et al. 2009). We note that our choice of  $\sigma_Z = 2$  dex was motivated by two factors: on the one hand,  $\sigma$  had to be large enough to allow *cecilia* to measure the true elemental abundances of a white dwarf, even if the latter deviated from chondritic values; and on the other hand, we also wanted  $\sigma$  to be small enough to down-weight unlikely regions of the parameter space in our training set, hence limiting the hypervolume of *cecilia*'s label parameter space.

As a final step, we generated the distribution  $\mathcal{D}$  of abundance ratios  $\log(Z/\text{He})$  for our training set by adding

$$\mathcal{D} = \mathcal{N}_{\log(Z/\text{Ca})} + \mathcal{U}_{\log(\text{Ca}/\text{He})} \quad (8)$$

In Table 1, we list the values of  $\mu_Z$  and  $\sigma_Z$  for Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni, rather than the properties of the (non-Gaussian) distributions resulting from evaluating Eq. 8.

Finally, for the remaining 14 heavy elements (C, N, Li, Na, Al, P, S, Cl, Ar, K, Sc, V, Co, and Cu), we chose to scale their abundances such that their abundance ratios relative to calcium corresponded

**Table 1.** *cecilia*'s ranges used to generate the training set of synthetic white dwarf properties (or stellar “labels”). All the atmospheric abundances are expressed as number abundance ratios in base 10 relative to helium. References (Ref.): [1] Chandra et al. (2020), [2] Coutu et al. (2019), [3] Asplund et al. (2009).

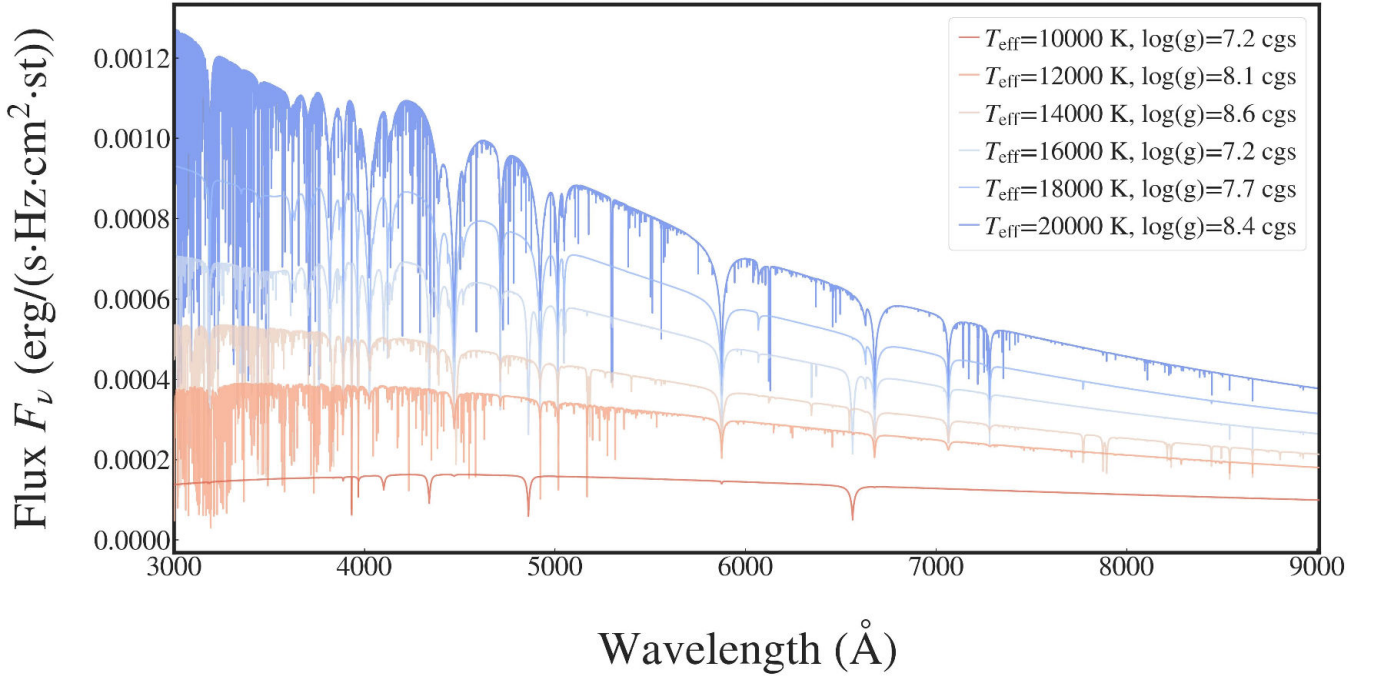
Sampled from a Random Uniform Distribution			
Label	Min. Bound	Max. Bound	Ref.
$T_{\text{eff}}$ [K]	10,000	20,000	-
$\log g$ [cgs]	7	9	[1]
$\log_{10}(\text{H}/\text{He})$	-7	-3	[2]
$\log_{10}(\text{Ca}/\text{He})$	-12	-7	[2]
Sampled from a Random Gaussian Distribution (Eq. 7)			
Label	$\mu_Z$ [dex]	$\sigma_Z$ [dex]	Ref.
$\log(\text{Mg}/\text{Ca})$	1.26	2	[3]
$\log(\text{Fe}/\text{Ca})$	1.16	2	[3]
$\log(\text{O}/\text{Ca})$	2.35	2	[3]
$\log(\text{Si}/\text{Ca})$	1.17	2	[3]
$\log(\text{Ti}/\text{Ca})$	-1.39	2	[3]
$\log(\text{Be}/\text{Ca})$	-4.96	2	[3]
$\log(\text{Cr}/\text{Ca})$	-0.70	2	[3]
$\log(\text{Mn}/\text{Ca})$	-0.91	2	[3]
$\log(\text{Ni}/\text{Ca})$	-0.12	2	[3]

to the values of CI chondrites reported in in Lodders (2003). This type of bodies are primitive carbonaceous meteorites with a pristine composition unaltered by melting and differentiation processes. With the exception of their volatile components, including their lithium and noble gas abundances, their chemical makeup —mostly dominated by O, Mg, Si, and Fe— closely resembles that of the solar photosphere (Palme et al. 2014) as well as that of bulk Earth (e.g. Zuckerman et al. 2007; Xu et al. 2019; Doyle et al. 2023). Therefore, CI chondrites are often used as a proxy for the average bulk composition of non-volatile objects in the Solar System (Lodders 2003; Jura & Young 2014). In the field of polluted white dwarfs, the assumption of chondritic ratios often constitutes a reasonable first-order approximation when there is no *a priori* knowledge of the stellar photospheric composition (Dufour et al. 2007; Coutu et al. 2019). In the context of this work, the use of chondritic abundances for the 14 “fixed” metals also allowed us to dramatically decrease the volume of the parameter space that *cecilia* had to learn during training.

### 3.1.2 Synthetic Atmosphere Models

Next, we used the 25,000 synthetic labels to generate their corresponding atmosphere models. To achieve this, we employed the code described in Dufour et al. (2007); Blouin et al. (2018a,b) and references therein. This code includes all the necessary physical principles to predict the emerging spectra of metal-polluted white dwarfs and it has been used successfully to reproduce ultraviolet (UV) and optical observations across a wide range of effective temperatures and metal abundances (Dufour et al. 2012; Giammichele et al. 2012; Vanderburg et al. 2015; Melis & Dufour 2017; Xu et al. 2017; Blouin et al. 2019a,b; Coutu et al. 2019; Blouin 2020; Kaiser et al. 2021; Klein et al. 2021; Caron et al. 2023; Doyle et al. 2023). More specifically, for each model, the code calculates a 1D thermodynamic structure first, and this stratification is used to generate a synthetic spectrum with no instrumental broadening and sampled with a pixel spacing

<sup>3</sup> In this work, all the abundances are expressed as logarithmic number abundance ratios (in base 10) relative to helium.



**Figure 2.** A selection of 6 atmosphere models for synthetic He-rich polluted white dwarfs with effective temperatures between 10,000 K and 20,000 K, and surface gravities between 7 cgs and 9 cgs. Cool and hot white dwarfs are shown, respectively, in red and blue.

equal to  $R \approx 50,000$  (i.e. 55,000 wavelength points between 3,000 Å and 9,000 Å). Figure 2 shows a selection of synthetic spectra for He-rich polluted white dwarfs spanning different effective temperatures, surface gravities, and photospheric compositions.

For *cecilia*, we independently varied 13 parameters for each model (i.e.  $T_{\text{eff}}$ ,  $\log g$ , and 11 elemental abundances; see Table 1), which we randomly selected for each calculation as described in Section 3.1.1. In total, we calculated 25,000 atmosphere models, but 2,176 of them (8.7%) had to be discarded due to computational issues, leaving 22,824 useful synthetic spectra for training, testing, and validation.<sup>4</sup> We note that exclusion of these models did not have any significantly impact *cecilia*’s predictive ability; this is due to our code’s ML-based spectral interpolation approach, which facilitates the computation of spectral models in regions of the parameter space where there are no observations. We also recall that while the abundances of 10 metals (Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, Ni) were varied independently between different synthetic spectra, the remaining 14 heavy elements (C, N, Li, Na, Al, P, S, Cl, Ar, K, Sc, V, Co, and Cu) were also included in our models with their abundance ratios  $Z/\text{Ca}$  fixed to their nominal CI chondritic abundance from [Lodders \(2003\)](#).

### 3.2 Data Normalisation

As is standard procedure in ML, we then normalised the synthetic spectra and their corresponding stellar labels between 0 and 1 in order to improve *cecilia*’s learning abilities during training. For our synthetic spectra, we experimented with different scaling techniques and ultimately adopted a “pixel-wise” (PW) normalisation

approach (also commonly known as *min-max* scaling in ML). For a given wavelength index  $\lambda_i$ , this approach performs the following linear transformation: first, it determines the minimum and maximum values (or “pixels”) of the synthetic flux across *all* the available spectra; and then, it uses these flux bounds to normalise the pixel at wavelength  $\lambda_i$  to a value between 0 and 1. Therefore, the normalisation is conducted with respect to the flux axis, rather than along the wavelength dimension.

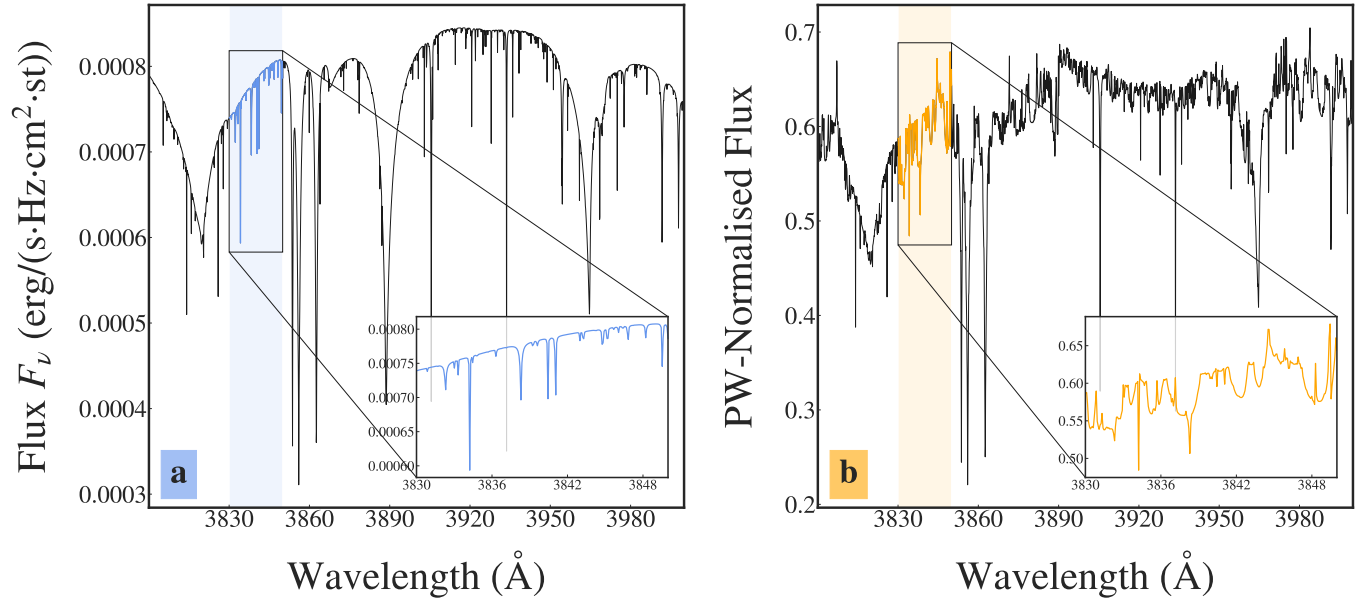
As illustrated in Figure 3, the PW approach may sometimes cause the normalised spectra to appear distorted to the human eye. Nevertheless, it is a powerful scaling technique because it treats each pixel as an independent spectral “feature,” thus allowing *cecilia* to learn both the small- and large-scale variations in the synthetic flux. In other words, pixels associated to large variations become less prominent, while those linked to smaller changes become magnified. In addition to this advantage, the PW normalisation differs from other ML scaling techniques (e.g. log-scaling) because it treats each pixel in the same way, therefore ensuring that *cecilia* does not have any preference for a specific absorption line.

## 4 METHODOLOGY

### 4.1 Design of *cecilia*’s ML Architecture

The central problem in this paper lies in accurately determining the main properties—or stellar labels—of polluted white dwarfs from their spectra. To achieve this goal, we have developed *cecilia*, an innovative ML pipeline capable of performing two tasks: (i) the generation of polluted white dwarf spectra from a set of stellar labels, i.e.  $S = \beta_1(\mathcal{L})$ , where  $\mathcal{L}$  is the vector space of the white dwarf labels (with dimension  $L$ ), and  $\beta_1$  is a generic non-linear function; and (ii) the prediction of polluted white dwarf labels from their corresponding spectra, or  $\mathcal{L} = \beta_2(S)$ , where  $S$  is the vector space of

<sup>4</sup> The missing models were not clustered in any specific region of the parameter space. Therefore, they did not negatively affect the robustness of *cecilia*.



**Figure 3.** Effect of applying the Pixel-Wise (PW) normalisation technique on a synthetic spectrum in the spectral window between 3,800 Å to 4,000 Å. This technique is designed to improve *cecilia*’s learning capabilities by amplifying any small variations in the stellar flux and attenuating the largest fluctuations. Panels (a) and (b) show, respectively, the raw and PW-normalised synthetic flux of the spectrum. The inset plots show the region between 3,330 Å and 3,350 Å.

the white dwarf spectrum (with dimension  $S$ ), and  $\beta_2$  is a function different to  $\beta_1$ .

Developed with the open-source Python `tensorflow` package (Abadi et al. 2015), *cecilia* is characterised by three deep neural networks: an Autoencoder, a Fully Connected Neural Network (FCNN1), and a replica of the FCNN1 for fine-tuning purposes (FT FCNN2). The Autoencoder is a ML tool consisting of two symmetrical networks (Figure 4): an *Encoder* and a *Decoder*. The *Encoder* identifies the main attributes of the input data and compresses the latter into a lower-dimensional representation. This operation takes place in the so-called “hidden” or “latent” space (denoted by  $\mathcal{H}$ ), where  $\mathcal{H}$  has a dimensionality equal to the number  $H$  of hidden/latent features used to encode the data (i.e.  $\mathcal{H} \in \mathbb{R}^H$ ). The second part of the Autoencoder is the *Decoder*, which is a network designed to reconstruct the original input data based on the *Encoder*’s latent model while trying to minimising the loss of information caused by the dimensionality reduction. We emphasise that the hidden features of the Autoencoder are fundamental properties of the ML model and are not intrinsically correlated to any stellar parameters. As a result, they have no astrophysical interpretation. In general, the choice of  $H$  constitutes a trade-off between the interpretability and the complexity of the ML model, characterised respectively by a low and a high  $H$  (Liang et al. 2023).

Following the Autoencoder, *cecilia* incorporates two deep neural networks: the FCNN1 and the FT FCNN2. As shown in Figures 5 and 6, these networks consist of an input layer, six hidden layers, and an output layer. Unlike the Autoencoder, which is used to perform compression and dimensionality reduction, the FCNN1 and FT FCNN2 architectures are designed to teach *cecilia* how to interpolate white dwarf spectra based on a set of 13 stellar labels. We provide a more detailed discussion about this subject in Section 4.2.

At the level of individual neurons, the choice of weights, biases, and activation functions can also shape the performance of a NN, as explained in Section 2. In this work, we followed common ML practice and set the weights of *cecilia* to random values before

optimising them during training. More specifically, we used the `tensorflow` functions `he_uniform` and `glorot_uniform` to initialise the weights of all the initial layers and the final output layer, respectively. These two functions draw values from random uniform distributions in the range between  $[-\text{limit}, +\text{limit}]$ . If  $\text{fan}_{\text{in}}$  and  $\text{fan}_{\text{out}}$  are, respectively, the number of input and output connections of a neuron, the parameter “limit” in `tensorflow` is defined as

$$\text{limit} = \begin{cases} \sqrt{6/\text{fan}_{\text{in}}} & (\text{he\_uniform}) \\ \sqrt{6/(\text{fan}_{\text{in}} + \text{fan}_{\text{out}})} & (\text{glorot\_uniform}) \end{cases}. \quad (9)$$

With respect to *cecilia*’s activation functions, we opted for using a Rectified Linear Unit (ReLU; Jarrett et al. 2009) and a sigmoid function (Goodfellow et al. 2016), as specified in Table 2. These two activation functions regulate the flow of information by introducing non-linearities across the networks, normalising the neurons’ output parameters, and enabling the identification of hidden and complex patterns in the training sample. For an input value  $x$ , the ReLU function returns 0 if  $x$  is negative, or  $x$  if  $x$  is positive,

$$\text{ReLU}(x) = \begin{cases} 0 & \text{if } x < 0 \\ x & \text{otherwise} \end{cases} = \max(0, x). \quad (10)$$

Therefore, its derivative is given by

$$\frac{d}{dx} [\text{ReLU}(x)] = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases}. \quad (11)$$

In contrast, the sigmoid activation function  $\sigma(x)$  transforms the input value  $x$  into a value between 0 and 1 with

$$\sigma(x) = \frac{1}{1 + e^{-x}}. \quad (12)$$

As a consequence, its derivative approaches zero when  $x \rightarrow \pm\infty$ . In neural networks with multiple hidden layers, the potential saturation of the sigmoid function towards zero values can lead to the so-called “vanishing gradient problem” (Hanin 2018). This issue

**Table 2.** Main properties of *cecilia*’s ML architecture. Its three neural networks are trained sequentially, but only the final one (i.e. the FT FCNN2) is used to model the spectra of He-rich polluted white dwarfs.

Parameter	Value
<i>Autoencoder</i>	
No. of layers in Encoder	5
No. of layers in Decoder	5
No. of neurons/layer	4,000
No. of training epochs	20,000
No. of latent features	100
Learning rate	$10^{-6}$
Batch size training	64
Batch size validation	64
Activation functions	ReLU (initial 4 layers) sigmoid (output layer)
Weight Initialisers	he_uniform (initial 4 layers) glorot_uniform (output layer)
<i>Fully-Connected Neural Network (FCNN1)</i>	
No. of total layers	6
No. of neurons/layer	2,000
No. of training epochs	10,000
Learning rate	$10^{-5}$
Batch size training	64
Batch size validation	64
Activation functions	ReLU (initial 5 layers) sigmoid (output layer)
Weight Initialisers	he_uniform (initial 5 layers) glorot_uniform (output layer)
<i>Fine-Tune Fully-Connected Neural Network (FT FCNN2)</i>	
No. of total layers	6
No. of neurons/layer	2,000
No. of training epochs	1,000
Learning rate	$10^{-4}$
Batch size training	64
Batch size validation	64
Activation functions	ReLU (initial 5 layers) sigmoid (output layer)
Weight Initialisers	he_uniform (initial 5 layers) glorot_uniform (output layer)

arises during backpropagation when the gradient of the loss function—which encapsulates the product of the derivatives of the activation functions with respect to their weights (see Section 2)—becomes extremely small. Due to this “vanishing” effect, the weights and biases of the initial layers may only be marginally updated, which can cause poor and slow model convergence. As evidenced by Eq. 11, ReLU functions are more robust to this problem because they have a more stable derivative (either 0 or 1). As a result, they often yield a better performance during training (Nair & Hinton 2010). In this work, we decided to employ ReLU activation functions for the initial layers of *cecilia*’s networks in order to avoid the vanishing gradient of the loss function. For the final layer of each network, we chose to use a sigmoid function with the goal of generating predicted values between 0 and 1.

In Table 2, we summarise the fundamental properties of *cecilia*’s Autoencoder, FCNN1, and FT FCNN2. These properties were de-

termined with the Python `tensorboard` library, which is a built-in optimisation module of `tensorflow` designed to execute and compare several configurations of the same ML architecture. In particular, `tensorboard` performs a short training of each configuration to identify the one with the most optimal performance. During *cecilia*’s development phase, we used `tensorboard` to optimise the number of layers, the neurons per layer, the batch size (i.e. the number of training instances processed together during a given iteration), and the learning rate of the Autoencoder and the two FCNNs.

## 4.2 Network Training

To train *cecilia*, we used a partition of MIT’s Satori Supercomputer with one NVidia V100 GPU.<sup>5</sup> First, we randomly split our collection of synthetic labels and spectra into three datasets for training (70%), testing (20%), and validation (10%) purposes. Next, we used the backpropagation technique to optimise the hyperparameters of *cecilia*’s neural networks. To this end, we implemented the Adaptive Moment Estimation (Adam) optimisation algorithm (Kingma & Ba 2017), which is an extension of the commonly used Stochastic Gradient Descent routine (Kiefer & Wolfowitz 1952). Finally, given that the synthetic spectra described in Section 3.1.2 represented a substantial amount of data (a total of about  $10^9$  flux pixels for 22,842 spectra), we chose to train our pipeline by dividing the synthetic spectra in 29 windows of 200 Å, rather than by considering the full wavelength range at once.<sup>6</sup> With this approach, the end-to-end training of *cecilia* took approximately three weeks.

For each spectral window of 200 Å, *cecilia*’s training consisted of three sequential learning phases: an initial one for the Autoencoder, another one for the FCNN1, and a final one for the FT FCNN2. In each phase, the networks employed the same datasets for training, validation, and testing purposes, with each dataset consisting of unique combinations of the 13 randomly generated labels listed in Table 1 (i.e.  $T_{\text{eff}}$ ,  $\log g$ , and 11 elemental abundances from H to Ni). It is important to clarify that the Autoencoder and the FCNN1 were *only* used to improve the learning of the fine-tuned neural network. After concluding *cecilia*’s training, our code only uses the FT FCNN2 to generate spectral models of polluted white dwarfs. Below, we summarise the three training phases in more detail.

- *Phase 1 (Figure 4)*: First, we trained the Autoencoder in order to compress and reconstruct the normalised synthetic spectra based on only 100 latent features. At the end of this process, we used the trained Encoder to transform our full database of 22,824 synthetic spectra into their latent space representations. If  $\mathcal{S}$  and  $\mathcal{H}$  represent, respectively, the vector spaces of the synthetic spectra and the hidden space (see Section 4.1), we can denote the operations of the trained Autoencoder as

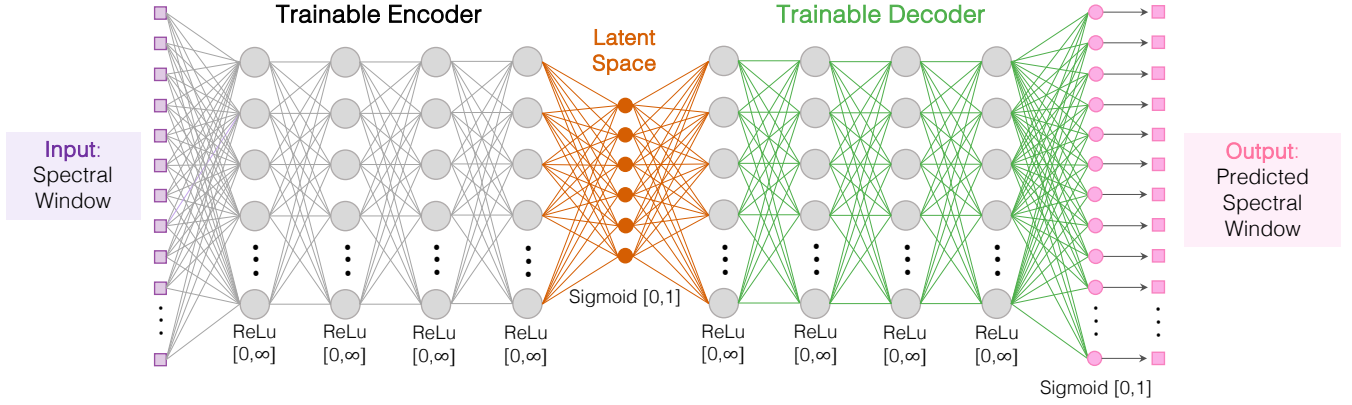
$$\mathcal{S} \rightarrow \mathcal{H} \rightarrow \mathcal{S}. \quad (13)$$

- *Phase 2 (Figure 5)*: Next, we trained the first Fully-Connected Neural Network (FCNN1) to compress the normalised stellar labels into their latent space representations. We then appended the trained Decoder to the output layer of the FCNN1 to convert the latent features of the normalised labels into predictions of normalised synthetic spectra. If  $\mathcal{L}$  is the vector space of the

<sup>5</sup> <https://mit-satori.github.io/>.

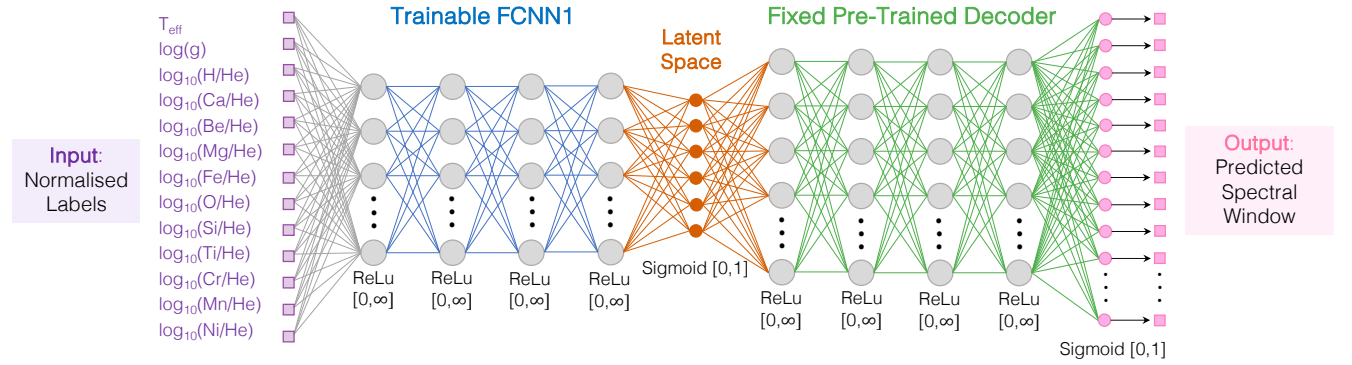
<sup>6</sup> We also evaluated the performance of *cecilia* with narrower and wider spectral windows. From this analysis, we found that 200 Å offered the highest computational efficiency and predictive accuracy.

## 1 Autoencoder



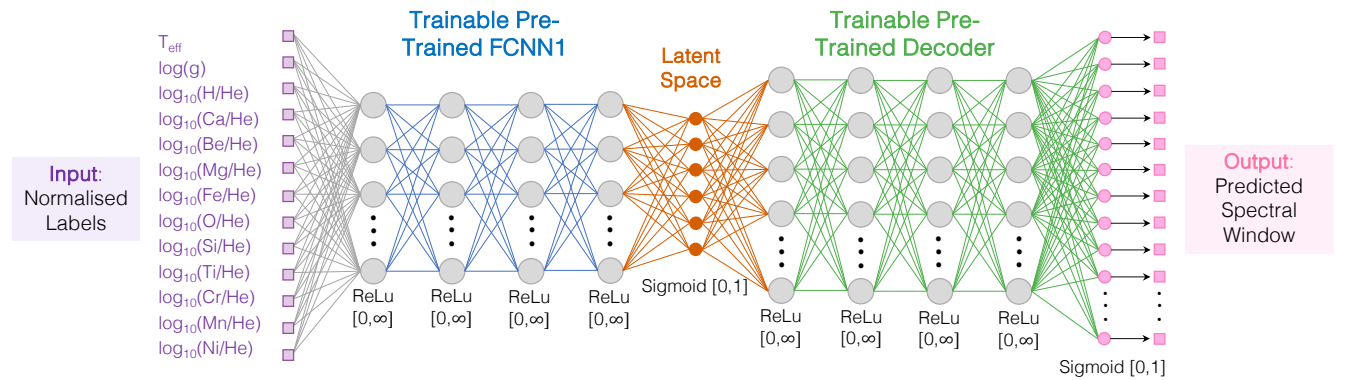
**Figure 4.** The first component of *cecilia*'s ML architecture is an Autoencoder, which consists of two networks: an Encoder (in grey) to compress input data into a small number of latent, hidden features (in orange); and a Decoder (in green) to reconstruct the encoded data from the latent space. In our work, the input and output data of the Autoencoder are spectral windows of 200 Å from our atmosphere models (in purple and pink, respectively).

## 2 FCNN1

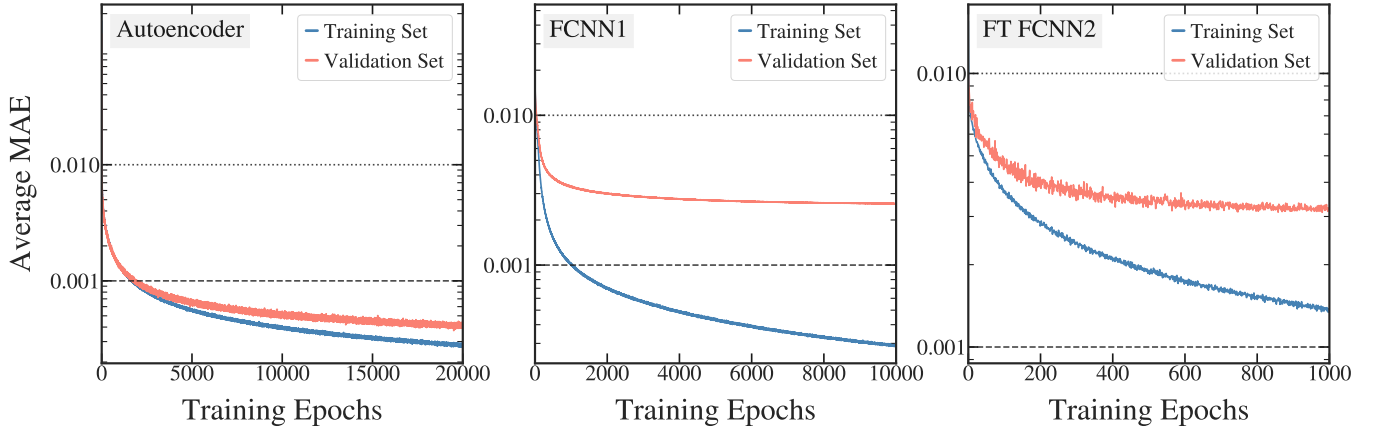


**Figure 5.** The second component of *cecilia*'s ML architecture is a feed-forward Fully Connected Neural Network (FCNN1; in blue). This network leverages the trained Decoder in Figure 4 (in green) to predict the spectrum of a polluted white dwarf (in pink) from its corresponding stellar labels (in purple).

## 3 Fine-Tuning (FT) FCNN2



**Figure 6.** The third component of *cecilia*'s ML architecture is a replica of the FCNN1 aimed at fine-tuning the final predictions of our pipeline (FT FCNN2; in blue). Similarly to the FCNN1 Figure 5, this network is connected to the trained Decoder (in green) to predict a metal-polluted spectrum (in pink) from its stellar labels (in purple).



**Figure 7.** Mean Absolute Error (MAE) statistic (or loss function) for the Autoencoder, FCNN1, and FT FCNN2, averaged for the 29 windows of 200 Å used to train *cecilia*. The training and validation sets are shown in blue and red. The dotted and dashed lines denote an average MAE of 0.01 and 0.001, respectively.

stellar labels, the FCNN1 and trained Decoder are thus in charge of performing the operation:

$$\mathcal{L} \rightarrow \mathcal{H} \rightarrow \mathcal{S}. \quad (14)$$

At this stage, we had two options for training the FCNN1 while keeping the trained Decoder fixed. On the one hand, we could directly determine the errors of the hidden layer with respect to the encoded spectra  $\mathcal{H}$ . This scenario required prior compression of the spectra with the Encoder and involved backpropagating the error through the FCNN1 to correct the weights and biases of the network. On the other hand, we could calculate the errors of the FCNN1's output layer with respect to the true synthetic spectra  $\mathcal{S}$ . This alternative option required backpropagating the errors through the fixed Decoder until reaching the hidden space, and then updating the weights and biases of the FCNN1 accordingly. In other words, the main difference between the two approaches was the location of the network used to calculate the loss function, with the first and second scenarios employing, respectively, the encoded synthetic spectra  $\mathcal{H}$  and the normalised synthetic spectra  $\mathcal{S}$ . In this work, we decided to follow the first approach in order to (i) avoid unnecessary backpropagation steps while traversing the Decoder, and (ii) prevent gradient vanishing effects of the loss function, which could result in a slower and poorer performance.

- *Phase 3 (Figure 6):* Lastly, we sought to improve and fine-tune *cecilia*'s final predictions by training a second Fully-Connected Neural Network (FT FCNN2) connected to the trained Decoder. This architecture replicates the structure of the FCNN1, but it differs from the latter because we allowed the parameters of the Decoder to be trainable. Similarly to Eq. 14, the joint FT FCNN2 and Trainable Decoder architecture performs the operation

$$\mathcal{L} \rightarrow \mathcal{H}^* \rightarrow \mathcal{S}, \quad (15)$$

where the latent space of this third network, represented by the term  $\mathcal{H}^*$  in Eq. 15, is now different than those of the Autoencoder and the FCNN1 ( $\mathcal{H}$ ).

Upon training *cecilia*'s Autoencoder and FCNNs, we calculated the error between the predictions of each network and their respective true synthetic models. To this end, we adopted the Mean Absolute Error (MAE) metric as our loss function—a popular choice in ML regression problems. In this work, we calculate the MAE of a spectral

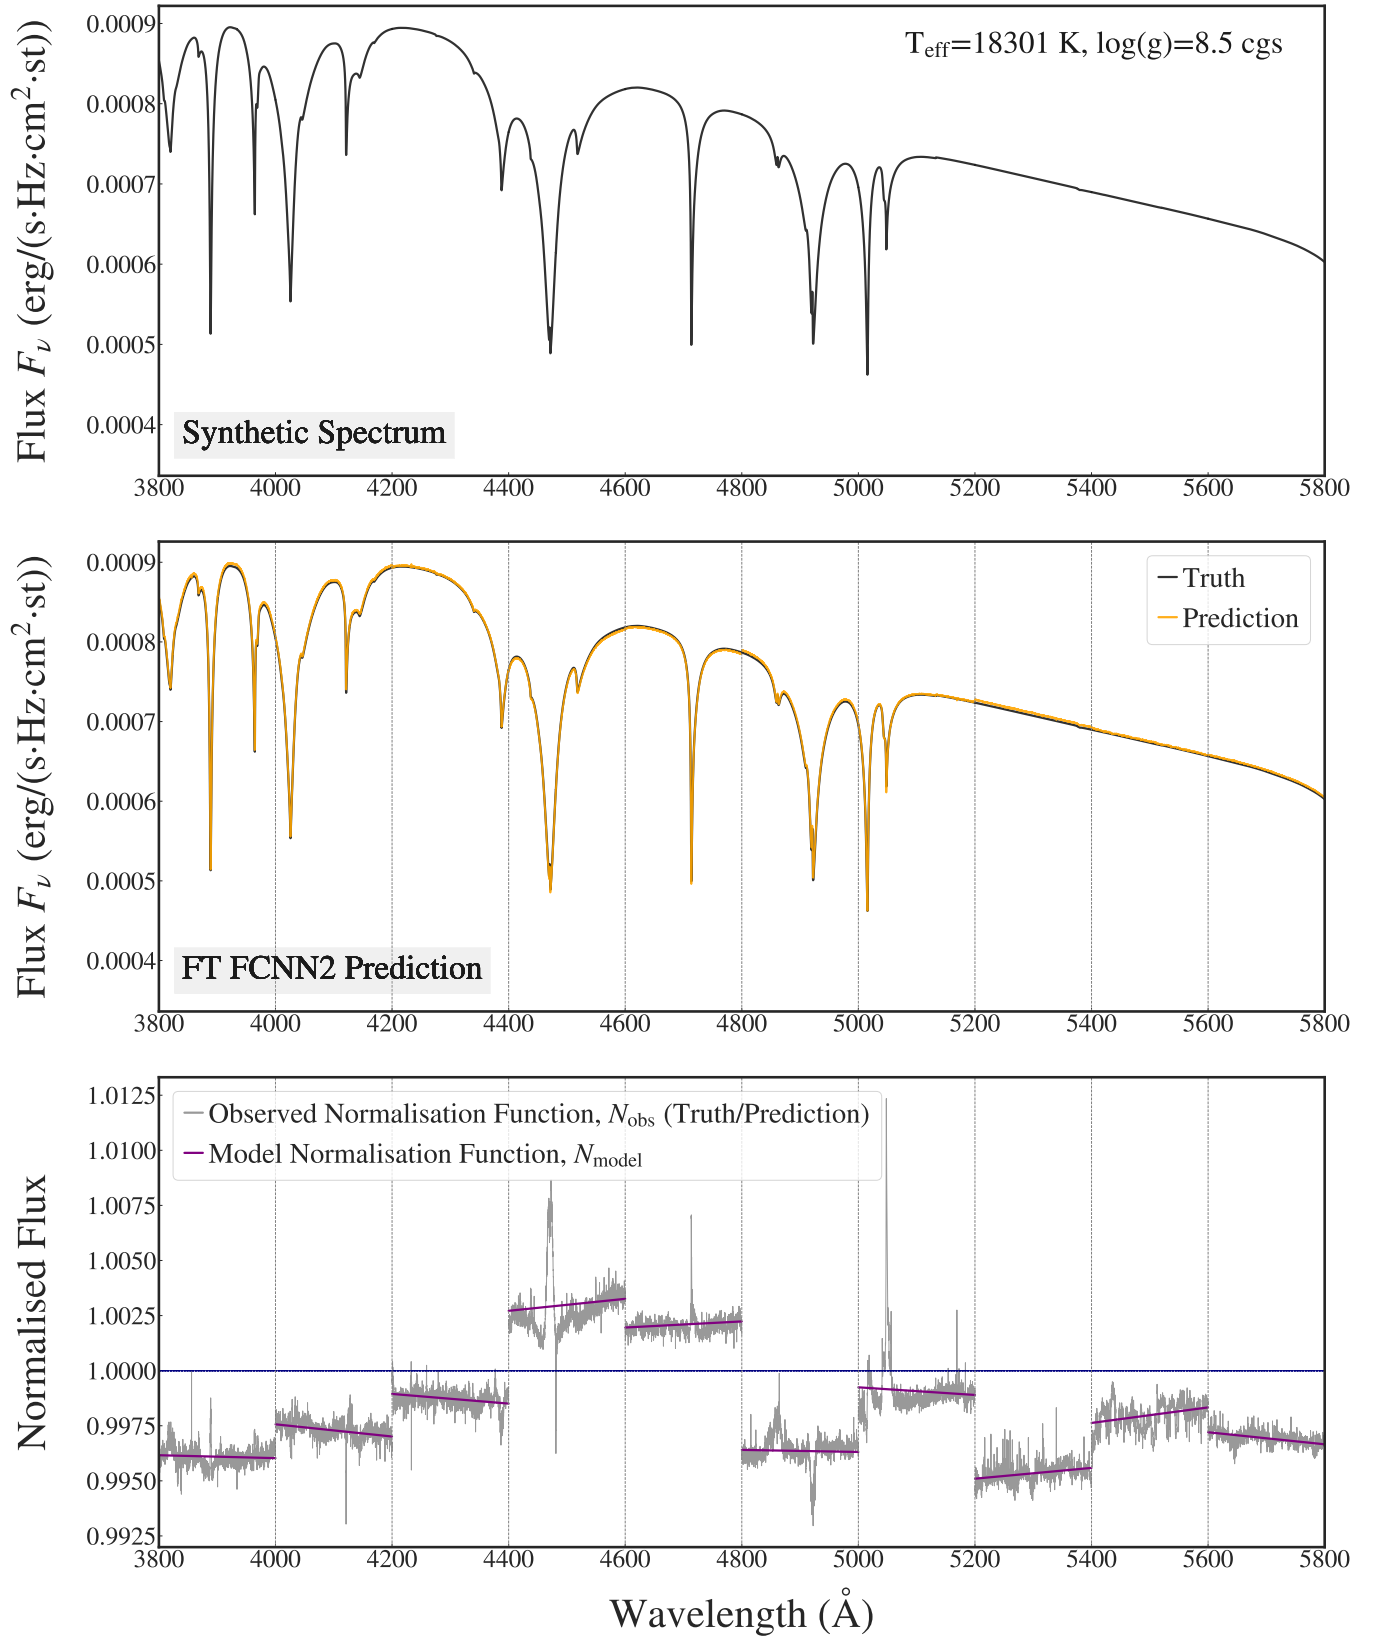
window as

$$\text{MAE} = \frac{1}{QP} \sum_{q=1}^Q \sum_{p=1}^P |y_{q,p} - \hat{y}_{q,p}|, \quad (16)$$

where  $q$  and  $Q$  represent, respectively, the index and the total number of instances—where we recall that an instance is a synthetic spectrum with its corresponding stellar labels—,  $p$  and  $P$  denote the index and total number of output pixels (i.e. the number of flux points in a given spectral window),  $y$  and  $\hat{y}$  correspond to the true and predicted synthetic flux. Figure 7 shows our loss function for each network, with the MAE averaged for the 29 windows in the wavelength range between 3,000 Å and 9,000 Å. From this figure, we find that the training error rates of *cecilia*'s networks decrease with each epoch, as should be the case during training. This behavior demonstrates the ability of *cecilia* to recognise abstract features in the training data and generalise its acquired knowledge to previously unseen observations. At the same time, Figure 7 shows that none of *cecilia*'s networks are suffering from the dangers of underfitting or overfitting:<sup>7</sup> on the one hand, the absence of underfitting can be demonstrated with the convergence of the validation error, which becomes asymptotically horizontal, thus leaving no substantial margin for further generalisation; on the other hand, the problem of overfitting can be discarded because the validation error does not increase over time. In Figure 8, we also present an example of *cecilia*'s predictions for several 200 Å-windows of a synthetic spectrum.

A general note about our training process is that Phases 2 and 3 *re-cycle* the trained Decoder to improve the performance of the FCNN1 and the FT FCNN2. After experimenting with less sophisticated ML architectures, we found this approach resulted in more accurate predictions than a simpler NN designed to predict stellar labels directly

<sup>7</sup> The concept of *underfitting* refers to a neural network that has not managed to learn the fundamental attributes of a training sample. This problem can arise from various reasons—e.g. an excessively simple ML model, a small training set, a short training duration—and it leads to poor predictions of new datasets. In contrast, *overfitting* describes a network that has memorised the underlying features of the training data too well but cannot make accurate predictions with a test dataset. This behaviour can also happen due to multiple reasons, such as an overly complex ML model or an unnecessarily long training process. An optimal ML model represents a balance between generalisation and specificity (Goodfellow et al. 2016).



**Figure 8.** *Top Panel:* Synthetic spectrum corresponding to a white dwarf with  $T_{\text{eff}}=18,301 \text{ K}$  and  $\log g=8.5 \text{ cgs}$ . To allow for a closer examination, we only show the wavelength range between  $3,800 \text{ \AA}$  and  $5,800 \text{ \AA}$ , which features 10 windows of  $200 \text{ \AA}$ . *Middle Top Panel:* Denormalised cecilia FT FCNN2 prediction. *Bottom Top Panel:* Slope correction with a linear least-squares fit. In each panel, we use dashed vertical lines to indicate the start/end of a  $200 \text{ \AA}$  spectral window.

from their spectra without any intermediary steps (i.e.  $L = f(S)$ ). The use of a pre-trained model to improve the training of another network is known as “Transfer Learning” and has been very successful in many scientific fields outside computational science (e.g. Tan et al. 2018).

### 4.3 Parameter Estimation

After training, *cecilia* can use its FT FCNN2 architecture to rapidly ( $<1$  second) generate a full high-resolution synthetic spectrum from 13 stellar labels ( $T_{\text{eff}}$ ,  $\log g$ ,  $\log_{10}(\text{H/He})$ , and 10 metal abundances relative to helium). Exploiting the fast and automated ML interpolation capabilities of our pipeline, we have developed a new fitting framework to simultaneously and accurately measure the main physical and chemical properties of He-rich polluted white dwarfs from their spectra. This section explains our methodology for achieving this goal using a step-by-step approach. In our discussion, we will employ the subscripts “synth” and “obs” to represent *cecilia*’s synthetic predictions and a real white dwarf spectrum, respectively. With this choice of notation, we will refer to their corresponding wavelength, flux, and flux error arrays as  $\lambda_X$ ,  $f_X$ , and  $f_{X,\text{err}}$ , where  $X$  will indicate the nature of the data (i.e. synthetic “synth,” or observed “obs”). Finally, we will denote their resolving power as  $R_X$ , where  $R_X \equiv \lambda_X / \Delta\lambda_X$  (unitless).

Broadly speaking, our fitting framework can be divided into three main phases (see Figure 9). First, *cecilia*’s trained FT FCNN2 architecture generates a preliminary atmosphere model based on the assumed astrophysical properties of the white dwarf. Second, our pipeline optimises this initial prediction with the non-linear least-squares Levenberg-Marquardt algorithm implemented by the fast Python `mpfit`<sup>8</sup> library (Moré 1978; Markwardt 2009). To conclude, we perform a global Bayesian fit with the differential evolution Markov Chain Monte Carlo (MCMC) sampler of the Python `edmc` package (Vanderburg 2021; Ter Braak 2006).<sup>9</sup> More specifically, our optimisation procedure consists of the following phases:

- (i) **Step 1: Input Spectrum and Stellar Parameters:** First, the user uploads the spectrum of their polluted white dwarf ( $\lambda_{\text{obs}}$ ,  $f_{\text{obs}}$ ,  $f_{\text{obs,err}}$ ) into *cecilia*, ensuring that the effective temperature and surface gravity of the star satisfy the ranges  $10,000 \leq T_{\text{eff}} \leq 20,000$  K and  $7 \leq \log g \leq 9$  cgs imposed in Section 3.1.1 (see Table 1). In parallel, the user provides a list of 14 white dwarf properties to our pipeline, either complete with initial guesses, or lacking the latter if unknown. These properties consist of the 13 stellar labels used during *cecilia*’s training,<sup>10</sup> as well as an additional Radial Velocity shift (or RV thereafter) to account for the barycentric motion and gravitational redshift of the white dwarf. If necessary, the RV term can also be used to handle potential wavelength calibration issues in the spectrum.
- (ii) **Step 2: Initial Guesses for Stellar Parameters:** Next, *cecilia*

reviews the user’s white dwarf properties and performs the following operations based on the availability (or absence) of initial guesses for the 13 stellar labels and the RV term:

- **Scenario 1: The user has prior knowledge on all stellar parameters:** In this optimal situation, *cecilia* examines the user’s initial guesses to confirm that they are within the ML bounds listed in Table 1. If the RV shift is not within the range given by  $-500 \leq \text{RV} \leq 500$  km/s, our code automatically sets this parameter to 0 km/s. Alternatively, if any of the abundances are outside of their corresponding ML ranges, *cecilia* attempts to correct this problem by generating a normal distribution centered on the user’s estimate with a scale of 0.1 dex, and subsequently adopting a random value from this distribution as an initial guess. If this modification is still insufficient, *cecilia*’s Gaussian estimate is then substituted by a chondritic abundance ratio based on the assumed  $\log_{10}(\text{Ca/He})$  of the white dwarf and the chondritic values of Asplund et al. (2009).
- **Scenario 2: The user has partial knowledge of some stellar parameters:** In this intermediate scenario, our code starts by verifying that the user’s initial guesses for  $T_{\text{eff}}$ ,  $\log g$ ,  $\log_{10}(\text{H/He})$  and  $\log_{10}(\text{Ca/He})$  are within *cecilia*’s allowed bounds in Table 1. If any of these parameters are unknown, *cecilia* uses the mean of their ML bounds as initial guesses for its optimisation routine (i.e.  $\bar{T}_{\text{eff}}=15,000$  K,  $\bar{\log g}=8$  cgs,  $\bar{\log(\text{H/He})}=-5$ , and  $\bar{\log(\text{Ca/He})}=-9.5$ ). For the remaining white dwarf parameters, *cecilia* adopts a zero RV shift and assumes chondritic abundances for Mg, Fe, O, Si, Ti, Be, Cr, Mn, Ni based on the estimated  $\log_{10}(\text{Ca/He})$  of the white dwarf and the chondritic ratios of Asplund et al. (2009).
- **Scenario 3: The user has no knowledge of the stellar parameters:** In this worst-case scenario, *cecilia* generates preliminary estimates of  $T_{\text{eff}}$ ,  $\log g$ ,  $\log_{10}(\text{H/He})$  and  $\log_{10}(\text{Ca/He})$  by taking the mean of their corresponding ML bounds, as in Scenario 2 above. For the remaining white dwarf properties, it also follows the same procedure described in Scenario 2, that is: it assumes a RV shift of 0 km/s and chondritic abundance ratios for those elements lacking a user-defined initial guess.
- (iii) **Step 3: Definition of *cecilia*’s Spectral Model:** After loading the observed spectrum into *cecilia* together with (adjusted) preliminary estimates of the 14 white dwarf parameters, the user is asked to choose the free and fixed parameters of their *cecilia* spectral model. If certain parameters are fixed, our code freezes their initial value to the user’s guess or, if not provided, to *cecilia*’s estimated value.
- (iv) **Step 4: Adjusting the Format of the Input Data:** The last step before initiating our fitting procedure is to adjust the white dwarf spectrum and its stellar labels to a format that *cecilia* can understand. To this end, our code automatically implements the following data processing techniques:

- **Normalisation of Initial Guess Labels:** First, it normalises the guess stellar labels using the minimum and maximum values of the synthetic labels considered during training (see Section 3.2). Given that *cecilia* was trained in normalised space, this step is crucial to ensure that our code can effectively process and interpret the user’s input data.
- **Check of Spectral Units:** Next, *cecilia* verifies that the observed stellar flux is expressed in terms of the calculated flux at the surface of the white dwarf, i.e.  $F_v$ ,

<sup>8</sup> <https://github.com/segasai/astrolibpy/blob/master/mpfit/mpfit.py>

<sup>9</sup> Differential evolution is a genetic algorithm designed to find an optimal distance and direction for a jumping distribution from one MCMC chain to another.

<sup>10</sup> A future improvement to *cecilia* will be to include a trained neural network to estimate the effective temperature and surface gravity of the polluted white dwarf from existing photometric observations (see Section 6). In the meantime, we assume that the user has good external constraints on  $T_{\text{eff}}$  and  $\log g$ .

where  $[F_\nu] = \text{erg}/(\text{s} \cdot \text{Hz} \cdot \text{cm}^2 \cdot \text{sr})$ . Although this check is not of paramount importance due to the slope and offset correction discussed in step (vi) below,  $F_\nu$  is the unit of choice of *cecilia*'s atmosphere models (e.g. see Figure 8) and is therefore the preferred notation of our pipeline. In reality, however, real spectra of white dwarfs are typically not expressed in units of  $F_\nu$ , but in  $f_\lambda$ , where  $[f_\lambda] = \text{erg}/(\text{s} \cdot \text{cm}^2 \cdot \text{\AA})$ . As a consequence, our code integrates the possibility of transforming the observed stellar flux from  $f_\lambda$  to  $F_\nu$  with

$$F_\nu = \left( \frac{f_\lambda \cdot w_{\text{obs}}^2}{c} \right) \cdot \left[ \frac{1}{4} \left( \frac{d}{r} \right)^2 \right], \quad (17)$$

where  $c$  is the speed of light (in  $\text{\AA}/\text{s}$ ),  $r$  is the radius of the white dwarf,  $d$  is its distance from Earth, and the second term represents the inverse solid angle of the white dwarf (in  $1/\text{sr}$ ). In Section 5, we use the *Gaia* DR3 database (Gaia Collaboration et al. 2016, 2023) to determine the values of  $r$  and  $d$ .

- **Spectrum Cropping:** Lastly, our code examines the observed wavelength array to make sure that it falls within *cecilia*'s spectral coverage between 3,000  $\text{\AA}$  and 9,000  $\text{\AA}$ . Any regions outside these boundaries are automatically excluded from our fitting analysis (e.g. see Figure 12).

(v) **Step 5: *cecilia*'s High-Resolution Spectral Prediction:**

Next, our code feeds the normalised stellar labels into *cecilia*'s trained FT FCNN2 architecture to quickly generate a normalised synthetic spectrum for the white dwarf ( $\lambda_{\text{synth}}, f_{\text{synth}}$ ). At this stage, *cecilia*'s synthetic spectrum shares the resolving power of the atmosphere models described in Section 3.1.2, i.e.  $R_{\text{synth}} \approx 50,000$ . Therefore, assuming that the resolving power of the observed white dwarf spectrum is lower than that of our atmosphere models ( $R_{\text{obs}} \ll R_{\text{synth}}$ ), *cecilia*'s prediction cannot be optimised yet.

(vi) **Step 6: Processing of *cecilia*'s High-Resolution Spectral Prediction:**

To enable a meaningful comparison between the output of *cecilia*'s FT FCNN2 architecture and the observed input spectrum, our code performs the following automated and sequential operations:

- **Denormalisation of *cecilia*'s Synthetic Prediction:** First, it denormalises *cecilia*'s prediction with the minimum and maximum flux values of the synthetic models used during training. This transformation brings *cecilia*'s synthetic spectrum to the non-normalised space of the observed white dwarf spectrum.
- **Smearing and Resampling:** Second, our code smears and resamples *cecilia*'s denormalised prediction to the resolving power ( $R_{\text{obs}}$ ) and wavelength grid ( $\lambda_{\text{obs}}$ ) of the input spectrum, respectively. To achieve this, it downgrades *cecilia*'s synthetic spectrum from  $R_{\text{synth}}$  to  $R_{\text{obs}}$  by convolution with a Gaussian smoothing function. Then, it resamples *cecilia*'s prediction by smearing the latter onto the wavelength grid of the observed spectrum using Python's one-dimensional linear interpolator `numpy.interp`.
- **Slope and Offset Correction:** Third, we correct for any spectral "jumps" in *cecilia*'s prediction caused by the training of our ML pipeline in independent windows of 200  $\text{\AA}$  (e.g. see Figure 8). For this task, we start by defining an *observed* normalisation function  $N_{\text{obs}}$  from the ratio between *cecilia*'s prediction and the observed flux,

$$N_{\text{obs}}^i = \frac{f_{\text{synth}}^i}{f_{\text{obs}}^i}, \quad (18)$$

where  $f_{\text{synth}}^i$  is *cecilia*'s synthetic flux at wavelength  $i$ , and  $f_{\text{obs}}^i$  is the observed stellar flux. Then, we use a weighted linear least-squares algorithm to fit  $N_{\text{obs}}$  as a function of the observed wavelength array,  $\lambda_{\text{obs}}$ . From this analysis, we obtain a *model* normalisation function  $N_{\text{model}}$  (e.g. see purple line in the second panel of Figure 8), which we then multiply to *cecilia*'s prediction to minimise the jumps between consecutive windows. We note that *cecilia* can do a slope and offset correction with an  $n$ -th degree polynomial as well.

- **Radial Velocity Shift:** Finally, our code applies a radial velocity shift to *cecilia*'s denormalised, smeared, resampled, and slope-corrected prediction. In particular, if  $c$  is the speed of light, and  $v_{\text{obs}}$  is the assumed RV shift of the white dwarf, the total wavelength shift of the spectral lines can be calculated with

$$\delta\lambda_{\text{synth}} = \frac{\lambda_{\text{synth}} \cdot v_{\text{obs}}}{c} \equiv \lambda'_{\text{synth}} - \lambda_{\text{synth}}. \quad (19)$$

From the above, we determine the shifted wavelength array of *cecilia*'s synthetic prediction with

$$\lambda'_{\text{synth}} = \lambda_{\text{synth}} \cdot \left( 1 + \frac{v_{\text{obs}}}{c} \right). \quad (20)$$

Using this new wavelength grid, we then perform a cubic spline interpolation to obtain the shifted flux and flux error arrays (Eq. 21). In the absence of flux points within a certain wavelength range, our function simply extrapolates the data. Therefore, for a shifted wavelength array  $\lambda'_{\text{synth}}$ , the final products of our interpolation are:

$$\begin{aligned} f_{\text{synth}}(\lambda_{\text{synth}}) &\rightarrow f'_{\text{synth}}(\lambda'_{\text{synth}}) \\ f_{\text{err,synth}}(\lambda_{\text{synth}}) &\rightarrow f'_{\text{err,synth}}(\lambda'_{\text{synth}}) \end{aligned} \quad (21)$$

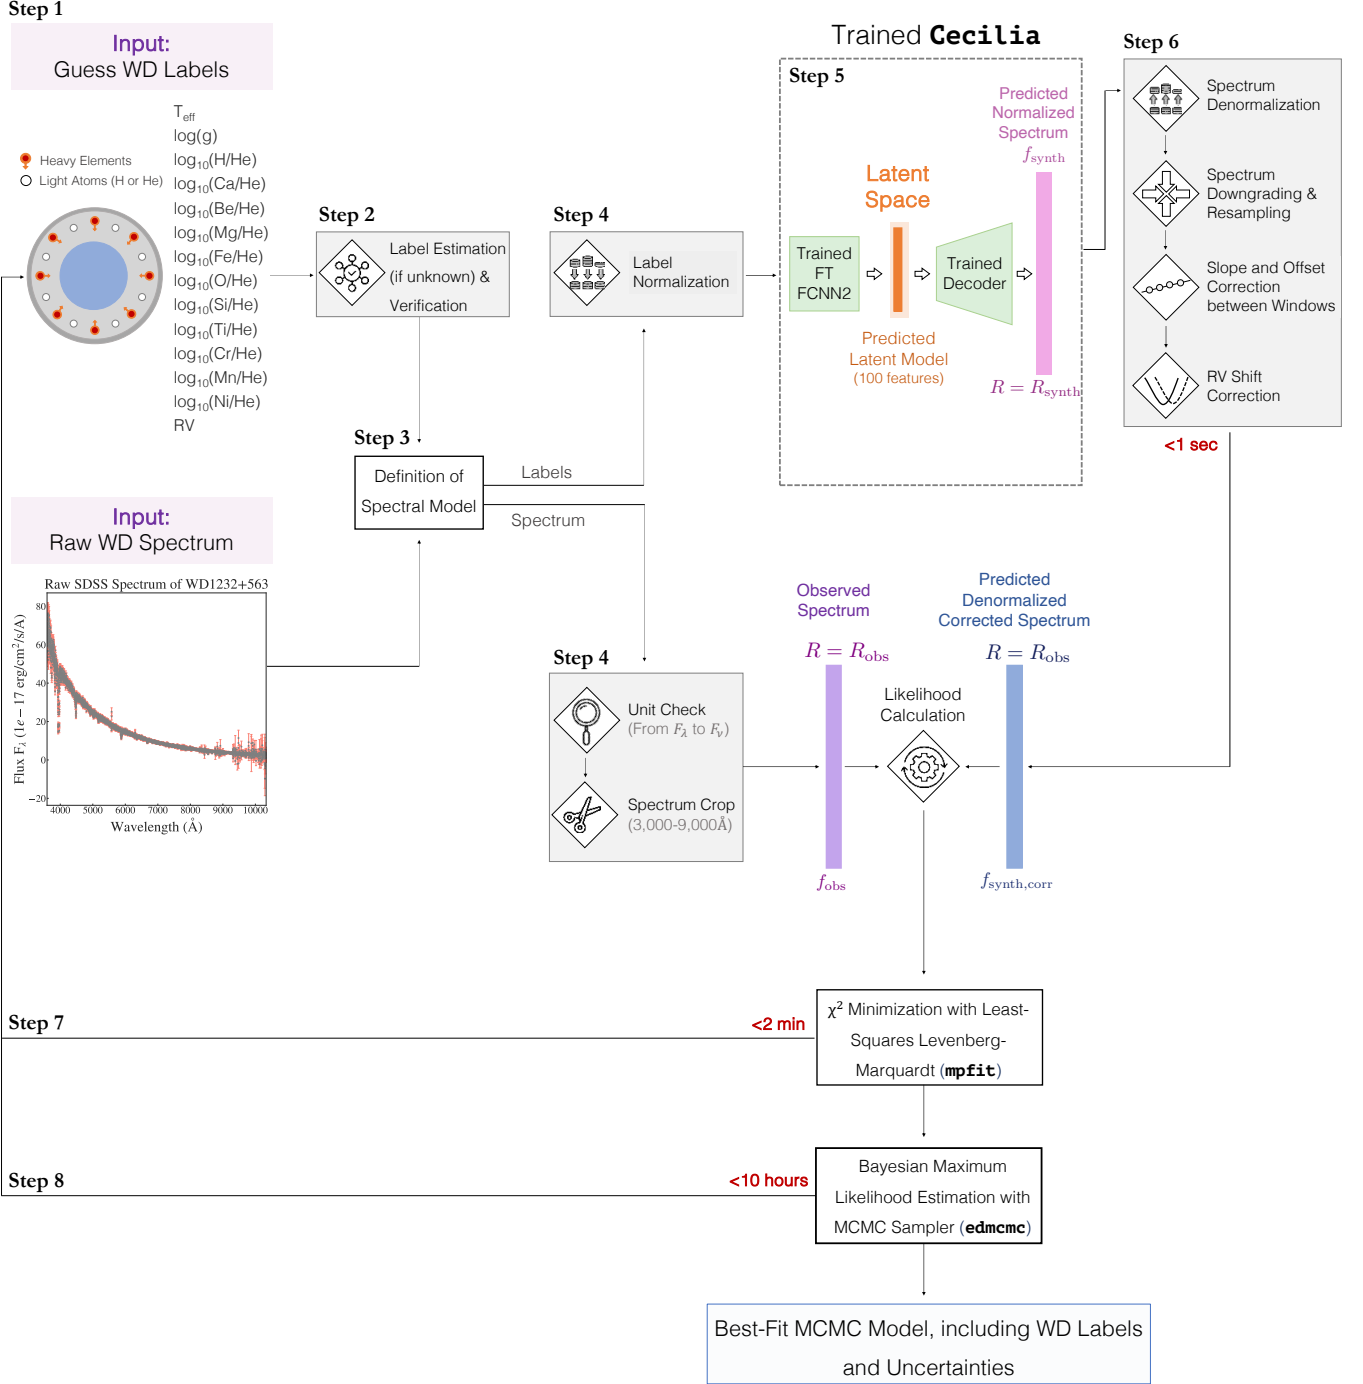
In the remaining discussion, we will denote *cecilia*'s smeared, resampled, slope-corrected, and RV-shifted synthetic wavelength and flux as  $w_{\text{synth,corr}}$  and  $f_{\text{synth,corr}}$ , respectively.

- (vii) **Step 7: Initial `mpfit`  $\chi^2$  Optimisation:** After concluding steps (i)-(vi), our code initiates the first phase of our optimisation procedure, namely: a preliminary non-linear least-squares fit implemented with the `mpfit` package. The goal of `mpfit` is to quickly ( $< 2$  minutes) identify good values for the model parameters which can then be used as initial guesses for our MCMC. To achieve this, `mpfit` minimises the chi-squared statistic ( $\chi^2$ ) of *cecilia*'s corrected prediction,

$$\chi^2 = \sum_{i=1}^{N_{\text{points}}} \left[ \frac{f_{\text{obs}}^i - f_{\text{synth,corr}}^i}{f_{\text{obs,err}}^i} \right]^2, \quad (22)$$

where  $N_{\text{points}}$  is the total number of points in the observed input spectrum,  $f_{\text{synth,corr}}^i$  is *cecilia*'s corrected synthetic flux at wavelength  $i$ ,  $f_{\text{obs}}^i$  is the observed stellar flux, and  $f_{\text{obs,err}}^i$  is the uncertainty of the latter.

As illustrated in Figure 9, `mpfit` computes the  $\chi^2$  metric in Eq. 22 as follows: first, it provides the normalised stellar labels to *cecilia*, which then invokes its trained FT FCNN2 architecture to rapidly generate a preliminary normalised, high-resolution ( $R_{\text{synth}}$ ), synthetic spectrum ( $\lambda_{\text{synth}}, f_{\text{synth}}$ ). Next, *cecilia*'s prediction is adjusted with the techniques described in step (vi), which allows `mpfit` to compare the observed input spectrum to



**Figure 9.** A schematic view of our fitting methodology using *cecilia* and a combination of frequentist and Bayesian statistical techniques. First, the user uploads the spectrum of their polluted white dwarf onto *cecilia*, together with a list of 14 stellar properties ( $T_{\text{eff}}$ ,  $\log g$ , 11 elemental abundances, and a RV shift; see step (i) in Section 4.3). If these properties are unknown, *cecilia* estimates initial guesses as explained in step (ii); afterwards, it verifies that the labels are within *cecilia*'s ML bounds. Next, the user defines the free and fixed parameters of their spectral model (step (iii)). Then, our code adapts the input spectrum and stellar properties to *cecilia*'s format (step (iv)); this includes normalising the guess labels and excluding any regions of the observed spectrum outside of *cecilia*'s coverage between 3,000 Å to 9,000 Å. After this data processing step, our code feeds the normalised labels to the trained FT FCNN2 network, which then takes <1 second to generate a normalised, high-resolution synthetic spectrum (step (v)). This prediction is subsequently denormalised and adjusted to the properties of the input spectrum (step (vi)) before being optimised with two techniques: a non-linear least-squares  $\chi^2$  minimisation algorithm (step (vii)), and a Bayesian MCMC (step (viii)). Using 1 GPU from the MIT Satori Supercomputer, our fitting procedure produces a full spectroscopic solution in less than about 10 hours, including the stellar labels of the polluted white dwarf with their corresponding uncertainties. Source of icons: *Flaticon*.

cecilia’s corrected prediction. From this comparison, `mpfit` calculates the  $\chi^2$  metric with Eq. 22 and repeats the entire optimisation procedure until identifying a set of optimal model parameters. For these best-fit labels, `mpfit` also provides their covariance matrix, which can then be used to derive initial estimates of their uncertainties.

- (viii) **Step 8: Final MCMC Exploration** The second phase of our fitting routine employs the best-fit results of `mpfit` to initialise an MCMC sampler and obtain more robust parameter uncertainties. In contrast to `mpfit`, which uses frequentist (or “classical”) statistics to minimise the  $\chi^2$  value of *cecilia*’s predictions and determine best-fit values with fixed error bars, an MCMC relies on Bayesian inference to maximise the likelihood  $\mathcal{L}$  of the model parameters and evaluate their full posterior probability distributions. These “posteriors” are proportional to the information contained in the observations as well as to any pre-conceived beliefs or assumptions about the model parameters (or “priors”). Mathematically, the posterior probability ( $\pi$ ) of observing the model parameters ( $\theta$ ) given an observed spectrum ( $f_{\text{obs}}$ ) is encoded in Bayes’ Theorem

$$\pi(\theta|f_{\text{obs}}) \propto \mathcal{L}(f_{\text{obs}}|\theta) \cdot \pi_0(\theta), \quad (23)$$

where  $\mathcal{L}$  is the likelihood function, and  $\pi_0(\theta)$  is the prior probability distribution given an array of model parameters  $\theta$ . Assuming that the likelihood of *cecilia*’s model parameters can be modelled with a Gaussian distribution, our MCMC is designed to explore and map the multidimensional likelihood function, given by

$$\begin{aligned} \ln(\mathcal{L}) = & -\frac{1}{2} \sum_{i=1}^{N_{\text{points}}} \left[ \frac{f_{\text{obs}}^i - f_{\text{synth,corr}}^i}{f_{\text{obs,err}}^i} \right]^2 + \\ & + \sum_{j=1} \pi_{0,\text{phot}_j} + \sum_{k=1} \pi_{0,\text{chondr}_k}. \end{aligned} \quad (24)$$

Above, we adopt the logarithmic scale for numerical and computational efficiency. In particular, the use of the log-likelihood is driven by three main factors: first, it simplifies and speeds up the calculation of derivatives; second, it penalises unlikely fits by assigning them a very high  $\chi^2$  value and thus a very small log-likelihood; and third, it does not entail any loss of information during the fit because both  $\mathcal{L}$  and  $\ln(\mathcal{L})$  have a maximum in the same location.

In Eq. 24, the last two terms represent the prior distributions in our *cecilia* model. The first term ( $\sum_{j=1} \pi_{0,\text{phot}_j}$ ) assumes that the user has reasonable estimates for the effective temperature and surface gravity of the white dwarf from an external photometric fit. The second term ( $\sum_{k=1} \pi_{0,\text{chondr}_k}$ ) is optional and can be used to impose chondritic abundance priors on any metal of interest based on the assumed calcium abundance of the star and the chondritic values of [Asplund et al. \(2009\)](#). These physically motivated priors are useful to constrain the abundances of those heavy elements that are barely visible or undetectable in the observed white dwarf spectrum.

After defining the prior distributions of our likelihood function, we execute our MCMC following the same structure presented in step (vii); this time, however, we do not minimise the  $\chi^2$  statistic of *cecilia*’s predictions, but maximise the log-likelihood function  $\ln(\mathcal{L})$ . To run our MCMC, we use an ensemble of  $n_{\text{walkers}}$  walkers and  $n_{\text{draws}}$  draws, discarding the first 20% of the draws in a “burn-in” phase to exclude non-stationary

values.<sup>11</sup> With 1 Satori GPU and suitable combinations of these three MCMC hyperparameters, we find that *cecilia* can produce an optimal spectroscopic model in less than 10 hours, with the time complexity of this process depending primarily on the computational time associated to the predictions of the FT FCNN2, rather than the quality and/or resolution of the observed input spectrum. In addition to providing a best-fit spectral solution, our MCMC also determines the conditional probabilities of the model parameters, also known as “posterior probability distributions.” From these posteriors, the MCMC computes the median values of each label, together with their 16th, 50th, and 84th percentiles. In this work, we report our MCMC results using the median statistic and its  $1\sigma$  confidence interval.

At the end of the MCMC, we also examine the quality of *cecilia*’s fit with a variety of diagnostic figures and performance metrics. Among them, we generate scatterplot matrices (or “corner plots”) to visualise the two-dimensional “marginal” posterior of a given pair of model parameters alongside their one-dimensional marginal histogram projection. For a parameter  $\theta_i$ , the marginal posterior  $\pi(\theta_i|f_{\text{obs}})$  can be computed with

$$\pi(\theta_i|f_{\text{obs}}) = \int \pi(\theta|f_{\text{obs}}) d\theta_{j \neq i}, \quad (25)$$

where  $j$  iterates over all the parameters. Finally, we test the convergence of the chains using the Gelman-Rubin (GR) potential scale reduction factor  $\hat{R}$  for each model parameter ([Gelman & Rubin 1992](#)). This metric compares the within-chain-variance with the scatter of the points between the chains. In this work, we assume convergence of a model parameter if its GR value is lower than 1.003 ([Vehtari et al. 2021](#); [Vats & Knudson 2018](#)).

## 5 RESULTS

In this section, we investigate the performance of *cecilia* using two types of observations: noiseless synthetic models, and a real spectrum of a well-characterised polluted white dwarf. By testing *cecilia* under ideal and realistic noisy conditions, we aim to understand its retrieval accuracy, its usage limits, and its suitability for the study of different white dwarfs and astronomical datasets.

### 5.1 Sensitivity Analysis with Synthetic Observations

One of the main disadvantages of ML techniques is the lack of their explainability, i.e. the difficulty of interpreting their predictions through a human lens. Despite notable progress in recent years to make ML models more transparent and easier to understand ([Saranya & Subhashini 2023](#); [Vilone & Longo 2020](#); [Gunning et al. 2019](#)), neural networks are still perceived as “black boxes,” which can sometimes hinder our ability to trust their predictions and their decision-making process. In the context of *cecilia*, we have sought to address this problem by evaluating the baseline performance of its trained architecture under well-defined and optimal conditions.

To conduct our study, we selected 1000 random synthetic spectra from our test set; in other words, we only considered atmosphere models that *cecilia* had *never* seen during training. We then used *cecilia*’s FT FCNN2 architecture and the `mpfit` library to model

<sup>11</sup> The walkers and draws represent, respectively, the number of chains in our MCMC, and the number of steps that are evaluated in each chain.

**Table 3.** Abundance accuracy of *cecilia* as determined from an analysis of 1,000 noiseless synthetic spectra from the test set (see Section 5.1). We also show *cecilia*’s allowed ML bounds as well as the abundance ranges detected in He-rich polluted white dwarfs with effective temperatures between  $10,000 \leq T_{\text{eff}} \leq 20,000$  K (source: default column in the Montreal White Dwarf Database, MWDD; Dufour et al. 2016). The dagger symbol ( $\dagger$ ) indicates that *cecilia*’s performance is stable over the entire ML range of the parameter, while the star symbol ( $\star$ ) is used to highlight that *cecilia* cannot constrain Be to within 0.2 dex for synthetic data, hence suggesting a potentially worse performance for real, noisy observations. References: [a]: Genest-Beaulieu & Bergeron (2019), [b]: Coutu et al. (2019), [c]: Wolff et al. (2002), [d]: Xu et al. (2013), [e]: Koester et al. (2005), [f]: Xu et al. (2016), [g]: Fortin-Archambault et al. (2020), [h]: Desharnais et al. (2008). Last MWDD Access: November 2023.

Label	Accuracy Bounds		<i>cecilia</i> ’s ML Bounds		MWDD	
	$\leq 0.1$ dex	$\leq 0.2$ dex	Min.	Max.	Min.	Max.
$\log_{10} (\text{H/He})$	$\geq -7^\dagger$	$\geq -7^\dagger$	-7	-3	$< -6.5^a$	$-2.9^e$
$\log_{10} (\text{Ca/He})$	$\geq -11.20$	$\ll -12^\dagger$	-12	-7	$-10.8^b$	$-6.6^f$
$\log_{10} (\text{Mg/He})$	$\geq -7.85$	$\geq -9.30$	-16.89	-0.17	$-8.5^c$	$-5.9^g$
$\log_{10} (\text{Fe/He})$	$\geq -8.35$	$\geq -11.20$	-18.20	0.18	$-8.2^c$	$-5.6^g$
$\log_{10} (\text{O/He})$	$\geq -5.80$	$\geq -7.15$	-17.46	1.25	$< -6.6^c$	$-5.1^g$
$\log_{10} (\text{Si/He})$	$\geq -7.40$	$\geq -8.65$	-17.42	-0.51	$-8.0^c$	$-5.9^g$
$\log_{10} (\text{Ti/He})$	$\geq -9.80$	$\geq -11.45$	-19.35	-2.32	$-9.5^d$	$-8.6^g$
$\log_{10} (\text{Be/He})$	$\gg -5.61^\star$		-23.85	-5.61	-	-
$\log_{10} (\text{Cr/He})$	$\geq -9.20$	$\geq -10.60$	-19.61	-1.71	$-9.3^d$	$< -7.7^h$
$\log_{10} (\text{Mn/He})$	$\geq -8.25$	$\geq -9.80$	-20.08	-1.65	$-9.8^d$	$-8.6^g$
$\log_{10} (\text{Ni/He})$	$\geq -6.60$	$\geq -8.10$	-18.96	-1.66	$-8.8^d$	$-7.0^g$

the spectra. For this task, we fixed  $T_{\text{eff}}$  and  $\log g$  to their true values,<sup>12</sup> and only fitted the 11 elemental abundances listed in Table 1 (i.e. from  $\log_{10} (\text{H/He})$  to  $\log_{10} (\text{Ni/He})$ ). To better understand the retrieval accuracy of *cecilia* under perfect conditions, we also decided to adopt the true stellar labels of each synthetic spectrum as our input guesses for our optimisation routine. We then executed *mpfit* with 30 iterations, obtaining a best-fit model in about 1 minute per spectrum. For computational efficiency purposes, we did not complement our  $\chi^2$  minimisation analysis with a full MCMC, and instead treated the *mpfit*’s results as our final solutions.

After performing our fits, we determined the residuals  $\vec{r}$  between their true labels and *cecilia*’s predictions. Our results are presented in Figure 10 as a function of photospheric composition, with the colorbar indicating the true effective temperature of the synthetic spectra. Recognising that the visibility of a metal is strongly dependent on  $T_{\text{eff}}$ , the general trend observed in Figure 10 is that *cecilia* performs very well for synthetic white dwarfs with high levels of metal pollution. As metal abundances decrease, the accuracy of our code is also lower because spectral signatures become less strong and harder for *cecilia* to mimic (typically less than roughly 1% with respect to the flux continuum, except for a few outliers). In practice, this problem is mitigated by the fact that subtle, shallow lines are challenging to detect observationally as well.

In Figure 10, we show all the independently varied elements during *cecilia*’s training except for Beryllium, which our code could not reliably constrain. This light metal has a relatively low abundance in the present-day solar photosphere and in meteoritic CI chondrites (Asplund et al. 2009; Lodders 2003), which could also be the case for most polluted white dwarfs. With only two prominent resonance lines of Be in the near ultraviolet at 3,130.42 Å and 3,131.07 Å (Kramida et al. 2022), its abundance is challenging to determine unless there is a substantial amount of this element in the photosphere of the white dwarf (e.g. Klein et al. 2021). Therefore, given the difficulty

of *cecilia* to model Be, we recommend fixing its abundance to a chondritic ratio before using *cecilia* to fit the spectrum of a polluted white dwarf. In the remaining sections, we follow this approach to generate our results.

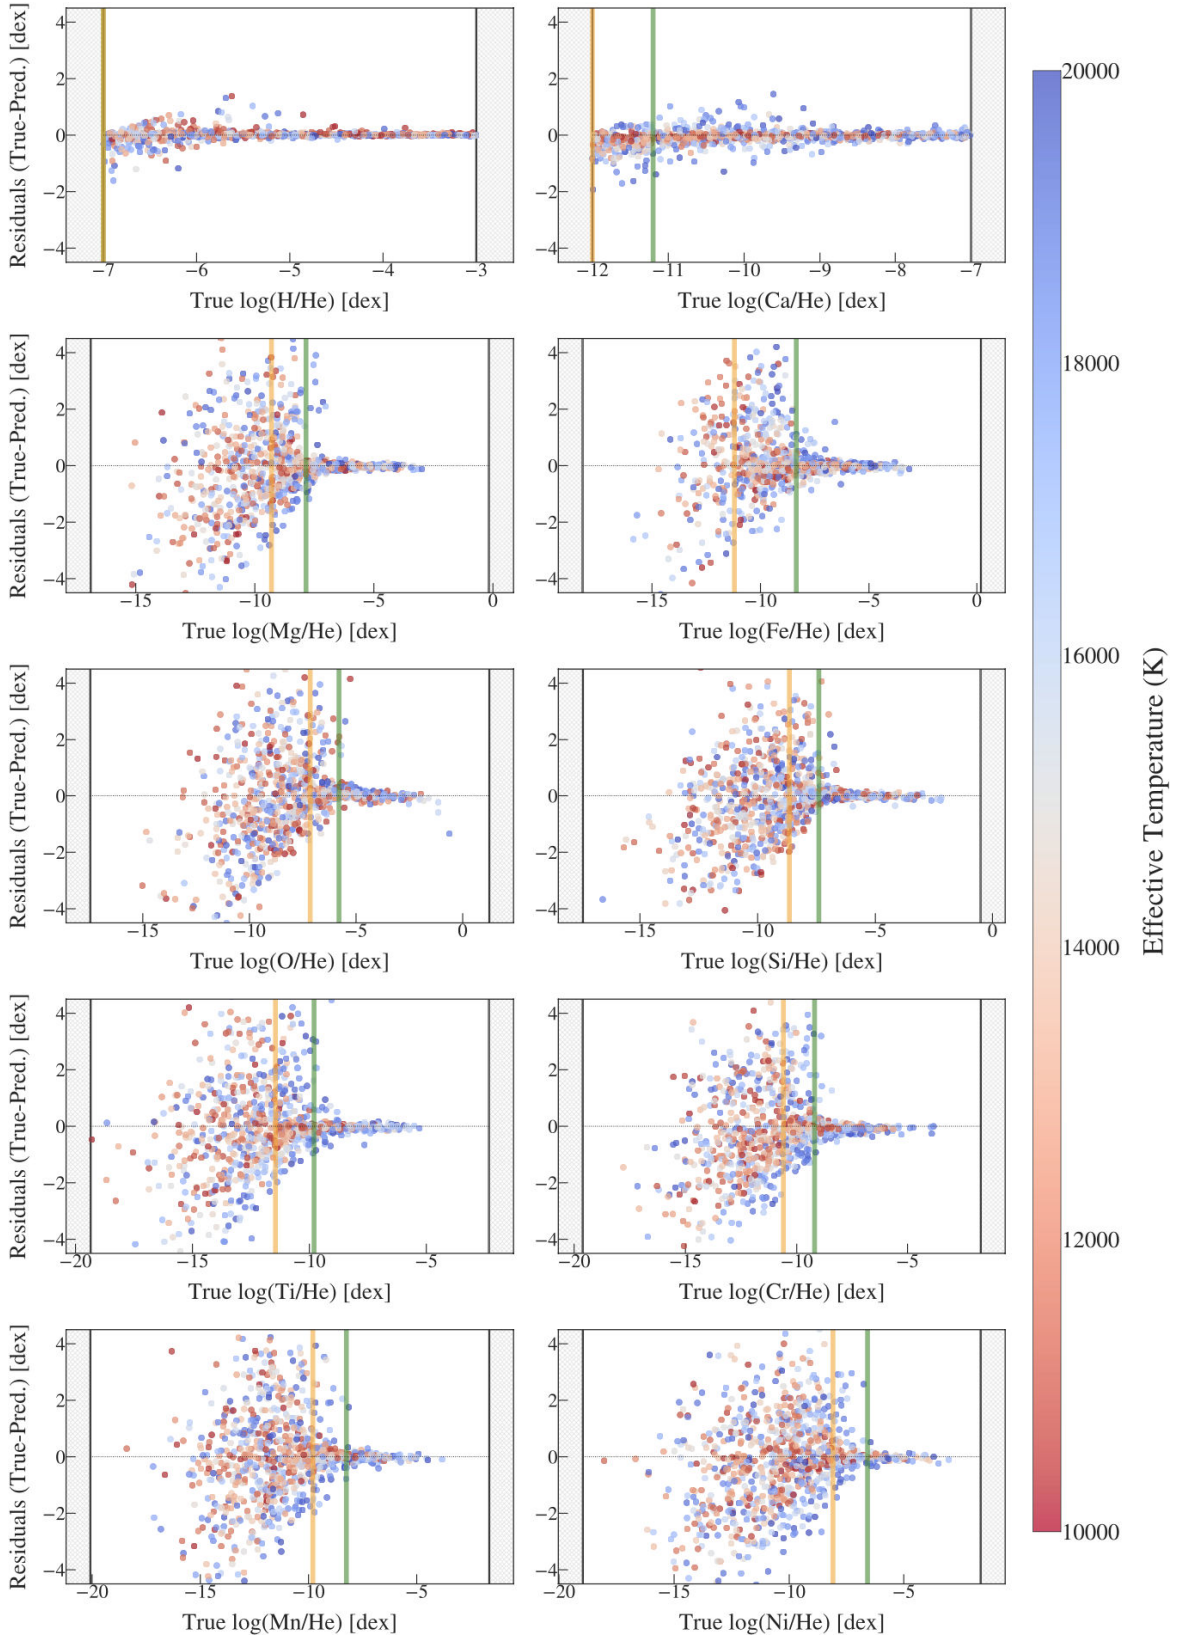
Using the residuals from Figure 10, we then calculated the systematic errors associated to *cecilia*’s *mpfit* predictions. To this end, we estimated the retrieval accuracy of our pipeline from the ratio between the Median Absolute Deviation (MAD) of the residuals and the median of a standard Gaussian distribution,

$$\text{Accuracy} = \frac{\text{median}(|r_i - \text{median}(r)|)}{0.67} \equiv \frac{\text{MAD}}{0.67}, \quad (26)$$

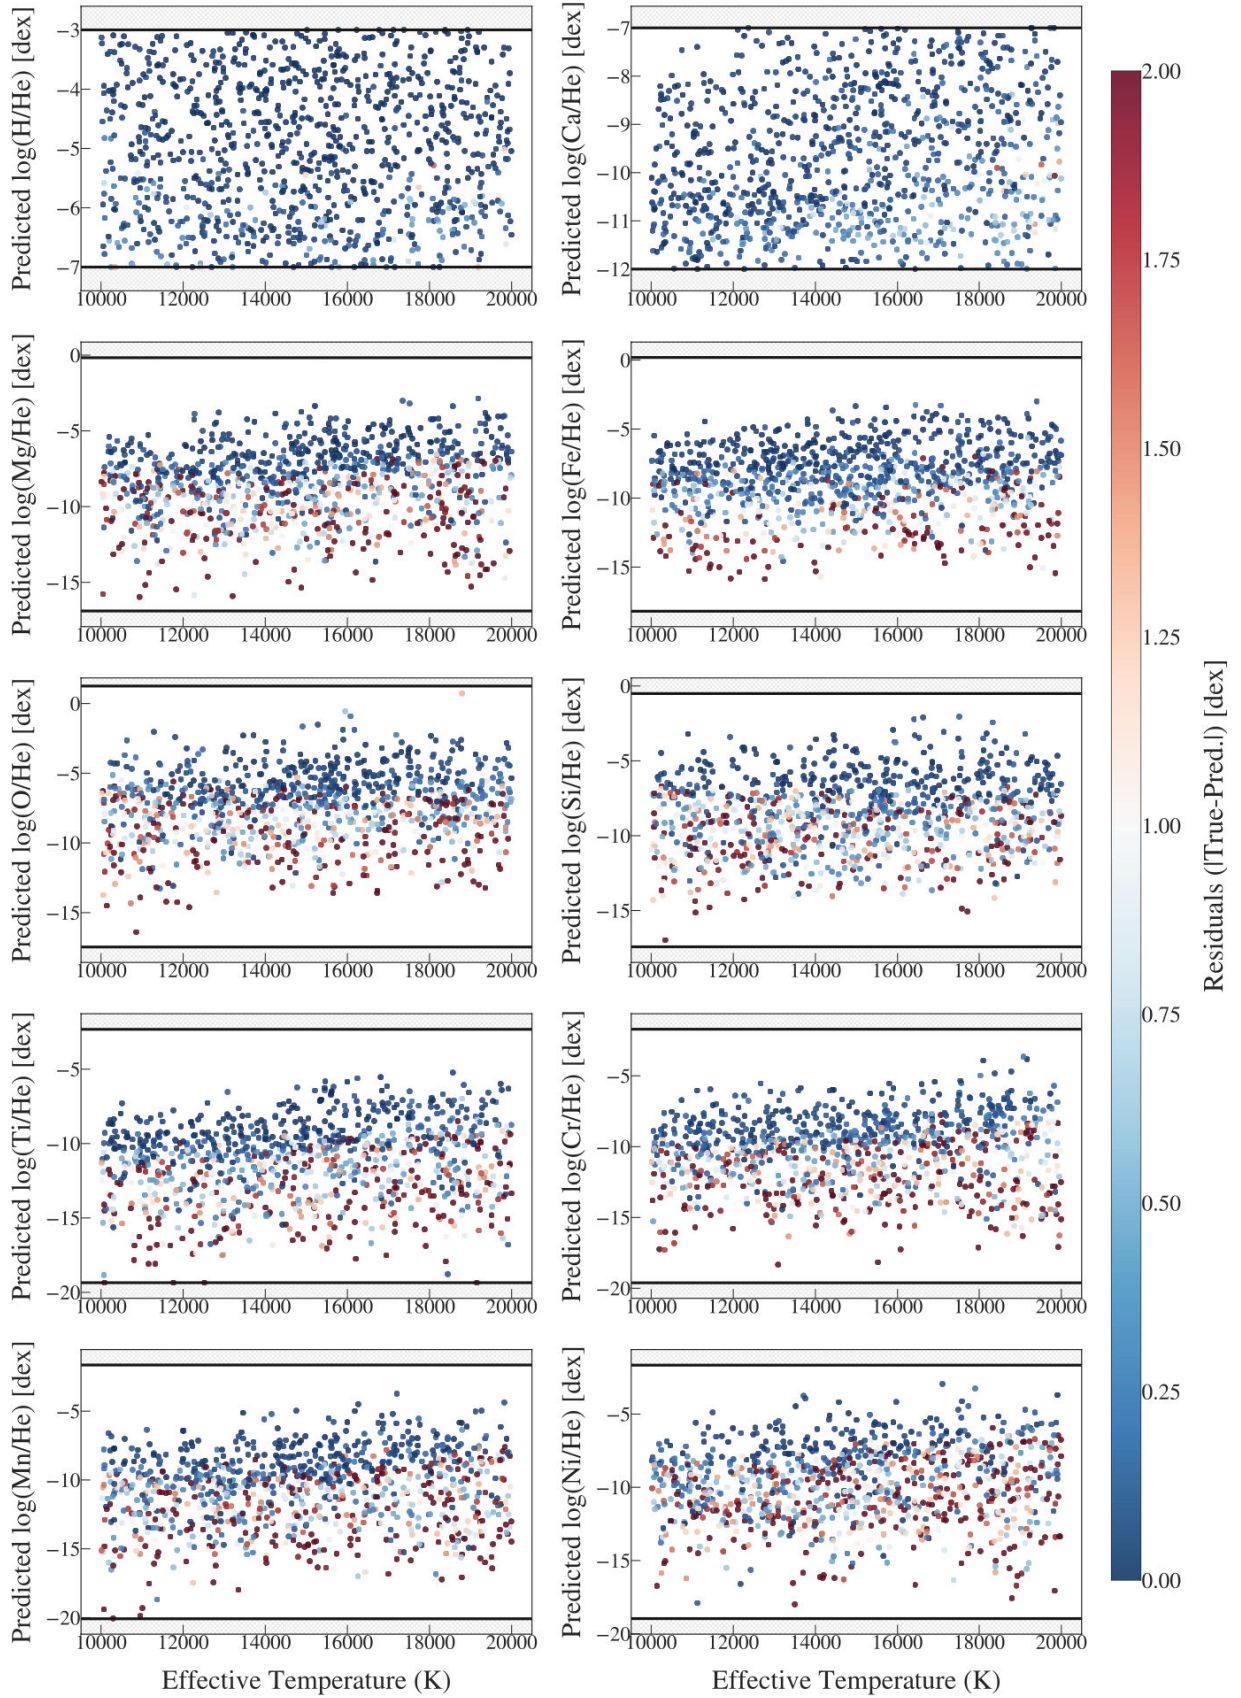
where the MAD provides a robust statistical metric against outliers. The green and orange vertical lines in Figure 10 show the lowest possible abundance values that *cecilia* can retrieve with an average accuracy lower than 0.1 dex and 0.2 dex, respectively. These bounds are summarised in Table 3 and can be used to assess the usefulness and robustness of *cecilia* for the study of a real polluted white dwarf.

In Figure 11, we also illustrate *cecilia*’s estimated abundances as a function of synthetic effective temperature, with the colorbar representing the residuals of its predictions. From this figure, we find that *cecilia*’s performance is overall quite stable over the range between 10,000 K and 20,000 K. However, its residuals are slightly higher for high-temperature white dwarfs, which is consistent with the fact that hotter sources are more opaque, thus making it more difficult to detect weak signs of metal pollution (e.g. Zuckerman et al. 2010). It is also interesting to highlight that *cecilia*’s predictive power is particularly good for  $\log_{10} (\text{H/He})$  and  $\log_{10} (\text{Ca/He})$ , as demonstrated by the low error in their predicted abundances. The level of accuracy for these two stellar labels is possibly due to the depth of their absorption lines compared to those of other metals (more than 1% deep relative to the flux continuum, even for the lowest possible abundances of H and Ca at the maximum and minimum allowed values of  $T_{\text{eff}}$  and  $\log g$ ).

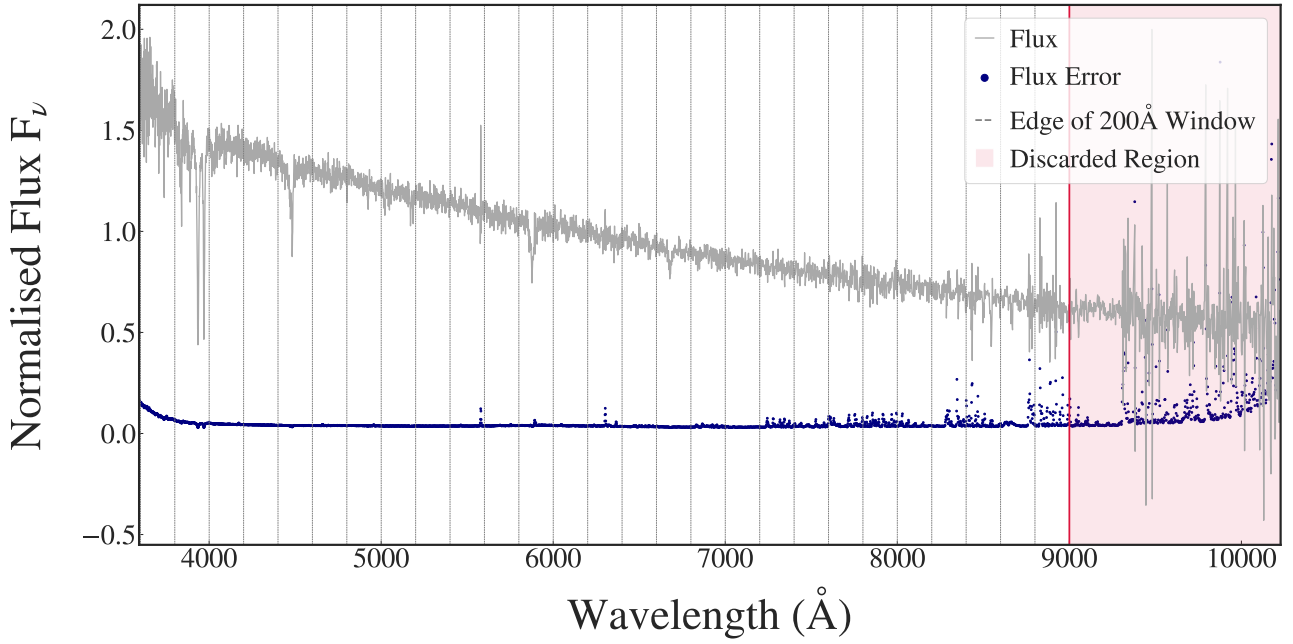
<sup>12</sup> We chose to keep  $T_{\text{eff}}$  and  $\log g$  fixed because these parameters are usually constrained from photometric observations.



**Figure 10.** Residuals of *cecilia* as a function of true elemental abundance for 1,000 noiseless atmosphere models from our test set (see Section 5.1). To conduct our analysis, we fixed  $T_{\text{eff}}$  and  $\log g$  to their true values, and only fitted the 11 abundances considered in this work (i.e. from  $\log_{10}(\text{H}/\text{He})$  to  $\log_{10}(\text{Ni}/\text{He})$  in Table 1). The vertical lines in green and orange denote, respectively, the boundaries where *cecilia* achieves a predictive accuracy lower than 0.1 dex and 0.2 dex as calculated by Eq. 26 (see Table 3). The black lines represent the minimum and maximum ML bounds of *cecilia* (see Table 1), with the grey hatched regions symbolising abundances outside of *cecilia*'s allowed parameter space. This figure does not show any results for  $T_{\text{eff}}$  and  $\log g$ , which were fixed parameters during our fits. It also does not include our residuals for Be due to the poor performance of our code with this metal.



**Figure 11.** cecilia's estimated elemental abundances as a function of synthetic effective temperature. The colorbar shows the residuals of its predictions, which are slightly higher for hotter white dwarfs. Similarly to Figure 10, the black lines denote the minimum and maximum allowed ML bounds of Cecilia (see Table 1), while the grey hatched areas represent photospheric compositions outside of cecilia's range. From this figure, we find that cecilia's predictions are particularly good for  $\log_{10}(\text{H/He})$  and  $\log_{10}(\text{Ca/He})$ .



**Figure 12.** Low-resolution SDSS spectrum of the He-rich, heavily polluted white dwarf WD 1232+563. The normalised  $F_\nu$  flux and its corresponding observational error are presented in grey and blue, respectively. The vertical lines denote the spectral windows of 200 Å used to train *cecilia* (see Section 4.2). The red shaded area covers the wavelength region excluded in our optimisation routine (see Section 5.2).

**Table 4.** Main astrophysical properties of WD 1232+563. References (Ref.): [1] *Gaia* Collaboration et al. (2016, 2023), [2] York et al. (2000); Almeida et al. (2023), [3] Martin et al. (2005), [4] Kaiser et al. (2010), [5] Kleinman et al. (2013), [6] Coutu et al. (2019).

Property	Value	Ref.
<i>Other Target Names</i>		
<i>Gaia</i> DR3	1571584539980588544	[1]
SDSS	J123432.65+560643.1	[2]
<i>GALEX</i>	J123432.6+560642	[3]
<i>Astrometric Properties</i>		
R.A. [J2020; h:m:s]	12:34:32.676	[1]
Dec [J2020; d:m:s]	+56:06:43.034	[1]
Parallax [mas]	5.817±0.095	[1]
Distance [pc]	171.92±3.00	[1]
$\mu_{R.A.}$ (mas y <sup>-1</sup> )	-72.884±0.089	[1]
$\mu_{Dec.}$ (mas y <sup>-1</sup> )	-0.223±0.094	[1]
<i>Photometric Properties</i>		
<i>Gaia</i> $G_{mag}$	18.05±0.01	[1]
SDSS $g_{mag}$	18.27±0.01	[2]
<i>Pan-STARRS</i> $g_{mag}$	17.95±0.01	[4]
<i>Physical and Chemical Properties</i>		
$T_{eff}$ [K]	11,787±423	[6]
log $g$ [cgs]	8.30±0.06	[6]
Mass [ $M_\odot$ ]	0.77	[6]
Cooling Age [Gyr]	0.64	[6]
Atm. Composition	He	[5, 6]
Spectral Type	DBZA	[5, 6]

## 5.2 Performance with Real Observations

In this section, we test and validate *cecilia* against real spectroscopic observations of the He-rich, heavily polluted white dwarf WD 1232+563. This star was discovered by Eisenstein et al. (2006), who compiled a catalog of 19,712 spectroscopically-confirmed white dwarfs from the Sloan Digital Sky Survey (SDSS, York et al. 2000). More recently, it was studied by Coutu et al. (2019) and Xu et al. (2019) at low- and high-resolution: the former conducted a spectroscopic analysis of 1,023 He-atmosphere polluted systems observed by SDSS, while the latter used data from the Keck HIRESb and ESI spectrographs to perform a detailed characterisation of WD 1232+563. From these literature studies, the physical and chemical properties of WD 1232+563 appear to fall well within the allowed bounds of *cecilia* (see Table 4). Therefore, this white dwarf constitutes an excellent candidate to understand the behavior of our pipeline with real astrophysical observations.

To perform our sensitivity analysis, we choose to use *cecilia* to fit the low-resolution, flux-calibrated spectrum of WD 1232+563 from the SDSS DR18 database (plate = 8232, fiber = 84; see Figure 12).<sup>13</sup> Given that SDSS observations have a variable resolving power of  $R=1,500$  at 3,800 Å and  $R=2,500$  Å at 9,000 Å, we adopt a mean value of  $R_{obs} \approx 2,000$  in our fitting procedure. Moreover, to evaluate *cecilia*'s behavior with unprocessed real data, we do not remove any problematic data points from the spectrum, such as bad flux points or instrumental artefacts.

As explained in Section 4.3, the first step in our fitting routine is to load the SDSS spectrum of WD 1232+563 into *cecilia*, together with reasonable estimates for the white dwarf's effective temperature, surface gravity, elemental abundances (H, Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni), and RV shift. To generate our list of stellar parameters, we adopt the  $T_{eff}$  and log  $g$  values obtained by Coutu

<sup>13</sup> <https://skyserver.sdss.org/dr18/>.

et al. (2019) from existing photometric observations; the calcium photospheric abundance derived by the same authors from the Ca II H&K absorption lines at air wavelengths 3,934 Å and 3,369 Å (Kramida et al. 2022); and the elemental abundances of Xu et al. (2019) for the remaining heavy elements. We also assume a zero RV shift as our initial guess for the RV term in our spectral model. In our analysis, we choose to treat all 14 stellar properties as free model parameters.

To begin with, we execute our `mpfit` least-squares  $\chi^2$  minimisation routine, restricting the label parameter space to `cecilia`'s ML bounds, requiring a maximum of 50 iterations, and using gradient step sizes<sup>14</sup> of 0.01 K for temperature, 0.01 cgs for surface gravity, 0.01 dex for all the elemental abundances, and 0.0001 km/s for the RV shift. With these hyperparameters, `mpfit` generates an optimal spectroscopic solution with its corresponding 14 stellar labels in less than 1.80 minutes. Next, we run a final MCMC with  $n_{\text{walkers}}=50$  walkers,  $n_{\text{draws}}=5,000$  draws, and  $n_{\text{burn}} = 0.2 \cdot n_{\text{draws}}=1,000$  draws. We initialize the walkers in Gaussian balls with standard deviations of 0.01 K for  $T_{\text{eff}}$ , 0.01 cgs for  $\log g$ , 0.01 dex for the 11 elemental abundances, and 0.0001 km/s for the RV shift. We then define our likelihood function with Eq. 24, incorporating two types of prior distributions: first, we adopt weak photometric priors of 500 K and 0.1 cgs on  $T_{\text{eff}}$  and  $\log g$  based on the uncertainties of these parameters in Xu et al. (2019); and second, we assume chondritic abundance priors for Be, Cr, Mn, and Ni, which are elements that we cannot visually and confidently detect in the SDSS spectrum.<sup>15</sup> For these chondritic labels, we impose a prior width of 0.5 for Be, Cr, and Ni, with a stricter width of 0.25 for Mn to ensure chain convergence.

With the exception of Be, Cr, Mn, and Ni, for which we use chondritic abundance ratios as initial guesses, we take the results of `mpfit` to initialise our MCMC sampler. This approach, combined with the `edmc` configuration mentioned above, generates a final spectral model in about 10 hours with a converged Gelman-Rubin statistic lower than 1.008 for all our model parameters. In Table 5, we also summarise the MCMC best-fit labels of WD 1232+563, excluding the four undetected heavy elements with chondritic priors (i.e. Be, Cr, Mn, and Ni). We also report the systematic ( $\sigma_{\text{sys,ML}}$ ) and statistical errors ( $\sigma_{\text{stat,MCMC}}$ ) of each model parameter, obtained respectively from our sensitivity analysis of synthetic spectra (see Section 5.1) and our MCMC fit. The former measure the errors of `cecilia`'s predictions arising from the ML-based spectral interpolation process, while the latter represent the errors of the SDSS data and are hence observational in nature. Assuming Gaussian and uncorrelated errors, we can approximate the total uncertainty of a model parameter ( $\sigma_{\text{tot}}$ ) by adding these two sources of errors in quadrature (together with any other systematic errors, or  $\sigma_{\text{other}}$ , coming from the atmosphere models, which we do not consider in this work),

$$\sigma_{\text{tot}} = \sqrt{\sigma_{\text{stat,MCMC}}^2 + \sigma_{\text{sys,ML}}^2 (+\sigma_{\text{other}}^2)}. \quad (27)$$

Using the above expression, we calculate the total uncertainty of each stellar label and provide it in Table 5. As a conservative measure, we assume that the MCMC has underestimated the statistical error of those abundance parameters with  $\sigma_{\text{tot}} < 0.15$  dex, where this

threshold represents an approximate SDSS noise floor.<sup>16</sup> For these abundances (flagged with a star symbol in Table 5), we replace the total uncertainty calculated in this work by  $\sigma_{\text{tot}}=0.15$  dex.

In Figure 13, we also show our `mpfit` and MCMC solutions as well as their corresponding “Observed-Calculated” (O-C) flux residuals. To complement this figure, Figure 14 presents the RV-shifted MCMC model for different regions of the SDSS spectrum with clear metallic absorption lines, together with two additional models obtained with the same  $T_{\text{eff}}$  and  $\log g$ , but with the abundances of the detected elements changed by  $3\sigma_{\text{stat,MCMC}}$ . Finally, Figure 15 presents a corner plot of the posterior probability distributions of the model parameters as a measure of MCMC chain convergence.

Based on Table 5, and taking the properties of WD 1232+563 from Xu et al. (2019) as our assumed ground truth, our results demonstrate that `cecilia` can accurately retrieve the majority of parameters in our model. This includes  $T_{\text{eff}}$  and  $\log g$  —for which we use weak photometric priors— as well as the elemental abundances of 6 heavy elements with visible absorption lines in the SDSS spectrum (Ca, Mg, Fe, O, Si, and Ti; see Figure 14). Among these metals, a few of them (e.g. Si, Ti, O) are weak detections, as evidenced by their small observational signatures and their high statistical uncertainties relative to those of other detected elements. For these elements, we might not have been able to claim a confident detection without knowledge of the higher-resolution and signal-to-noise spectra from Xu et al. (2019). However, given that our current knowledge of this star confirms their existence, it is promising to see that our pipeline can estimate abundance values consistent with those in the literature. In the future, `cecilia`'s detection of weak lines in low-resolution data can be a sign that a white dwarf deserves spectroscopic follow-up at high resolution and high signal-to-noise.

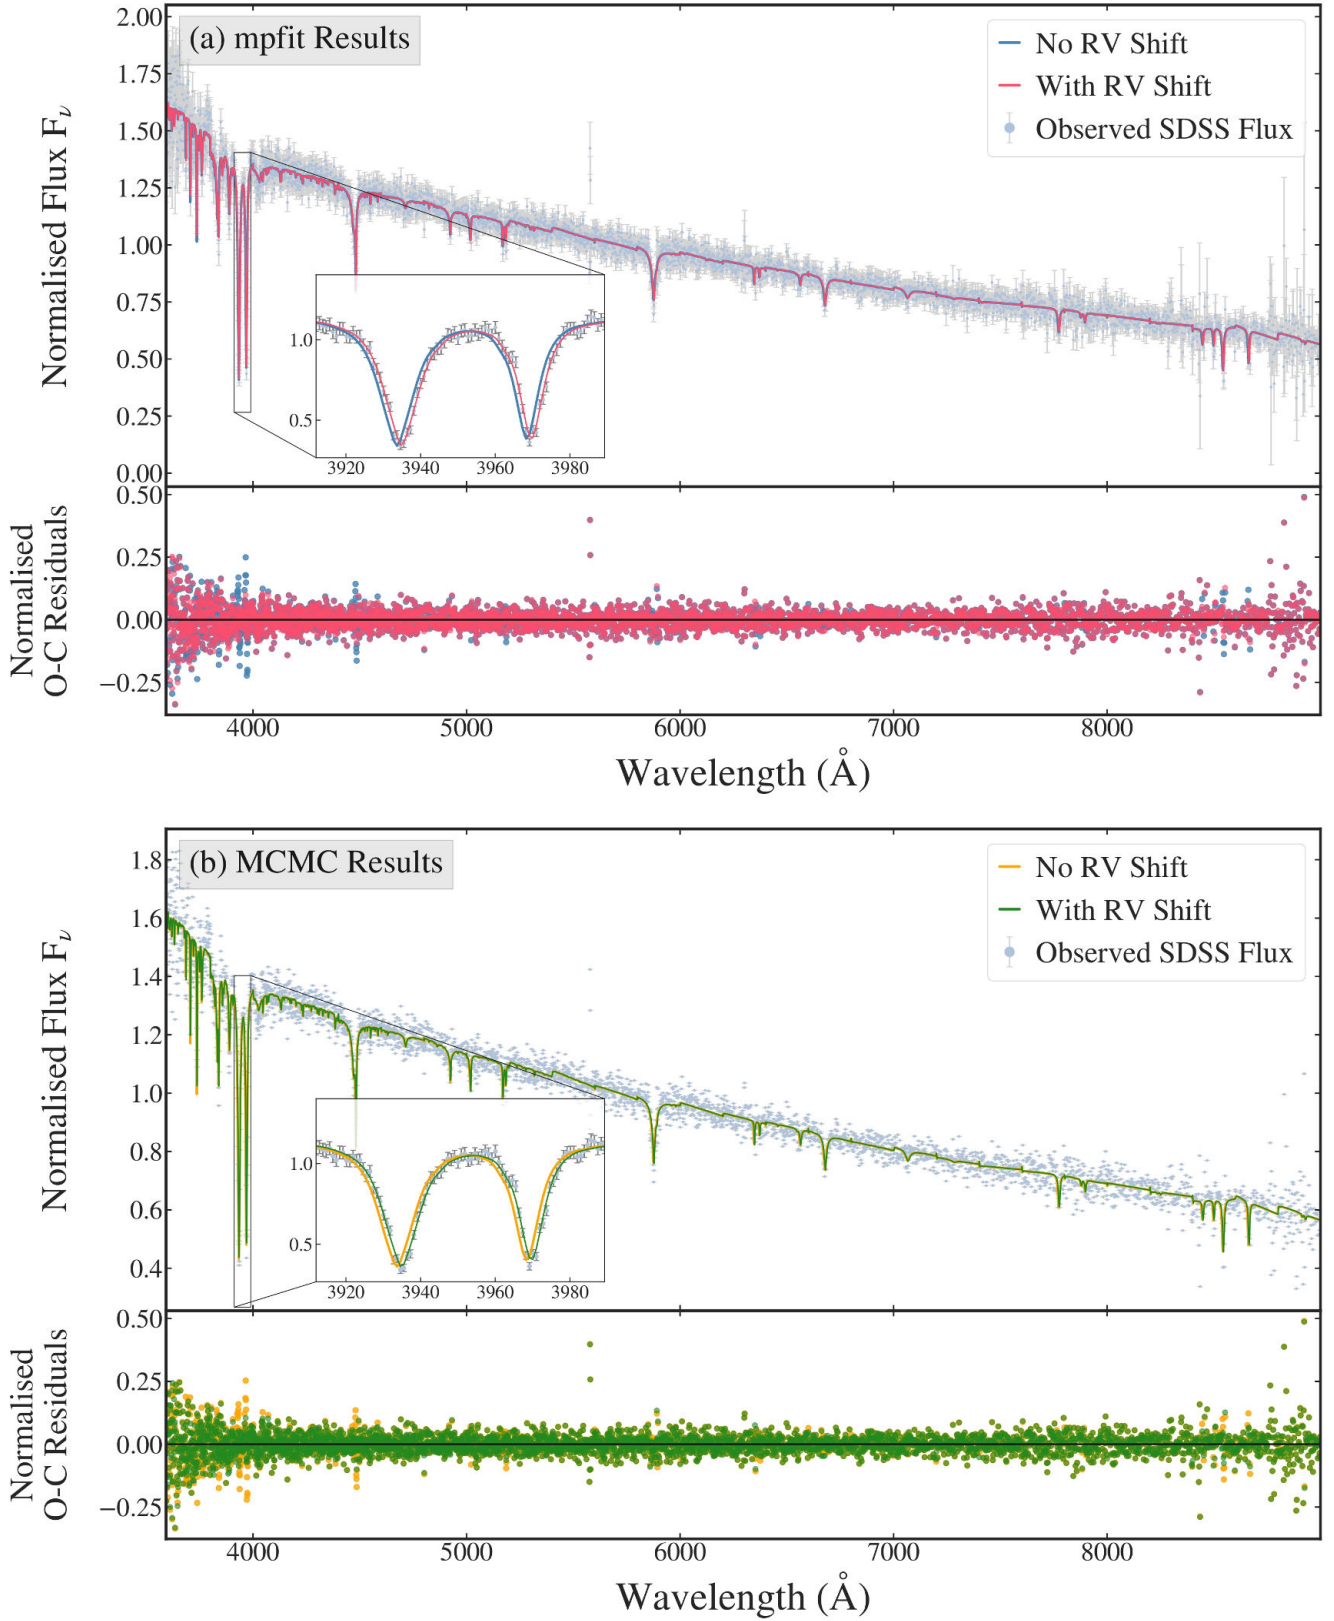
With respect to those elements with no clear spectral signatures (Be, Cr, Mn, and Ni), `cecilia` can marginalise over the uncertainty in the true values of their abundances thanks to the use of chondritic priors in the likelihood function. These chondritic priors limit the values of the parameters explored in our fit, which in turn facilitates MCMC convergence. We note that the assumption of chondritic priors in our MCMC is effectively equivalent to the classical convention of including undetectable metals in atmosphere models with their abundances fixed to a certain value (e.g. chondritic levels).

Another interesting result from Table 5 is `cecilia`'s retrieved calcium abundance ( $\log_{10}(\text{Ca/He}) \approx -7.44$ ), which is more consistent with that of Coutu et al. (2019) ( $\log_{10}(\text{Ca/He})_{\text{Coutu}} \approx -7.41$ ) than with the value reported by Xu et al. (2019) ( $\log_{10}(\text{Ca/He})_{\text{Xu}} \approx -7.69$ ). This discrepancy is understandable given the different optimisation strategies used in each paper. More specifically, Coutu et al. (2019) fitted the same SDSS spectrum of WD 1232+563 considered in this work, focusing exclusively on the Ca II H&K doublet and using a maximum likelihood approach similar to that of our code. In contrast, Xu et al. (2019) implemented a line-by-line  $\chi^2$  minimisation technique of the Ca II H&K doublet and the near-infrared (IR) Ca II triplet in their high-resolution Keck spectrum. For this task, they individually fitted all the Ca lines and weighted them equally to obtain a mean calcium abundance, hence giving more importance to the IR triplet than `cecilia` does. Indeed, our code fits all the Ca lines simultaneously —in combination with the full spectrum— and subsequently weighs them based on their observed strengths, uncertainties, and signal-to-noise. As Figure 14 shows, the Ca II H&K lines in the SDSS data are significantly stronger than the near-IR Ca

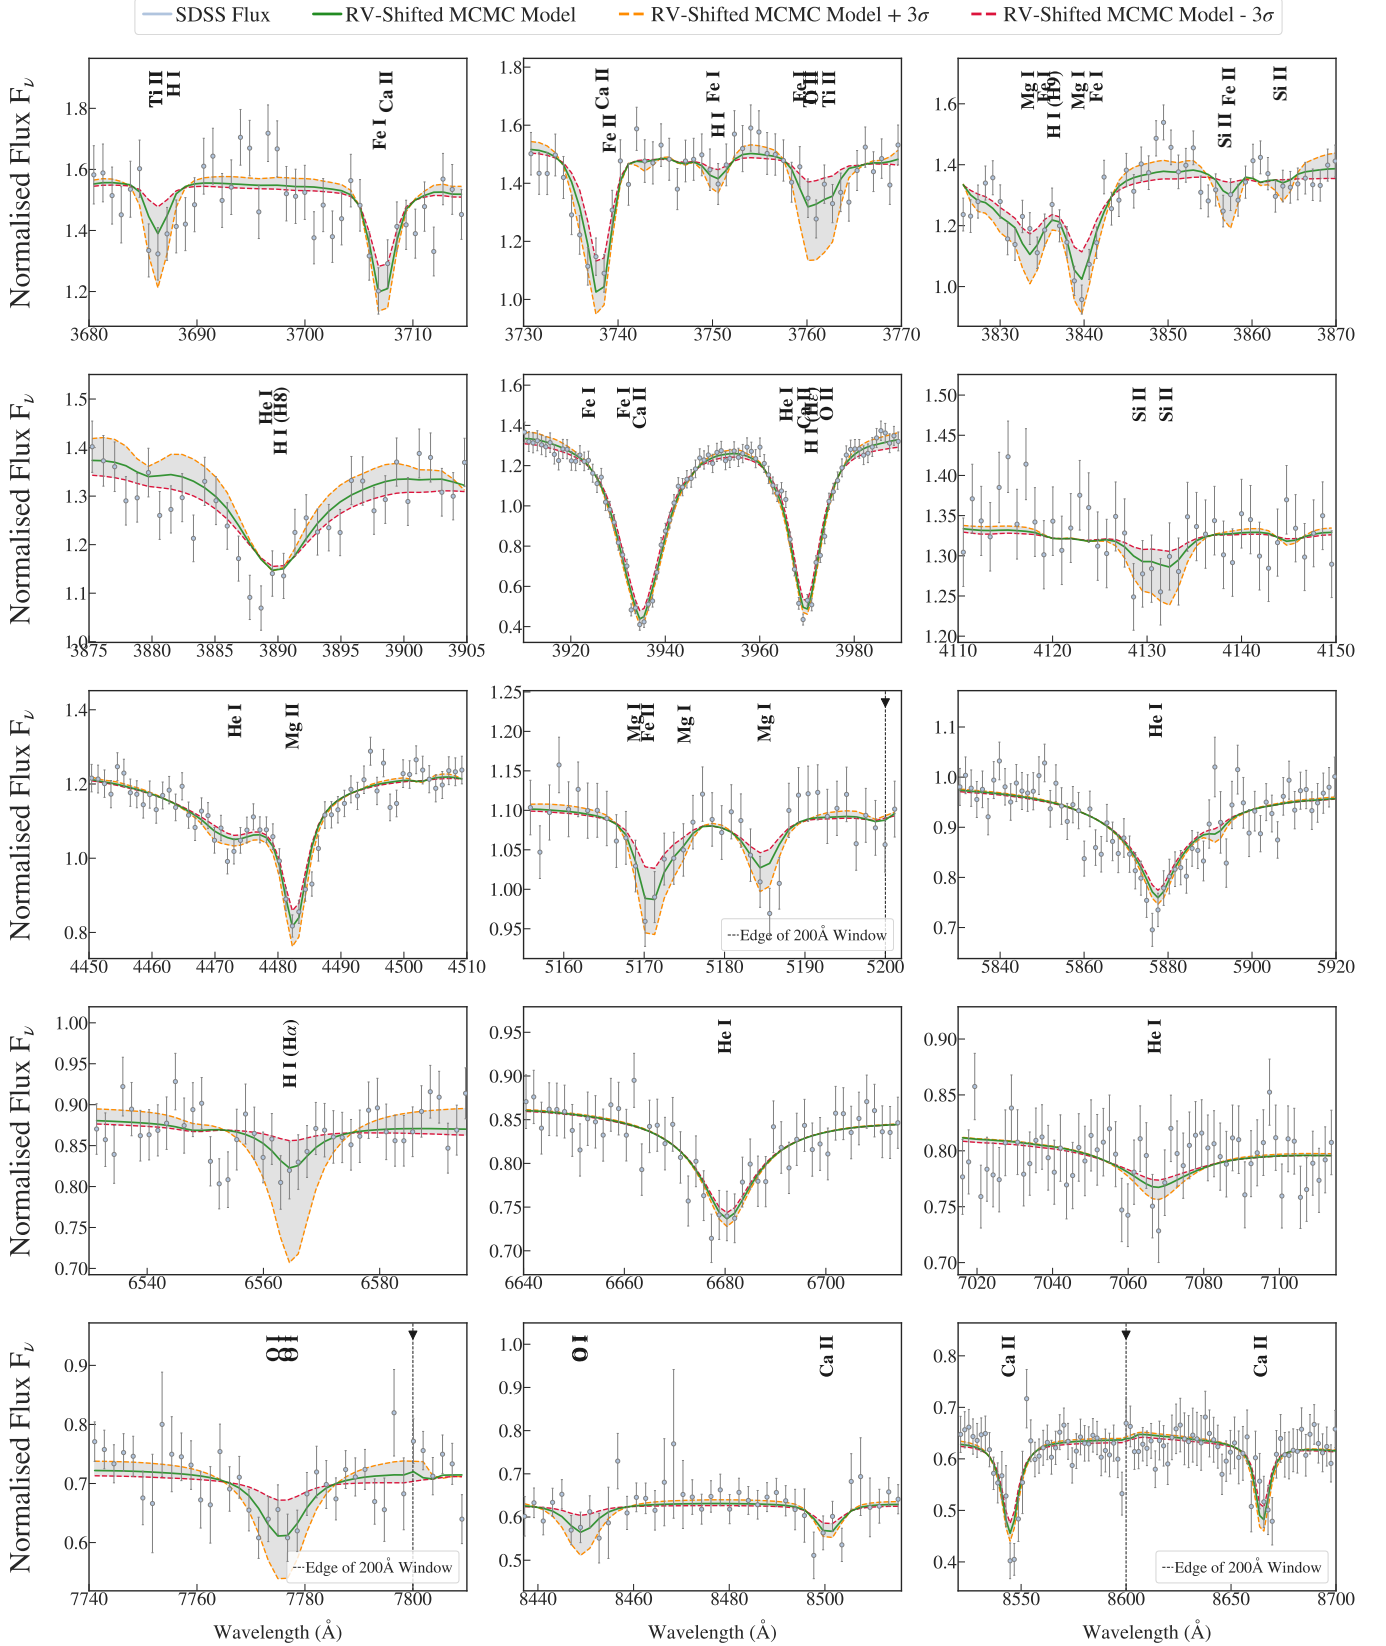
<sup>14</sup> In the `mpfit` code, the “step size” corresponds to the step size used to calculate gradients during the fitting process. It does not refer to the size of the parameter update at each iteration of the algorithm for a given model parameter.

<sup>15</sup> We note that we also considered using a chondritic prior on Ti, but we ended up discarding this option after realising that `cecilia` could constrain this metal relatively well.

<sup>16</sup> Our assumed SDSS noise floor of 0.15 dex is a reasonable value based on the authors' experiences with white dwarf spectral modelling.



**Figure 13.** Normalised best-fit models for the SDSS spectrum of WD 1232+563 (in light blue). *Panel (a):* Best-fit `mpfit` models using an RV shift of 0 km/s and 94.42 km/s (in red and blue, respectively). The inset plot shows the spectral window between 3,800 Å and 4,000 Å, where there are two important Ca II lines at about 3,934 Å and 3,969 Å. The bottom plot shows the normalised flux residuals (Observed-Calculated; O-C) associated to each prediction. *Panel (b):* Best-fit `edmc` MCMC models using an RV shift of 0 km/s and 94.74 km/s (in green and orange, respectively). Similarly to Panel (a), the inset plot shows the spectral window between 3,800 Å and 4,000 Å, while the bottom panel shows the normalised O-C residuals for each spectroscopic solution. We note that the MCMC value of the RV shift is different from that reported in Table 5 because it also accounts for the barycentric velocity and gravitational redshift correction.



**Figure 14.** A selection of spectral windows illustrating our best-fit, RV-shifted MCMC model (in green) for the raw SDSS spectrum of WD 1232+563 (in light blue with grey error bars). In each panel, we also show a grey shaded area bounded by an orange and red dashed line, which correspond, respectively, to a  $\pm 3\sigma_{\text{stat,MCMC}}$  variation on the best-fit model (using the same  $T_{\text{eff}}$  and  $\log g$ , and only changing the abundances of the detected elements by  $3\sigma_{\text{stat,MCMC}}$ ). In black letters, we indicate some of the strongest spectral signatures of hydrogen, helium, and multiple heavy elements. Finally, we represent the start/end of a 200  $\text{\AA}$  spectral window with a dashed vertical line and a black inverted triangle symbol.

**Table 5.** MCMC best-fit stellar parameters for the SDSS spectrum of WD 1232+563, together with their uncertainties and assumed ground truths. The systematic ( $\sigma_{\text{sys,ML}}$ ) and statistical uncertainties ( $\sigma_{\text{stat,MCMC}}$ ) measure the errors associated to *cecilia*’s spectral interpolation (see Section 5.1) and to the MCMC fit, respectively. For abundance studies of exoplanetary material, we can approximate the total uncertainty of a model parameter ( $\sigma_{\text{tot}}$ ) by computing their quadrature sum (see Eq. 27). In this work, if the total error of an abundance is lower than an SDSS noise floor of 0.15 dex, we assume that the MCMC statistical errors are underestimated and impose  $\sigma_{\text{tot}} \approx 0.15$  dex (see star  $\star$  symbol). The dagger symbol ( $\dagger$ ) flags the model parameters dominated by photometric priors ( $T_{\text{eff}}$ ,  $\log g$ ). In this table, we do not present the best-fit abundances of the four heavy elements with chondritic priors (Be, Cr, Mn, and Ni); these metals were included in *cecilia*’s MCMC, but are not detected in the SDSS spectrum. Finally, we do not report the systematic errors for  $T_{\text{eff}}$  and  $\log g$  because these two parameters were fixed during our sensitivity study of synthetic spectra (see Section 5.1). References (Ref.): [1] [Coutu et al. \(2019\)](#), [2] [Xu et al. \(2019\)](#).

Label <sup>a</sup>	cecilia’s MCMC Best-Fit Model		Systematic <sup>b</sup> and Total Uncertainty		Assumed Ground Truth	
	Value	$\sigma_{\text{stat,MCMC}}$	$\sigma_{\text{sys,ML}}$	$\sigma_{\text{tot}}$	Value	Ref.
$T_{\text{eff}}^{\dagger}$ [K]	11777.80	+86.33 −84.81	-	±85.57	11787±423	[1]
$\log g^{\dagger}$ [cgs]	8.24	+0.06 −0.06	-	±0.06	8.30±0.06	[1]
RV <sup>c</sup> [km/s]	36.09	+5.41 −5.41	±10	±11.37	19.0±2.0	[2]
$\log_{10}$ (H/He)	-5.92	+0.15 −0.19	±0.03	±0.17	-5.52±0.10 -5.90±0.15	[1] [2]
$\log_{10}$ (Ca/He)	-7.44	+0.05 −0.05	±0.07	±0.15 $\star$	-7.41±0.06 -7.69±0.05	[1] [2]
$\log_{10}$ (Mg/He)	-6.11	+0.04 −0.04	±0.05	±0.15 $\star$	-6.09±0.05	[2]
$\log_{10}$ (Fe/He)	-6.51	+0.07 −0.08	±0.06	±0.15 $\star$	-6.45±0.11	[2]
$\log_{10}$ (O/He)	-5.19	+0.12 −0.13	±0.09	±0.15 $\star$	-5.14±0.15	[2]
$\log_{10}$ (Si/He)	-6.45	+0.10 −0.12	±0.07	±0.15 $\star$	-6.36±0.13	[2]
$\log_{10}$ (Ti/He)	-8.92	+0.11 −0.13	±0.08	±0.15 $\star$	-8.96±0.11	[2]

[a]: All abundances are given as logarithmic number abundance ratios relative to Helium, i.e.  $\log_{10}(n(Z)/n(\text{He}))$ .

[b]: We note that there is at least another source of systematic errors coming from the atmosphere models themselves; these errors should be of the order of 0.10 dex and are not considered in this work.

[c]: We assume a SDSS systematic RV error of 10 km/s based on the estimates provided by e.g. [Abazajian et al. \(2009\)](#).

II triplet at around 8,498 Å, 8,542 Å and 8,662 Å. Therefore, similarly to the study of [Coutu et al. \(2019\)](#), they dominate *cecilia*’s fitting procedure, leading to a best-fit  $\log_{10}$  (Ca/He) abundance somewhat higher than the result of [Xu et al. \(2019\)](#). In addition to the different fitting methodologies employed in each paper, it is also possible that the uncertainties on the SDSS flux, which are higher in the red region of the spectrum (likely due to telluric and/or instrumental effects), may have brought our Ca abundance closer to that of ([Coutu et al. 2019](#)) than to the result of ([Xu et al. 2019](#)).

Besides providing the main physico-chemical properties of WD 1232+563, *cecilia* generates the posterior probability distributions of its model parameters and shows their interdependence in the form of a corner plot (see Figure 15). With this Bayesian approach, *cecilia* differs from conventional white dwarf analysis techniques, which can only use a limited number of points in their interpolation grid to explore the relationship between two model parameters (e.g. [Klein et al. 2011](#); [Xu et al. 2019](#)). In general terms, corner plots reveal possible correlations between two labels if their posterior distributions have a slanted and oval-like contour shape. For WD 1232+563, this seems to be the case for multiple stellar properties, such as  $T_{\text{eff}}$  and  $\log g$ ,  $T_{\text{eff}}$  and  $\log_{10}$  (Ca/He),  $\log g$  and  $\log_{10}$  (Mg/He), or  $\log_{10}$  (Ca/He) and  $\log_{10}$  (Mg/He). To qualitatively investigate the presence of these correlations, we computed the Pearson Coefficient (PC) for all our model parameters,

$$\text{PC} = \frac{\sum(\pi_x - \mu_{\pi_x})(\pi_y - \mu_{\pi_y})}{\sqrt{\sum(\pi_x - \mu_{\pi_x})^2 \sum(\pi_y - \mu_{\pi_y})^2}} \in [-1, 1], \quad (28)$$

where  $\pi_x$  and  $\pi_y$  are the posterior distributions of labels  $x$  and  $y$ , and  $\mu_{\pi_x}$  and  $\mu_{\pi_y}$  are the arithmetic means of their posteriors. A PC greater or lower than 0 denotes, respectively, a positive or negative

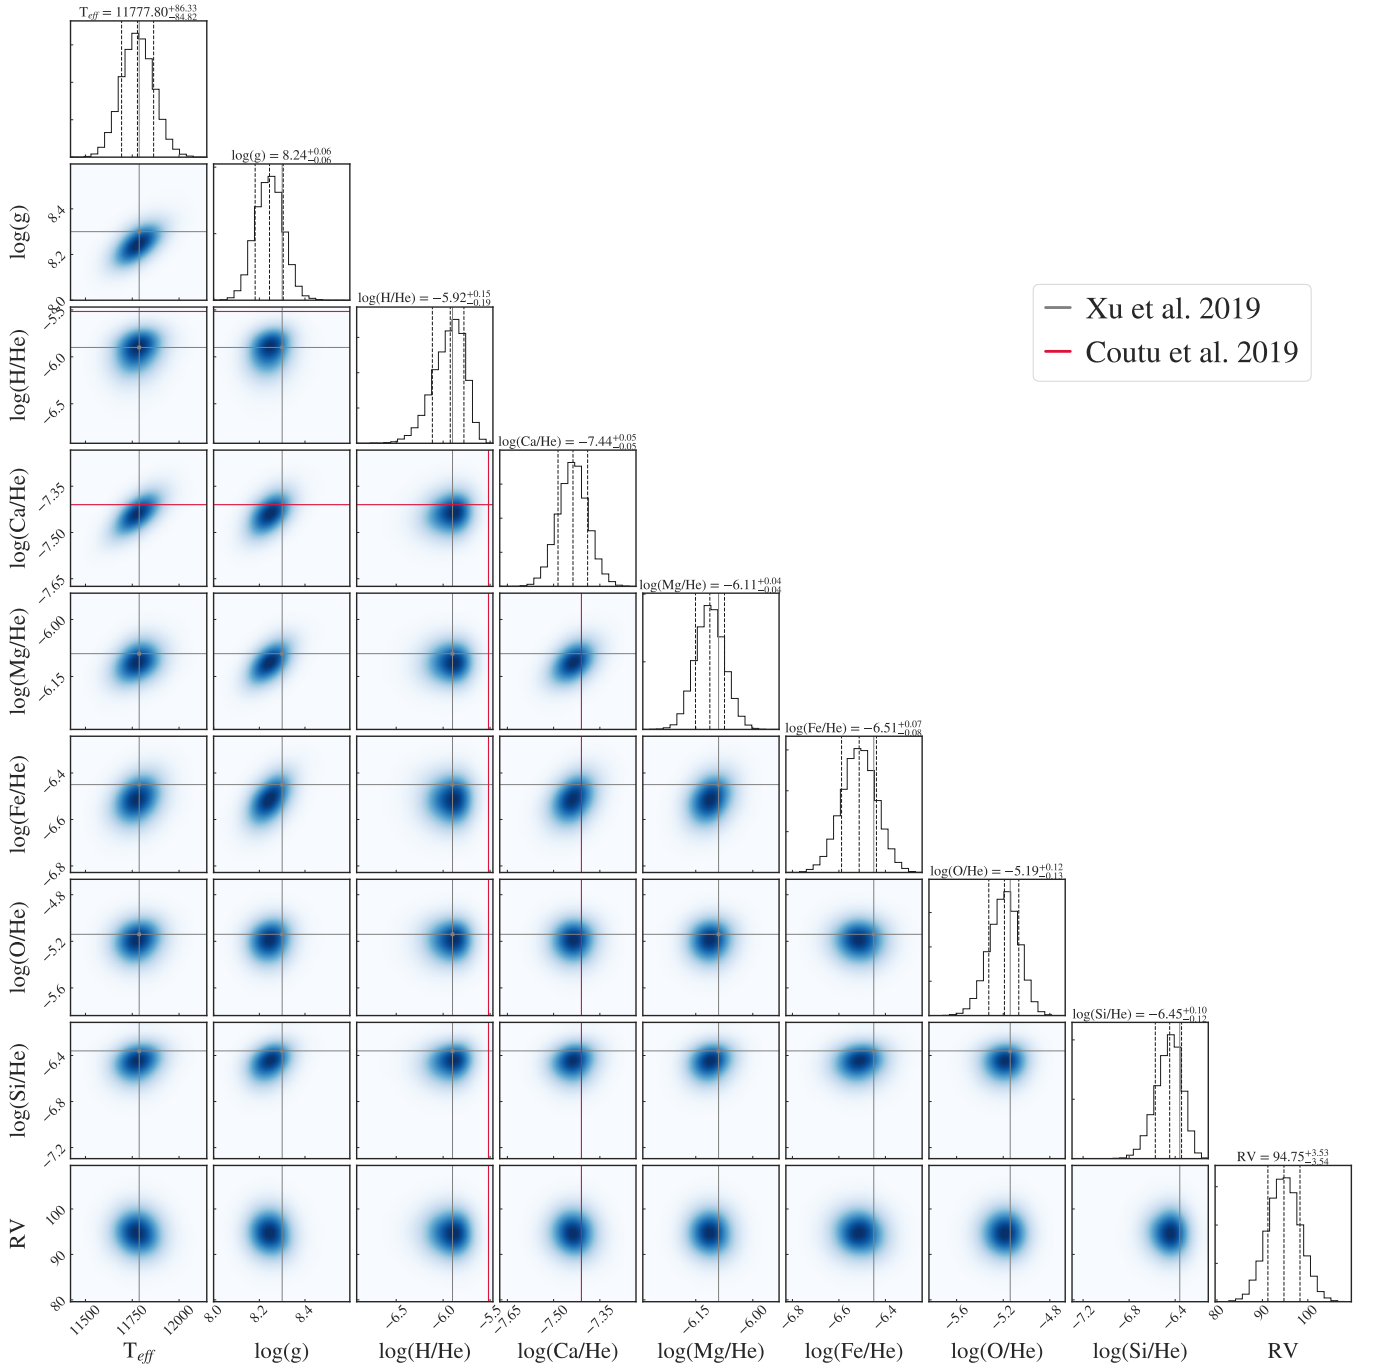
correlation, while a  $\text{PC} = \pm 1$  indicates that the data can be perfectly modelled by a line.

In Figure 16, we present the Pearson coefficients for *cecilia*’s best-fit MCMC model parameters. Assuming a strong positive correlation if  $\text{PC} \geq 0.5$ , our results suggest that certain stellar labels are linearly dependent on one another (e.g.  $T_{\text{eff}}$  and  $\log g$ , or  $\log_{10}$  (Ca/He) and  $\log_{10}$  (Mg/He)), with an increase in one label resulting in an increase of its corresponding label pair. Based on the corner plot and the Pearson coefficients alone, it is hard to determine whether the observed correlations can be generalised to the entire population of polluted white dwarfs. Although such an investigation is outside the scope of this paper, it is possible that certain stellar properties are linearly dependent on one another due to observational or physical reasons. Observationally, some of the metal correlations may be caused by the presence of blended absorption lines. From a physical perspective, changes in  $T_{\text{eff}}$  and  $\log g$  (a proxy for pressure) modify the temperature-pressure conditions at the stellar photosphere, which in turn control the strength and the width of the absorption lines. In addition, different values of  $T_{\text{eff}}$  and pressure result in different excitation, ionisation, or dissociation states of the atoms in the stellar photosphere (e.g. Saha equation, Boltzmann equation), which could also potentially affect the metal abundances of the white dwarf.

## 6 DISCUSSION

### 6.1 Implications for White Dwarf Science

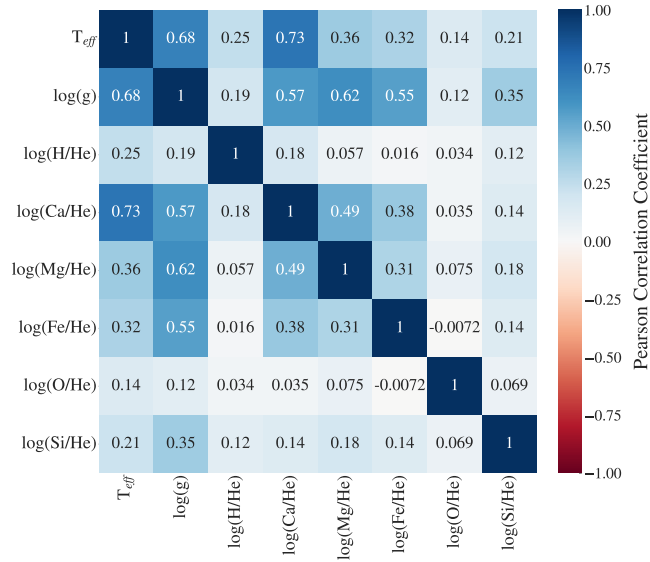
As demonstrated in Section 5, *cecilia* is the first pipeline capable of simultaneously estimating 10 elemental abundances with a retrieval accuracy of  $\lesssim 0.1$  dex. This performance compares favorably to that of conventional techniques, which in turn depend on a variety of



**Figure 15.** MCMC corner plot for the SDSS spectrum of WD 1232+563. The off-diagonal plots show the two-dimensional histograms of the posteriors of the model parameters, marginalised over the rest of parameter values. In contrast, the histograms along the diagonal illustrate their one-dimensional marginalised distributions, together with their corresponding  $1\sigma$  confidence intervals. The grey and red lines show the best-fit values of Xu et al. (2019) and Coutu et al. (2019), respectively, as defined in Table 5. In this figure, we exclude the 4 heavy elements with chondritic priors (i.e. Be, Cr, Mn, and Ni) and show the RV shift without subtracting the barycentric motion and gravitational redshift of the white dwarf.

factors, such as the main properties of the star (e.g.  $T_{\text{eff}}$ ,  $\log g$ , level of metal contamination), the quality of the spectroscopic observations (e.g. resolution, signal-to-noise), or the constitutive physics of their model atmosphere codes. In particular, conventional methods can reach typical uncertainties of about 0.10-0.20 dex for He-rich polluted white dwarfs, both across the UV and the optical range, and irrespective of their underlying atmosphere models (e.g. Doyle et al.

2023; Klein et al. 2021; Izquierdo et al. 2020; Raddi et al. 2015; Wilson et al. 2015; Jura et al. 2012; Zuckerman et al. 2007). Therefore, cecilia’s performance is close to—if not comparable to—that of classical approaches, only introducing small uncertainties due to the accumulation of errors in its ML predictions. With a promising performance and an automated architecture that eliminates the need for manual, time-consuming, and iterative work, cecilia stands as



**Figure 16.** Pearson Correlation Coefficients (PC) for the stellar properties of WD 1232+563. Some label pairs (e.g.  $T_{\text{eff}}$  and  $\log g$ ,  $T_{\text{eff}}$  and Ca, H and Mg) exhibit a strong positive correlation ( $\text{PC} \geq 0.5$ ), as suggested by their slanted, oval-like corner plots in Figure 15. In general, however, most labels have a weak or moderate correlation ( $\text{PC} < 0.5$ ).

a promising tool for the study of polluted white dwarfs and the bulk composition of their accreted material.

With the advent of wide-field, multi-object astronomical surveys like SDSS-V (Kollmeier et al. 2017), the Dark Energy Spectroscopic Instrument (DESI; Cooper et al. 2023; DESI Collaboration et al. 2016), or WEAVE (Dalton et al. 2014), the volume of spectroscopic observations is poised to grow exponentially (Smith & Geach 2023). These surveys are expected to acquire a large number of white dwarf spectra, so they represent a unique opportunity to improve our fundamental knowledge of metal pollution. Unfortunately, as of today, the exploitation of these datasets remains a challenge for conventional techniques due to their time-consuming and human-supervised nature. With *cecilia*, we have developed a practical solution to characterise the spectra of polluted white dwarfs in a fast, automated, and accurate way. For instance, using 1 GPU of the MIT Satori supercomputer and limiting our optimisation procedure to *mpfit*, we estimate that *cecilia* would take about 14 and 17 days to process the 40,000 and 50,000 white dwarf spectra in the cumulative DESI and WEAVE databases, respectively. In reality, *cecilia*’s performance should be much faster than two weeks, as only a small fraction of these observations will exhibit signatures of metal pollution. To quickly identify these polluted spectra, *cecilia* could benefit from other efficient ML-based classifiers, such as the deep neural networks of Vincent et al. (2023) or the Random Forest approach of Garcia-Zamora et al. (2023).

Another valuable point of comparison is the SDSS-V survey, which has monitored more 6,000 unique *Gaia* white dwarfs as of 2021, and will acquire a total of 100,000 spectra of these stars in the coming years (Kollmeier et al. 2017; Chandra et al. 2021). Assuming, unrealistically, that all these spectra will be metal-polluted, they would represent an intractable amount of data for traditional white dwarf characterisation methods. Nonetheless, with *mpfit* only, it would only take about a month to process all these observations. Therefore, in light of these capabilities, our ML-based pipeline is a powerful tool to harness the information content of massive spectroscopic

databases and scale up individual studies of polluted white dwarfs, thus opening the door to a statistically-informed understanding of ancient extrasolar geochemistry. Ultimately, we aspire to use *cecilia* to bring the field of white dwarf science into the era of big data, allowing scientists to overcome the scalability problem associated to “human-in-the-loop” conventional analysis techniques.

In addition to *cecilia*’s promising retrieval accuracy and speed, our pipeline offers a innovative Bayesian framework to estimate the elemental abundances of polluted white dwarfs. More specifically, our inference approach differs from that of classical methods because (i) it integrates prior knowledge about the spectra and the main properties of the white dwarf, (ii) it provides robust parameter uncertainties based on their full posterior probability distributions, and (iii) it illustrates possible degeneracies between different model parameters thanks to the use of corner plots. Through these added capabilities, *cecilia* aims to provide a more detailed and nuanced perspective on the properties of white dwarf pollutants and their parent bodies.

## 6.2 Limitations and Future Work

The sensitivity analyses presented in Section 5 show that *cecilia* can effectively determine the main properties of He-rich polluted white dwarfs, achieving an accuracy comparable to that of conventional techniques without the need for human supervision. This performance has the potential to unlock population-wide studies of metal pollution, but it has several limitations that deserve further investigation. In this Section, we highlight the main caveats of our pipeline, discuss possible mitigation strategies, and propose new opportunities for future work.

### 6.2.1 General ML Design Choices

First, *cecilia*’s ML architecture can be improved along three important dimensions: versatility, speed, and accuracy. In terms of versatility, our work has revolved around He-dominated polluted white dwarfs, which—for a given abundance level—exhibit stronger metallic lines than H-rich stars due to their lower photospheric opacity (Dufour et al. 2012; Klein et al. 2021; Saumon et al. 2022). Nevertheless, *cecilia*’s potential extends beyond He-dominated systems, as its ML architecture can always be retrained with the model atmospheres and astrophysical properties of other stellar objects.

With respect to speed, our full optimisation procedure takes between 4 and 10 hours to provide a best-fit spectroscopic solution, depending on the user’s choice of MCMC hyperparameters (see Section 4.3). This behavior makes *cecilia* an efficient pipeline capable of modelling several stars at the same time, while running uninterruptedly on a high-performance supercomputer cluster. The scalability of *cecilia* represents an improvement over the performance of classical methods, but it can still be optimised further in preparation for large-scale studies of polluted white dwarfs. A possible solution to making *cecilia* more efficient would be to explore whether simpler and/or faster ML architectures can achieve similar performance.

Lastly, *cecilia*’s retrieval accuracy is better than 0.1 dex for 10 heavy elements, which places our pipeline at the level of many conventional techniques. Nevertheless, the errors accumulated in *cecilia*’s networks introduce a small amount of noise into its final synthetic prediction, which in turn challenges the modeling of very shallow lines and introduces limits below which *cecilia* performs poorly. As model atmospheres improve with our understanding of white dwarf physics, larger training sets and/or better ML architectures might help improve *cecilia*’s predictive abilities, obtaining

a better match between its best-fit solutions and the observed metal abundances of a polluted white dwarf. Regardless of whether we improve *cecilia*'s ML predictions, the errors are small enough that most absorption lines detected in the data would be well modelled by our code.

### 6.2.2 Training Improvements

As explained in Section 3, the size and quality of a training dataset can highly affect the performance of a neural network. In this work, we have trained *cecilia* with more than 22,000 synthetic labels and atmosphere models, assuming that white dwarfs can be characterised with a set of 13 randomly varied independent parameters ( $T_{\text{eff}}$ ,  $\log g$ ,  $\log_{10}(\text{H/He})$ , and 10 metal abundances). As we design future iterations of *cecilia*, a larger and more comprehensive training dataset might help our pipeline to achieve a better level of precision. Another important update of *cecilia* would involve changing our training strategy to prevent distortions of the lines near the edges of the 200 Å spectral windows. A possible solution to this problem would be to re-train *cecilia* using a system of overlapping windows.

In relation to the stellar labels used in this work, we have identified three improvement opportunities that can be easily implemented in the future. To begin with, we can extend the abundance ranges listed in Table 1 —upwards and downwards— so that *cecilia* can gain sensitivity to both stronger and weaker levels of metal pollution. This is particularly relevant for Ca and Be. For instance, the inclusion of higher Ca abundances in our training set would allow us to validate our code against several heavily polluted white dwarfs whose astrophysical properties are currently outside of *cecilia*'s allowed ML bounds (e.g. GD 40, Ton 345, GD 362, although the latter would also require extending *cecilia* to lower temperatures, as discussed below). Similarly, expanding the abundance ranges of Be to higher values would make our pipeline more sensitive to stronger Be lines. This would facilitate the analysis of several Be-polluted systems such as those described in Klein et al. (2021), which have a  $\log(\text{Be/Ca})$  ratio ( $\approx -2.50$  dex) considerably higher than that of the Sun ( $\approx -4.96$  dex, Asplund et al. 2009). Second, we can make our pipeline more useful by retraining it with more heavy elements. With this idea in mind, we aim to include at least three new elements: carbon (C), sodium (Na), and aluminium (Al). These species have been detected in the atmospheres of multiple white dwarfs and can provide valuable insights into the volatile (C) and rocky (Na, Al) compositions of white dwarf pollutants (Greenstein 1976; Holberg et al. 1998; see review by Klein et al. 2021). Finally, we can try to improve *cecilia*'s performance by exploring different ways to generate our stellar labels. This is especially important for the 14 metals with fixed chondritic abundances (C, N, Li, Na, Al, P, S, Cl, Ar, K, Sc, V, Co, and Cu), which might have introduced biases in our pipeline and caused *cecilia*'s spectral predictions to favor chondritic values. Although the results presented in Section 5 do not show any evidence of this behavior, it will be important to investigate this potential problem in future versions of *cecilia*.

Regarding our atmosphere models, there are two possible changes that might make *cecilia* more useful and versatile. On the one hand, we can retrain our pipeline using a broader wavelength coverage. At present, *cecilia* can analyze spectroscopic observations in the optical range between 3,000 Å and 9,000 Å, but it is not prepared to model data in the UV ( $\sim 900 \leq \lambda \leq 3,000$  Å). The UV contains important spectral signatures of volatile species —e.g. oxygen, carbon, nitrogen, sulfur, phosphorous— and is thus a crucial region for understanding the volatile budget of a white dwarf pollutant. From

a practical point of view, the integration of the UV into our pipeline will also make *cecilia* particularly relevant for the study of UV observations from space-based facilities such as the Hubble Space Telescope.

Moreover, we can also retrain *cecilia* with synthetic spectra covering a broader range of effective temperatures. For example, it could be useful to include models of slightly cooler ( $T_{\text{eff}} < 10,000$  K) and hotter white dwarfs ( $20,000 \leq T_{\text{eff}} \leq 25,000$  K). At temperatures higher than  $T_{\text{eff}} > 25,000$  K, the origin of metal pollution is unclear due to the outward effects of radiative levitation pressure, which can bring metals to the photosphere from the deep stellar interior (Chayer et al. 1995). For cooler white dwarfs, it is still possible to observe traces of metal pollution in their spectra (e.g. Elms et al. 2022; Blouin 2020; Hollands et al. 2017). While most of these cool systems only show a handful of metallic absorption lines (because of the insufficient thermal energy to populate atomic states responsible for many transitions in the optical), they also allow the detection of elements that cannot be easily observed in warmer objects where they are ionized (e.g. Na, K, Li).

### 6.3 Fitting Analysis

In addition to exploring possible improvements to *cecilia*'s ML architecture and training datasets, we have also identified four opportunities to refine our fitting procedure. First, we can reconfigure our code to predict upper abundance limits when the absorption features of a metal are ambiguous or unobservable. This option is currently unavailable because our pipeline was trained in logarithmic space, which implies that *cecilia* can never predict a zero elemental abundance (i.e. the complete lack of an element). Introducing the possibility of quantifying upper limits would involve a reformulation of our code, but would constitute a valuable addition to *cecilia*.

A second enhancement to our pipeline would be to allow for a simultaneous spectroscopic and photometric fit of a polluted white dwarf spectrum. Given that our code does not currently use photometric observations in its optimisation routine, *cecilia* requires the user to have accurate estimates of the effective temperature and surface gravity of the star. These two parameters are also dependent on the amount of heavy elements in the stellar photosphere, so their prediction becomes an iterative problem.<sup>17</sup> Moreover, even if these two properties are well-defined, they might have been estimated with a set of atmosphere models different to those employed in this work, which could potentially introduce minor errors in *cecilia*'s predictions. Therefore, to ensure that our parameter estimation routine is entirely self-consistent, it would be beneficial to incorporate a simple neural network capable of estimating values for  $T_{\text{eff}}$  and  $\log g$  based on the metal abundances predicted by *cecilia* at each iteration.

Third, we hope to improve *cecilia*'s MCMC fitting procedure by introducing a jitter term in our log-likelihood function. This parameter would allow us to fit unknown sources of noise (e.g. instrumental noise, stellar activity), and could thus be particularly important to model weak absorption lines in real spectroscopic observations. Finally, given that MCMC pipelines can sometimes be quite sensitive to

<sup>17</sup> We note that this issue also affects traditional white dwarf characterisation methods. As of today, several iterative approaches have been developed to mitigate it (e.g. Dufour et al. 2005, 2007; Coutu et al. 2019). These approaches perform an initial fit to the stellar photometry, followed by a spectroscopic fit. Afterwards, this optimisation process is repeated multiple times —each time with a better set of initial parameters— until a best-fit solution is found *both* for the existing photometric and spectroscopic observations.

the initialisation of the model parameters, we hope to make *cecilia* capable of calculating the RV shift of the white dwarf when no estimates are provided by the user.

## 7 CONCLUSIONS

In this work, we have presented *cecilia*, the first ML pipeline designed to measure the main physical and chemical properties of intermediate-temperature, He-rich polluted white dwarfs from their spectra. In particular, our code exploits the power of neural-network-based-interpolation to (i) rapidly generate synthetic spectral models of polluted white dwarfs in high-dimensional space, (ii) map the latter to a set of 13 underlying stellar properties ( $T_{\text{eff}}$ ,  $\log g$ ,  $\log_{10}(\text{H/He})$ , and 10 metal abundances); and (iii) generate a final spectral prediction through frequentist and Bayesian fitting techniques (see Sections 3 and 4).

With the development of *cecilia*, we have sought to tackle the dependence of conventional spectral analysis techniques on manual and iterative work, which in turn make them time-intensive, prone to human errors, and impossible to scale up to large samples of polluted white dwarfs. As explained in Section 5, *cecilia*'s architecture speeds up and automates these conventional methods by providing preliminary metal abundances in less than 2 minutes and a final spectroscopic solution in less than 10 hours. After testing the performance of our pipeline with noiseless synthetic data and with real spectroscopic observations of WD 1232+563, we have shown that *cecilia* can retrieve metal abundances with uncertainties better than 0.1 dex for up to 10 elements, with only beryllium proving hard to constrain. This level of accuracy is similar to that of state-of-the-art techniques, and it can only be improved in the future as we incorporate larger/better training datasets and integrate user feedback.

Acknowledging the main limitations of our pipeline (see Section 6), *cecilia* has the potential to enable population-wide studies of metal pollution without entailing any manual work. These studies will significantly expand the number of polluted white dwarfs with well-characterised abundances, thus playing an important role in improving our knowledge of the geology and chemistry of extrasolar material. Unveiling the diversity of extrasolar compositions is especially relevant in the context of upcoming wide-field astronomical surveys like DESI, WEAVE, or SDSS-V, which will acquire a large volume of spectroscopic observations in the coming years. While conventional white dwarf characterisation techniques are too slow to process these massive datasets, our ML pipeline offers a fast, automated, and reliable solution to measuring trace elements in the atmospheres of polluted white dwarfs.

Finally, as a feature extraction tool with a versatile ML architecture, *cecilia* can also be replicated and transferred to other disciplines beyond the realm of white dwarf and of planetary science. As we step into the world of Big Data, where the ability to uncover meaningful insights from observations will increasingly become more important, *cecilia* exemplifies the immense potential of combining powerful ML tools with state-cutting-edge scientific knowledge and analysis techniques.

## 8 DATA AVAILABILITY

The SDSS spectrum of WD 1232+563 can be downloaded from the SDSS DR18 online database.

## ACKNOWLEDGEMENTS

MBA thanks Vedant Chandra, Oriol Abril-Pla, Dr. Kishalay De, Dr. Ignasi Ribas, and Dr. Amy Bonsor for useful discussions. MBA is supported by the MIT Department of the Earth, Atmospheric, and Planetary Sciences, NASA grants 80NSSC22K1067 and 80NSSC22K0848, and the MIT William Asbjornsen Albert Memorial Fellowship. SB is a Banting Postdoctoral Fellow and a CITA National Fellow, supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). SX is supported by NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

This work has made use of the Montreal White Dwarf Database (Dufour et al. 2016).

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

This work has also made use of data from the Sloan Digital Sky Survey (SDSS), which is funded by the Alfred P. Sloan Foundation, the Heising-Simons Foundation, the National Science Foundation, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is [www.sdss.org](http://www.sdss.org). SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration, including the Carnegie Institution for Science, Chilean National Time Allocation Committee (CNTAC) ratified researchers, the Gotham Participation Group, Harvard University, Heidelberg University, The Johns Hopkins University, L'École polytechnique Fédérale de Lausanne (EPFL), Leibniz-Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Extraterrestrische Physik (MPE), Nanjing University, National Astronomical Observatories of China (NAOC), New Mexico State University, The Ohio State University, Pennsylvania State University, Smithsonian Astrophysical Observatory, Space Telescope Science Institute (STScI), the Stellar Astrophysics Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Illinois at Urbana-Champaign, University of Toronto, University of Utah, University of Virginia, Yale University, and Yunnan University.

This work has employed the following open-source software packages: Python (G. van Rossum 1995), numpy (Oliphant 2015), scipy (Virtanen et al. 2020), matplotlib (Hunter 2007), pandas (McKinney 2010), mpfit (Markwardt 2009), edmc (Vanderburg 2021), tensorflow (Abadi et al. 2015), and corner (Foreman-Mackey 2016; Luger et al. 2017).

## REFERENCES

- Abadi M., et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, <https://www.tensorflow.org/>
- Abazajian K. N., et al., 2009, *ApJS*, **182**, 543
- Almeida A., et al., 2023, *ApJS*, **267**, 44
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, **47**, 481
- Blouin S., 2020, *MNRAS*, **496**, 1881
- Blouin S., Dufour P., Allard N. F., 2018a, *ApJ*, **863**, 184

- Blouin S., Dufour P., Allard N. F., Kilic M., 2018b, *ApJ*, **867**, 161
- Blouin S., Dufour P., Allard N. F., Salim S., Rich R. M., Koopmans L. V. E., 2019a, *ApJ*, **872**, 188
- Blouin S., Dufour P., Thibeault C., Allard N. F., 2019b, *ApJ*, **878**, 63
- Campbell B., Walker G. A. H., Yang S., 1988, *ApJ*, **331**, 902
- Caron A., Bergeron P., Blouin S., Leggett S. K., 2023, *MNRAS*, **519**, 4529
- Chandra V., Hwang H.-C., Zakamska N. L., Budavári T., 2020, *MNRAS*, **497**, 2688
- Chandra V., et al., 2021, *ApJ*, **921**, 160
- Charbonneau D., Brown T. M., Latham D. W., Mayor M., 2000, *ApJ*, **529**, L45
- Chayer P., Fontaine G., Wesemael F., 1995, *ApJS*, **99**, 189
- Cooper A. P., et al., 2023, *The Astrophysical Journal*, 947, 37
- Coutu S., Dufour P., Bergeron P., Blouin S., Loranger E., Allard N. F., Dunlap B. H., 2019, *ApJ*, **885**, 74
- DESI Collaboration et al., 2016, *arXiv e-prints*, p. [arXiv:1611.00037](https://arxiv.org/abs/1611.00037)
- Dalton G., et al., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. p. 91470L ([arXiv:1412.0843](https://arxiv.org/abs/1412.0843)), doi:10.1117/12.2055132
- Desharnais S., Wesemael F., Chayer P., Kruk J. W., Saffer R. A., 2008, *ApJ*, **672**, 540
- Dorn C., Khan A., Heng K., Connolly J. A. D., Alibert Y., Benz W., Tackley P., 2015, *A&A*, **577**, A83
- Doyle A. E., et al., 2023, *arXiv e-prints*, p. [arXiv:2303.00063](https://arxiv.org/abs/2303.00063)
- Dufour P., Bergeron P., Fontaine G., 2005, *ApJ*, **627**, 404
- Dufour P., et al., 2007, *ApJ*, **663**, 1291
- Dufour P., Kilic M., Fontaine G., Bergeron P., Melis C., Bochanski J., 2012, *ApJ*, **749**, 6
- Dufour P., Blouin S., Coutu S., Fortin-Archambault M., Thibeault C., Bergeron P., Fontaine G., 2016, The Montreal White Dwarf Database: a Tool for the Community, doi:10.48550/ARXIV.1610.00986, <https://arxiv.org/abs/1610.00986>
- Eisenstein D. J., et al., 2006, *ApJS*, **167**, 40
- Elms A. K., Tremblay P.-E., Gänsicke B. T., Koester D., Hollands M. A., Gentile Fusillo N. P., Cunningham T., Apps K., 2022, *MNRAS*, **517**, 4557
- Farihi J., Brinkworth C. S., Gänsicke B. T., Marsh T. R., Girven J., Hoard D. W., Klein B., Koester D., 2011, *ApJ*, **728**, L8
- Farihi J., Gänsicke B. T., Koester D., 2013, *Science*, **342**, 218
- Foreman-Mackey D., 2016, *The Journal of Open Source Software*, **1**, 24
- Fortin-Archambault M., Dufour P., Xu S., 2020, *ApJ*, **888**, 47
- G. van Rossum 1995, Technical report, Python tutorial, Technical Report CS-R9526. Amsterdam
- Gaia Collaboration et al., 2016, *A&A*, **595**, A1
- Gaia Collaboration et al., 2023, *A&A*, **674**, A1
- Garcia-Zamora E. M., Torres S., Rebassa-Mansergas A., 2023, White dwarf Random Forest classification through Gaia spectral coefficients ([arXiv:2308.07090](https://arxiv.org/abs/2308.07090))
- Gelman A., Rubin D. B., 1992, *Statistical Science*, **7**, 457
- Genest-Beaulieu C., Bergeron P., 2019, *ApJ*, **882**, 106
- Giammichele N., Bergeron P., Dufour P., 2012, *ApJS*, **199**, 29
- Goodfellow I. J., Bengio Y., Courville A., 2016, Deep Learning. MIT Press, Cambridge, MA, USA
- Greenstein J. L., 1976, *ApJ*, **207**, L119
- Gunning D., Stefik M., Choi J., Miller T., Stumpf S., Yang G.-Z., 2019, Science robotics, **4**, eaay7120
- Hanin B., 2018, *arXiv e-prints*, p. [arXiv:1801.03744](https://arxiv.org/abs/1801.03744)
- Holberg J. B., Barstow M. A., Sion E. M., 1998, *ApJS*, **119**, 207
- Hollands M. A., Koester D., Alekseev V., Herbert E. L., Gänsicke B. T., 2017, *MNRAS*, **467**, 4970
- Hoskin M. J., et al., 2020, *MNRAS*, **499**, 171
- Hunter J. D., 2007, *Computing in Science & Engineering*, **9**, 90
- Izquierdo P., Toloza O., Gänsicke B. T., Rodríguez-Gil P., Farihi J., Koester D., Guo J., Redfield S., 2020, *Monthly Notices of the Royal Astronomical Society*, **501**, 4276
- Jarrett K., Kavukcuoglu K., Ranzato M., LeCun Y., 2009, in 2009 IEEE 12th International Conference on Computer Vision. pp 2146–2153, doi:10.1109/ICCV.2009.5459469
- Jura M., 2003, *ApJ*, **584**, L91
- Jura M., Young E. D., 2014, *Annual Review of Earth and Planetary Sciences*, **42**, 45
- Jura M., Xu S., Klein B., Koester D., Zuckerman B., 2012, *ApJ*, **750**, 69
- Kaiser N., et al., 2010, The Pan-STARRS wide-field optical/NIR imaging survey. p. 77330E, doi:10.1117/12.859188
- Kaiser B. C., Clemens J. C., Blouin S., Dufour P., Hegedus R. J., Reding J. S., Bédard A., 2021, *Science*, **371**, 168
- Kiefer J., Wolfowitz J., 1952, The Annals of Mathematical Statistics, **23**, 462
- Kingma D. P., Ba J., 2017, Adam: A Method for Stochastic Optimization ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980))
- Klein B., Jura M., Koester D., Zuckerman B., 2011, *ApJ*, **741**, 64
- Klein B. L., Doyle A. E., Zuckerman B., Dufour P., Blouin S., Melis C., Weinberger A. J., Young E. D., 2021, *ApJ*, **914**, 61
- Kleinman S. J., et al., 2013, *ApJS*, **204**, 5
- Koester D., 2009, *A&A*, **498**, 517
- Koester D., Napiwotzki R., Voss B., Homeier D., Reimers D., 2005, *A&A*, **439**, 317
- Koester D., Gänsicke B. T., Farihi J., 2014, *Astronomy & Astrophysics*, **566**, A34
- Kollmeier J. A., et al., 2017, *arXiv e-prints*, p. [arXiv:1711.03234](https://arxiv.org/abs/1711.03234)
- Kramida A., Yu. Ralchenko Reader J., and NIST ASD Team 2022, NIST Atomic Spectra Database (ver. 5.10), [Online]. Available: <https://physics.nist.gov/asd> [2016, January 31]. National Institute of Standards and Technology, Gaithersburg, MD.
- LeCun Y., Bengio Y., Hinton G., 2015, *Nature*, **521**, 436
- Liang Y., Winn J. N., Melchior P., 2023, *arXiv e-prints*, p. [arXiv:2311.18326](https://arxiv.org/abs/2311.18326)
- Lodders K., 2003, *ApJ*, **591**, 1220
- Luger R., Lustig-Yaeger J., Agol E., 2017, *ApJ*, **851**, 94
- Markwardt C. B., 2009, in Bohlender D. A., Durand D., Dowler P., eds, Astronomical Society of the Pacific Conference Series Vol. 411, Astronomical Data Analysis Software and Systems XVIII. p. 251 ([arXiv:0902.2850](https://arxiv.org/abs/0902.2850)), doi:10.48550/arXiv.0902.2850
- Martin D. C., et al., 2005, *ApJ*, **619**, L1
- McCarthy J., Minsky M. L., Rochester N., Shannon C. E., 2006, AI magazine, **27**, 12
- McKinney W., 2010, in van der Walt S., Millman J., eds, Proceedings of the 9th Python in Science Conference. pp 51–56
- Melis C., Dufour P., 2017, *ApJ*, **834**, 1
- Moré J. J., 1978, in , Vol. 630, Lecture Notes in Mathematics, Berlin Springer Verlag. pp 105–116, doi:10.1007/BFb0067700
- Nair V., Hinton G. E., 2010, in Proceedings of the 27th International Conference on International Conference on Machine Learning. ICML10. Omnipress, Madison, WI, USA, pp 807–814
- Ness M., Hogg D. W., Rix H. W., Ho A. Y. Q., Zasowski G., 2015, *ApJ*, **808**, 16
- Oliphant T. E., 2015, Guide to NumPy, 2nd edn. CreateSpace Independent Publishing Platform, North Charleston, SC, USA
- Palme H., Lodders K., Jones A., 2014, in Davis A. M., ed., , Vol. 2, Planets, Asteroids, Comets and The Solar System. pp 15–36
- Paquette C., Pelletier C., Fontaine G., Michaud G., 1986, *ApJS*, **61**, 197
- Raddi R., Gänsicke B. T., Koester D., Farihi J., Hermes J. J., Scaringi S., Breedt E., Girven J., 2015, *Monthly Notices of the Royal Astronomical Society*, **450**, 2083
- Rogers L. A., Seager S., 2010, *ApJ*, **712**, 974
- Samuel A. L., 1959, *IBM Journal of Research and Development*, **3**, 210
- Saranya A., Subhashini R., 2023, *Decision Analytics Journal*, **7**, 100230
- Saunon D., Blouin S., Tremblay P.-E., 2022, *Phys. Rep.*, **988**, 1
- Seager S., Kuchner M., Hier-Majumder C. A., Militzer B., 2007, *ApJ*, **669**, 1279
- Shallue C. J., Vanderburg A., 2018, *AJ*, **155**, 94
- Smith M. J., Geach J. E., 2023, *Royal Society Open Science*, **10**, 221454
- Tan C., Sun F., Kong T., Zhang W., Yang C., Liu C., 2018, A Survey on Deep Transfer Learning ([arXiv:1808.01974](https://arxiv.org/abs/1808.01974))
- Ter Braak C. J. F., 2006, *Statistics and Computing*, **16**, 239
- Ting Y.-S., Conroy C., Rix H.-W., Cargile P., 2019, *ApJ*, **879**, 69

- Turing A. M., 1950, *Mind*, 59, 433
- Vanderburg A., 2021, `avanderburg/edmcmc`: v1.0.0, doi:10.5281/zenodo.5599854, <https://doi.org/10.5281/zenodo.5599854>
- Vanderburg A., et al., 2015, *Nature*, 526, 546
- Vats D., Knudson C., 2018, *arXiv e-prints*, p. arXiv:1812.09384
- Vehtari A., Gelman A., Simpson D., Carpenter B., Bürkner P.-C., 2021, *Bayesian Analysis*, 16, 667
- Veras D., 2021, Planetary Systems Around White Dwarfs ([arXiv:2106.06550](https://arxiv.org/abs/2106.06550))
- Vilone G., Longo L., 2020, *arXiv e-prints*, p. arXiv:2006.00093
- Vincent O., Bergeron P., Dufour P., 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 760
- Virtanen P., et al., 2020, *Nature Methods*,
- Wilson D. J., Gänsicke B. T., Koester D., Toloza O., Pala A. F., Breedt E., Parsons S. G., 2015, *Monthly Notices of the Royal Astronomical Society*, 451, 3237
- Wolff B., Koester D., Liebert J., 2002, *A&A*, 385, 995
- Xiang M., et al., 2022, *A&A*, 662, A66
- Xu S., Jura M., Klein B., Koester D., Zuckerman B., 2013, *ApJ*, 766, 132
- Xu S., Jura M., Dufour P., Zuckerman B., 2016, *ApJ*, 816, L22
- Xu S., Zuckerman B., Dufour P., Young E. D., Klein B., Jura M., 2017, *ApJ*, 836, L7
- Xu S., Dufour P., Klein B., Melis C., Monson N. N., Zuckerman B., Young E. D., Jura M. A., 2019, *AJ*, 158, 242
- York D. G., et al., 2000, *AJ*, 120, 1579
- Zeng L., Seager S., 2008, *Publications of the Astronomical Society of the Pacific*, 120, 983
- Zeng L., Sasselov D. D., Jacobsen S. B., 2016, *ApJ*, 819, 127
- Zuckerman B., Koester D., Reid I. N., Hünsch M., 2003, *ApJ*, 596, 477
- Zuckerman B., Koester D., Melis C., Hansen B. M., Jura M., 2007, *The Astrophysical Journal*, 671, 872
- Zuckerman B., Melis C., Klein B., Koester D., Jura M., 2010, *ApJ*, 722, 725

This paper has been typeset from a  $\text{\LaTeX}$  file prepared by the author.