A Machine-Learning Compositional Study of Exoplanetary Material Accreted Onto Five Helium-Atmosphere White Dwarfs with `cecilia`

Mariona Badenas-Agusti,[1,2,3]* Siyi Xu (许偲艺),[4] Andrew Vanderburg,[2] Kishalay De,[2]
Patrick Dufour,[5] Laura K. Rogers,[1,4] Susana Hoyos,[2,6] Simon Blouin,[7] Javier Viaña,[2]
Amy Bonsor,[1] Ben Zuckerman[8]

[1]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK*
[2]*Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[3]*Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[4]*Gemini Observatory/NSF's NOIRLab, 950 N. Cherry Ave, Tucson, AZ, 85719, USA*
[5]*Département de Physique, Université de Montréal, Montréal, Québec H3C 3J7, Canada*
[6]*Department of Earth, Planetary, and Space Sciences, University of California, Los Angeles, CA 90095-1567, USA*
[7]*Department of Physics and Astronomy, University of Victoria, Victoria, BC V8W 2Y2, Canada*
[8]*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095-1562, USA*

ABSTRACT

We present the first application of the Machine Learning (ML) pipeline `cecilia` to determine the physical parameters and photospheric composition of five metal-polluted He-atmosphere white dwarfs without well-characterised elemental abundances. To achieve this, we perform a joint and iterative Bayesian fit to their *SDSS* (R=2,000) and *Keck/ESI* (R=4,500) optical spectra, covering the wavelength range from about 3,800 Å to 9,000 Å. Our analysis measures the abundances of at least two —and up to six— chemical elements in their atmospheres with a predictive accuracy similar to that of conventional WD analysis techniques (≈0.20 dex). The white dwarfs with the largest number of detected heavy elements are SDSS J0859+5732 and SDSS J2311−0041, which simultaneously exhibit O, Mg, Si, Ca, and Fe in their *Keck/ESI* spectra. For all systems, we find that the bulk composition of their pollutants is largely consistent with those of primitive CI chondrites to within 1-2σ. We also find evidence of statistically significant (> 2σ) oxygen excesses for SDSS J0859+5732 and SDSS J2311−0041, which could point to the accretion of oxygen-rich exoplanetary material. In the future, as wide-field astronomical surveys deliver millions of public WD spectra to the scientific community, `cecilia` aspires to unlock population-wide studies of polluted WDs, therefore helping to improve our statistical knowledge of extrasolar compositions.

Key words: stars: white dwarfs - stars: atmospheres - stars: abundances - techniques: spectroscopic - methods: data analysis - planets and satellites: composition

1 INTRODUCTION

For several decades, the search and characterisation of exoplanets has primarily relied on two indirect observational techniques: transit photometry, which measures periodic changes in the flux of a star due to a transiting planet (Charbonneau et al. 2000), and high-precision radial velocities, which track the Doppler shift of a star's spectral lines induced by a planet's gravitational pull (Campbell et al. 1988). These techniques provide, respectively, the radius and mass of the planet, which can then be combined to estimate its bulk density and infer its most plausible interior composition. At the heart of this characterisation process are theoretical Mass-Radius (MR) diagrams in which the planet's radius and mass are compared to synthetic interior models (Seager et al. 2007; Fortney et al. 2007; Zeng & Seager 2008; Zeng & Sasselov 2013; Dressing et al. 2015). These

diagrams have proved useful to broadly differentiate between rocky or gaseous planets (Zeng et al. 2016). However, they are strongly degenerate as they can yield more than one possible core composition for an identical mass and radius measurement (Rogers & Seager 2010; Dorn et al. 2015). This ambiguity is further complicated by the limited number of exoplanets with well-measured masses and radii (Jontof-Hutter 2019), which makes it difficult to distinguish between different planetary interiors.

As we seek to improve our knowledge of extrasolar bodies, it will be crucial to mitigate the degeneracies of conventional exoplanet characterisation techniques. Fortunately, spectroscopic observations of "polluted" White Dwarfs (WDs) provide a solution to this problem by facilitating accurate measurements of the bulk composition of exoplanetary material (e.g. Zuckerman et al. 2007; Klein et al. 2010; Gänsicke et al. 2012; Rogers et al. 2024a). WDs are the degenerate cooling remnants of low- and intermediate-mass (≤ 8 $M_\odot$) Main-Sequence (MS) stars (Weidemann & Koester 1983). Their radius is

*E-mail: mbadenas@mit.edu

similar to that of the Earth ($R_{WD} \approx 1\ R_{\oplus}$), but their mass is half that of the Sun ($M_{WD} \approx 0.6\ M_{\odot}$). As a result, they are extremely dense and compact, with surface gravities on the order 10^8 cm/s². Due to their strong gravitational fields, WDs should have pristine atmospheres composed only of lightweight elements (i.e. H and/or He), with elements heavier than He (or "metals;" atomic number Z>2) sinking rapidly towards the unobservable stellar interior in timescales much shorter than the evolutionary age of the star (Fontaine & Michaud 1979; Paquette et al. 1986).[1] Nevertheless, contrary to this expectation, observations suggest that between 25% to 50% of isolated WDs are contaminated with traces of metals, such as calcium (Ca), magnesium (Mg), oxygen (O), silicon (Si), or iron (Fe) (e.g. Zuckerman et al. 2003, 2010; Koester et al. 2014). At temperatures lower than $T_{eff} \lesssim 25,000$ K, this phenomenon is likely due to the recent or ongoing accretion of tidally disrupted material from a planetary system that survived the post-MS evolution of its host star (e.g. see reviews by Jura & Young 2014; Farihi 2016; Veras 2021; Xu et al. 2024). In hotter systems, the outward flow of radiative levitation pressure can also bring heavy elements to the surface (Chayer et al. 1995), while in cool ($T_{eff} \lesssim 10,000$ K), helium-dominated WDs, the presence of carbon can be attributed to convective dredge-up from the deep stellar interior (Pelletier et al. 1986; Camisassa et al. 2017; Bédard et al. 2022).

Since the discovery of three metals in the atmosphere of the He-rich white dwarf "van Maanen 2" (van Maanen 1917; Weidemann 1960), the field of polluted WDs has consolidated into a valuable discipline to infer the geology and chemistry of extrasolar bodies. This inference process typically consists of two stages: first, detailed atmosphere models are used to fit a WD spectrum and obtain the star's photospheric abundances (e.g. Koester 2009, 2010; Dufour et al. 2007). Next, the observed stellar abundances are used to constrain the bulk composition of the polluting body via WD accretion and diffusion equations (Koester 2009; Jura & Young 2014). In general, the most polluted WDs only exhibit one or two metals in their spectra (Williams et al. 2024), especially Ca and Mg in the optical (Coutu et al. 2019), and Si in the Ultraviolet (UV) (Koester et al. 2014). However, there are several dozen well-characterised polluted WDs with multiple heavy elements in their atmospheres (e.g. Xu et al. 2019; Putirka & Xu 2021), including GD 362 with a total of 19 metals observed at the same time (Becklin et al. 2005; Zuckerman et al. 2007; Xu et al. 2013, 2017; Melis & Dufour 2017). With the exception of a few systems rich in water ices (e.g. Farihi et al. 2013; Raddi et al. 2015), these discoveries have revealed that most WD pollutants are relatively dry and rocky (Jura 2006; Zuckerman et al. 2007; Jura & Xu 2012a; Xu et al. 2019; Swan et al. 2019), with compositions similar to those of bulk Earth or the CI chondrites —i.e. the most primordial type of meteorites in the Solar System.

As of today, spectroscopic analyses of polluted WDs have enabled the detection of more than 20 metals, revealing a diverse landscape of extrasolar compositions (e.g. Klein et al. 2021; Xu et al. 2019; Doyle et al. 2019). Despite these exciting discoveries, the scarcity of polluted WDs with multiple heavy elements in their spectra, combined with the low number of polluted WDs with well-measured elemental abundances, has impeded a statistical study of the pollution phenomenon. In this paper, we address this problem by exploiting the fast and automated Machine Learning (ML) pipeline `cecilia` (Badenas-Agusti et al. 2024, or BA24 thereafter). This pipeline is the first Neural-Network (NN)-based spectral interpolator capable of rapidly estimating the elemental abundances[2] of intermediate-temperature He-atmosphere polluted WDs ($10,000 \leq T_{eff} \leq 20,000$ K) from optical spectra covering the wavelength range between 3,000 Å and 9,000 Å. More specifically, we use `cecilia` to measure the atmospheric composition of five polluted WDs with existing *SDSS* spectra, newly acquired *Keck/ESI* observations, and no well-measured abundances in the literature.

This paper is organised as follows. In Section 2, we motivate the study of He-rich polluted WDs using `cecilia`'s ML and Bayesian framework. Section 3 details our target sample and the spectroscopic observations analysed in this work. In Section 4, we provide an overview of `cecilia` and describe our methodology for estimating the atmospheric composition of our targets. Section 5 presents `cecilia`'s best-fit models and the geochemical properties of the WD pollutants. In Section 6, we investigate the limitations of our compositional analysis and discuss potential improvements to our code. Finally, we summarise our work and conclude in Section 7.

2 MOTIVATION

Our choice to study He-atmosphere polluted WDs is driven by two key observational advantages. First, their lower photospheric opacity compared to H-rich WDs makes it easier to observe low levels of metal pollution (Dufour et al. 2012; Klein et al. 2021; Saumon et al. 2022). Second, their extended convection zones often result in longer metal diffusion timescales, which also facilitates the detection of metal pollution (e.g. Zuckerman et al. 2010). From a computational perspective, we leverage the fast and automated interpolation capabilities of `cecilia` to address the limitations of conventional WD characterisation methods. These "classical" tools have underpinned the field of WDs for several decades, offering unique insights into the bulk composition of extrasolar material. Nevertheless, they involve time-intensive and manual work, so they would be too impractical and prohibitively expensive to analyse large samples of polluted WDs. For example, in the coming years, multiple wide-field spectroscopic surveys will deliver an unprecedented amount of data to the WD community. This includes the Sloan Digital Sky Survey V (SDSS-V; Kollmeier et al. 2017; Chandra et al. 2021, the Dark Energy Spectroscopic Instrument (DESI; Cooper et al. 2023; DESI Collaboration et al. 2016a,b), or WEAVE (Dalton et al. 2014), which are expected to acquire spectra of about 100,000, 40,000, and 50,000 WDs, respectively, in the near future —some of which may exhibit signs of metal pollution. This vast amount of data would be intractable with conventional, "human-in-the-loop" methods. However, `cecilia` can obtain preliminary abundance measurements in less than a month using a single GPU (Badenas-Agusti et al. 2024), therefore offering a scalable solution to mine large databases with minimal human supervision.

3 DATA

3.1 Target Selection

The five He-atmosphere polluted WDs considered in this work are SDSS J0231+2512, SDSS J0859+5732, SDSS J1109+1318, SDSS J1333+6364, and SDSS J2311−0041. In Table 1, we present

---

[1] Diffusion timescales of heavy elements can vary from days to thousands of years in warm, H-atmosphere WDs (spectral type DA), to millions of years in cool, He-dominated systems (spectral type DB and DZ) (Koester 2009).

[2] In this paper, abundances are expressed as logarithmic number abundance ratios in base 10 relative to Helium, i.e. $\log_{10}(n(Z)/n(He))$.

Table 1. Main astrophysical properties of the five He-atmosphere polluted WDs studied in this work. References (Ref.): [1] The *Gaia* Mission (Gaia Collaboration et al. 2016, 2023), [2] The *SDSS* spectroscopic survey (York et al. 2000; Almeida et al. 2023), [3] The *GALEX* database (Martin et al. 2005), [4] The *Pan-STARRS* database (Kaiser et al. 2010), [5] This paper. In particular, $T_{\rm eff}$ and log g were obtained from an external fit to *Pan-STARRS* and *SDSS* photometry, while the WD mass and the cooling age were derived from the MWDD evolutionary models of Bédard et al. 2020 (see Section 4.3), [6] "MWDD HE" column in the MWDD (Dufour et al. 2016; see Section 4.3), [7] Spectral type from Coutu et al. 2019.

Property	Ref.	SDSS J0231+2512	SDSS J0859+5732	SDSS J1109+1318	SDSS J1333+6364	SDSS J2311−0041
Other Target Names						
Gaia DR3 ID	[1]	1023508230108687736	1037518722660955392	3965233688795064832	1665473315344805760	2650975899537545728
SDSS ID	[2]	J023154.82+251259.5	J085957.20+573249.9	J110957.82+131827.9	J133306.98+634936.4	J231141.58−004100.7
GALEX ID	[3]	J023154.8+251259	J085957.2+573249	J110957.8+131827	-	J231141.6−004100
Astrometric Properties						
R.A. [J2020; h:m:s]	[1]	02:31:54.82	08:59:57.19	11:09:57.83	13:33:06.99	23:11:41.58
Dec. [J2020; d:m:s]	[1]	+25:12:59.51	+57:32:50.01	+13:18:28.08	+63:49:36.38	−00:41:00.62
Parallax [mas]	[1]	3.16±0.24	3.14±0.27	3.35±0.22	3.39±0.12	4.23±0.18
Distance [pc]	[1]	316.59±24	318.94±27	298.36±20	294.94±11	236.46±10
$\mu_{\rm R.A.}$ [mas y^{-1}]	[1]	0.41±0.27	-4.32±0.23	-18.64±0.27	-32.54±0.17	-18.27±0.20
$\mu_{\rm Dec.}$ [mas y^{-1}]	[1]	8.22±0.27	-17.99±0.21	-32.79±0.21	-1.36±0.13	-23.86±0.19
Photometric Properties						
Gaia G$_{\rm mag}$	[1]	18.89±0.04	19.11±0.03	18.73±0.03	18.52±0.01	18.71±0.01
SDSS g$_{\rm mag}$	[2]	18.84±0.01	18.97±0.01	18.59±0.01	18.39±0.01	18.62±0.01
Pan-STARRS g$_{\rm mag}$	[4]	18.840±0.004	19.03±0.01	18.63±0.01	18.44±0.01	18.650±0.004
Physical Properties						
$T_{\rm eff}$ [K]	[5]	12620±503	12677±722	15112±1688	14762±1340	12023±544
log g [cgs]	[5]	$7.76^{+0.14}_{-0.13}$	$7.95^{+0.16}_{-0.15}$	$8.09^{+0.20}_{-0.19}$	$7.95^{+0.14}_{-0.13}$	$8.07^{+0.10}_{-0.09}$
Mass [$M_\odot$]	[5]	$0.45^{+0.07}_{-0.06}$	$0.56^{+0.10}_{-0.08}$	$0.64^{+0.13}_{-0.11}$	$0.56^{+0.08}_{-0.07}$	$0.62^{+0.06}_{-0.06}$
Cooling Age [Gyr]	[5]	$0.24^{+0.06}_{-0.04}$	$0.30^{+0.10}_{-0.07}$	$0.23^{+0.13}_{-0.09}$	$0.19^{+0.08}_{-0.06}$	$0.42^{+0.09}_{-0.07}$
Luminosity [$L_\odot$]	[5]	$0.005^{+0.002}_{-0.002}$	$0.004^{+0.002}_{-0.002}$	$0.007^{+0.005}_{-0.004}$	$0.007^{+0.004}_{-0.004}$	$0.003^{+0.001}_{-0.001}$
Spectral Type[a]		DB [6]	DBZ [7]	DBAH [6]	DBZA [7]	DBZ [7]
SDSS Observational Log						
Date [MJD]	[2]	2005 Nov 04	2003 Jan 23	2003 Dec 25	2000 Apr 05	2003 Nov 20
Spectrograph	[2]	*SDSS*	*SDSS*	*SDSS*	BOSS	*SDSS*
Plate ID	[2]	2399	483	1751	6822	381
MJD ID	[2]	53764	51924	53377	56711	51811
Fiber ID	[2]	157	463	9	196	103
Wavelength Coverage [Å]	[2]	3802-9191	3809-9215	3810-9206	3565-10341	3794-9198
Mean S/N	[5]	17.03	8.60	10.45	27.85	10.91
Keck/ESI Observational Log						
Date [UT]	[5]	2016 Nov 19	2016 Nov 18, 2017 Mar 7	2017 Mar 7	2016 Mar 28	2016 Jun 8, 2016 Nov 19
Time [s]	[5]	4,800	21,000	6,000	3,000	9,000
Standard Star	[5]	BD+28°4211	G191-B2B, Feige 34	HZ 44	BD+28°4211	BD+28°4211
Wavelength Coverage [Å]	[5]	3915-9299	3915-9299	3915-9299	3915-9299	3915-9299
Mean S/N	[5]	28	65	47	34	45

[a]: The estimated mass for SDSS J0231+2512 is below $0.5M_\odot$, which could point to the binary nature of this system (Rebassa-Mansergas et al. 2011). This scenario was also proposed by the ML study of Vincent et al. (2023), where SDSS J0231+2512 was autonomously classified as a DB+M dwarf binary based on its *SDSS* spectrum. Although we cannot fully rule out this hypothesis, we see no visual evidence of it in our spectra (e.g. no large RV variations, or M-dwarf-like flux modulations in the red part of the spectrum). Moreover, the Renormalized Unit Weight Error (`ruwe`) of this object is 1.05, which falls below the *Gaia* DR3 threshold for unresolved binaries (`ruwe>1.25`; Penoyre et al. 2020).

their full *Gaia* IDs and summarise their main astrophysical properties and spectroscopic observations.

Our targets were selected from the sample of Koester & Kepler (2015), who originally identified them as He-rich polluted WDs in the Sloan Digital Sky Survey (*SDSS*; York et al. 2000). From this sample, we chose about 20 objects for follow-up at higher resolution and signal-to-noise ratio (S/N) using the Echellette Spectrograph and Imager (ESI) on the Keck II Telescope (Bigelow & Nelson 1998; Sheinis et al. 2002) (programs: U067E, U131E, U153E, and U059; PI: B. Zuckerman). Among the observed systems, our five WDs exhibited the most interesting spectral features and were thus ideal candidates for a detailed abundance analysis with cecilia. Beyond their polluted nature, we also selected our targets based on two additional criteria. First, their effective temperatures (T_{eff}) and surface gravities (log g) satisfied cecilia's allowed parameter bounds (see Table 1 in BA24). Second, they all lacked well-measured photospheric compositions, with only estimates of their H and Ca abundances available in the Montreal White Dwarf Database (MWDD; Dufour et al. 2016) from the study of Koester & Kepler (2015).[3]

In our final target sample, only SDSS J1109+1318 had an estimated Ca abundance from Koester & Kepler (2015) (−6.46±0.5) slightly above cecilia's maximum bound for this metal (−7.00). However, we chose to include this object in our analysis because cecilia revises its Ca abundance to be well within its acceptable range (see Table 3). Morever, cecilia produces accurate fits to the observed Ca spectral lines (Figure A2) and yields a fully converged MCMC posterior distrubution that is not truncated at the upper limit (Figure A6).[4] For the remaining WDs, their published astrophysical properties, including their H and Ca abundances, satisfied cecilia's allowed parameter space at the time of this writing.

3.2 Spectroscopic Observations

3.2.1 SDSS

All five WDs were observed with the *SDSS* spectrograph (SDSS J0231+2512, SDSS J1109+1318, SDSS J0859+5732, SDSS J2311−0041) or the upgraded *BOSS* spectrograph (SDSS J1333+6364) mounted on the Sloan Foundation Telescope at Apache Point Observatory (York et al. 2000; Gunn et al. 2006). Their spectra (in vacuum) were downloaded as FITS files from the *SDSS* DR18 online database[5] and have a variable resolving power ($R \equiv \lambda/\Delta\lambda$) of R≈1,500 at 3,800 Å and R≈2,500 at 9,000 Å. In our optimisation routine, we assumed a fixed R=2,000 for each SDSS dataset (see Section 4.3).

3.2.2 Keck/ESI

The five polluted WDs in our sample were also observed with the *Keck/ESI* spectrograph (Sheinis et al. 2002) across 10 overlapping orders covering the wavelength range 3,900−11,000 Å. All datasets (in air) were obtained with a slit width of 1 arcseconds ($''$), resulting

in a resolving power of about R=4,500. The observing logs of the *Keck/ESI* observations are presented in Table 1.

After acquiring the data, we carried out basic reductions using the makee package,[6] including bias subtraction, flat fielding, spectrum extraction, wavelength calibration, and barycentric correction. We performed flux calibration on the spectra by computing their corresponding instrumental sensitivity function. To this end, we used *Keck/ESI* observations of spectrophotometric standards listed in the database of the European Southern Observatory,[7] obtained spatially and temporally close to the science target observations. For each WD, we computed the wavelength-dependent flux response of the instrument using STIS spectra from the CALSPEC STScI database (Bohlin et al. 2014).[8] The instrumental response function was then applied to the extracted spectra of the science targets to obtain a flux-calibrated spectrum. We then stacked all the orders together to generate a continuous spectrum by binning the wavelength, flux, and flux uncertainties over an output wavelength grid spanning the full wavelength range, with wavelength spacing defined by the lowest resolution order to conserve flux. For WDs that were observed multiple times (either during the same night or on different days; see Table 1), we followed a similar stacking strategy and combined all their existing observations into a final spectrum with an improved S/N.

3.3 Data Treatment

Next, we visually inspected the *SDSS* and reduced *Keck/ESI* spectra to minimise the impact of spurious effects, instrumental noise, and telluric contamination. We first eliminated scattered outliers, skylines, infinite (*NaN*) flux points, and flux error measurements equal to zero. We then cleaned the *SDSS* and *Keck/ESI* spectra from tellurics by clipping the spectral bands of $O_2\gamma + O_4$ (air wavelengths: 6,270-6,330 Å), O_2B (6,860-6,965 Å), H_2O (7,143-7,398 Å), O_2A (7,586-7,703 Å), and H_2O (8,085-8,420 Å and 8,915-9,000 Å) (Buton et al. 2013). In addition to these processing steps, we also addressed a well-known limitation of atmosphere models for He-rich WDs cooler than $T_{eff} \lesssim 16,000$ K (as is the case for the five targets in our sample) —namely, their inability to accurately reproduce the broadening of He I lines caused by collisions with neutral particles, such as He and H atoms (Bergeron et al. 2011; Genest-Beaulieu & Bergeron 2019a; Cukanovaite et al. 2020; Saumon et al. 2022). Given that our pipeline was trained on these atmosphere models, any inaccuracies in the treatment of He I broadening could propagate to cecilia's predicted elemental abundances (since the amount of metals in the WD affects the stellar electron density and photospheric opacity, and hence the shapes of the spectral lines; Hollands et al. 2017). Therefore, we sought to mitigate this issue with two complimentary strategies:

(i) *Clipping of prominent He I absorption lines*: First, we removed the strongest He I features from each spectrum. To perform this task, we applied a symmetric clipping window of 20 Å around the center of the lines, focusing on those at air wavelengths 3,819 Å, 4,026 Å, 4,387 Å, 4,713 Å, 4,921 Å, 5,015 Å, 5,047 Å, 5,875 Å, 6,678 Å, 7,065 Å, and 7,281 Å (Kramida et al. 2022). For the He I feature at 4,471 Å, we used an asymmetrical window of 20 Å and 13 Å to the left

---

[3] The MWDD (https://www.montrealwhitedwarfdatabase.org/) is the largest database of spectroscopically confirmed WDs to date.

[4] In line with the case of SDSS J1109+1318, systems that would *a priori* be discarded for having T_{eff}, log g, or abundance estimates near or beyond cecilia's bounds can be retained for atmospheric analysis if their MCMC posterior distributions do not pile up at the limits, but instead remain well inside cecilia's boundaries.

[5] https://skyserver.sdss.org/dr18/

[6] https://sites.astro.caltech.edu/~tb/makee/

[7] https://www.eso.org/sci/observing/tools/standards/spectra/stanlis.html

[8] https://archive.stsci.edu/hlsps/reference-atlases/cdbs/calspec/

Figure 1. Median-normalised optical spectra of the five polluted WDs in our sample (*left*: *SDSS* with an assumed fixed resolving power of R=2,000 (in vacuum); *right*: *Keck/ESI* with an assumed fixed R=4,500 in air). The stellar fluxes and their corresponding uncertainties are shown, respectively, in blue and red, while `cecilia`'s best-fit models are presented in dark green. The data gaps correspond to the discarded wavelength regions described in Section 3.3, including the strongest He I lines. The grey dashed vertical lines denote the 200 Å spectral windows used during `cecilia`'s training and optimisation routine.

and to the right of the center, respectively, in order to preserve a potential Mg absorption line near 4,481 Å. The final *SDSS* and *Keck/ESI* observations used in our analysis are shown in Figure 1.[9]

(ii) *Fixed photometric values for T_{eff} and* log g: Second, we froze the effective temperature and surface gravity of each WD to their photometric solution during `cecilia`'s fitting procedure (see Section 4.3). This decision was not only motivated by the unreliable nature of He-rich WD models below $T_{\text{eff}} \lesssim 16,000$ K, but also by the strong effect of He I lines in driving `cecilia`'s spectroscopic results for T_{eff} and log g. In particular, although we removed the strongest He I features from our spectra, our clipping method did not eliminate the full extent of their wings, allowing residual He I to affect `cecilia`'s predictions. In addition to this leftover He I absorption, we also observed that prominent metal lines (e.g. Ca II at about 3,934 Å and 3,969 Å) could similarly influence the inferred values for T_{eff} and log g. However, unlike the He I lines, these metal features could not be removed without losing important information about the WD's elemental abundances. Therefore, given that both residual He I absorption and strong metal lines could bias `cecilia`'s spectroscopic solution for T_{eff} and log g, we adopted a more conservative approach by fixing these two parameters to their photometric estimates. We note that we also explored leaving T_{eff} and log g free, but even with informative photometric priors based on external fits to the stellar photometry, `cecilia`'s results would still deviate significantly from their photometric solution. This behaviour further reinforced our decision to rely on the fixed photometric values for T_{eff} and log g in our optimisation routine.

4 METHODOLOGY

In this Section, we describe our procedure for estimating the elemental abundances of the five polluted WDs. First, we provide a brief summary of `cecilia` and present several new functionalities recently added to our code. We then discuss `cecilia`'s optimisation procedure.

4.1 An Overview of `cecilia`

In recent years, ML has improved our ability to address complex scientific problems, while also helping to reduce our dependence on time-intensive and manual data analysis techniques. These advancements have impacted all branches of astrophyhsics, offering new ways to exploit large astronomical datasets and learn about the fundamental properties of planets and stars. Despite this progress, however, the complex spectral features of polluted WDs have not previously been modeled with ML-based techniques. Therefore, we chose to develop `cecilia`, the first ML system capable of measuring the main astrophysical properties of He-rich polluted WDs from their spectra. Implemented with the open-source ML package `tensorflow`[10] (Abadi et al. 2015, 2016), and trained with MIT's Satori Supercomputer and one NVidia V100 GPU,[11] `cecilia`'s ML architecture is composed of three unsupervised, multi-layered (i.e. "deep") Neural Networks (NN): an Autoencoder, a Fully Connected Neural Network (FCNN1), and a Fine-Tuning Fully Connected Neural Network (FT FCNN2).[12] These networks use thousands of computational units (or "neurons") and non-linear "activation" functions to transform a set of input features (WD properties) into useful output parameters (a polluted WD spectrum).

To achieve a good balance between predictive accuracy and computational efficiency, we trained `cecilia` sequentially using 29 windows of 200 Å in the wavelength range between 3,000 Å and 9,000 Å. Our training, validation, and testing sets involved more than 22,000 randomly generated synthetic WD properties (or "labels"), together with their corresponding synthetic spectra (in air wavelengths). As described in BA24, `cecilia`'s labels consist of 13 independently varied parameters: the effective temperature T_{eff} and surface gravity log g of the WD; its logarithmic hydrogen abundance relative to He, $\log_{10}$ (H/He); and 10 additional number abundances for Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni. We also incorporated 14 additional elements (C, N, Li, Na, Al, P, S, Cl, Ar, K, Sc, V, Co, Cu) with their calcium abundance ratio (Z/Ca) scaled to their CI chondritic abundance from Lodders (2003). In addition to our choice of 13 stellar labels, we trained `cecilia` with high-resolution (R≈50,000) synthetic spectra featuring 55,000 points between 3,000 Å and 9,000 Å. These models were generated with the atmosphere code of Dufour et al. 2007; Blouin et al. 2018a,b, which is a local thermodynamic equilibrium code that self-consistently considers H, He, and metallic species in its equation of state and in the calculation of opacities. Metal lines are included using the Vienna Atomic Line Database (VALD; Piskunov et al. 1995; Kupka et al. 1999; Ryabchikova et al. 2015).

Once trained, `cecilia` exploits the speed of NN-based interpolation to rapidly produce high-resolution model spectra in a fully autonomous manner. For example, in the context of this work, `cecilia` can predict a synthetic *SDSS* or *Keck/ESI* spectrum in 0.17-0.19 seconds, with individual wavelength windows generated in less than 6 milliseconds on average. When factoring in additional computational overheads (e.g. estimation of RV shifts), the total processing time per spectrum reaches 0.35-0.37 seconds, including the evaluation of the MCMC log-likelihood function. This efficiency makes `cecilia` several orders of magnitude faster than conventional WD atmosphere codes, which require one to three hours on a single CPU core to produce an equivalent synthetic spectrum. As `cecilia` continues to evolve, further refinements in its ML architecture could help improve its speed, making it more efficient for large-scale WD spectral analysis. We also note that `cecilia`'s interpolation time is largely unaffected by variations in the input stellar labels (i.e. T_{eff}, log g, and elemental abundances). This stability is expected, as `cecilia`'s architecture leverages the learned latent-space representations of its training dataset, therefore bypassing the need for computationally intensive radiative transfer calculations. A summary of `cecilia`'s

---

[9] In a few cases, we intentionally retained a small number of weaker He I absorption lines when they were visible in some (but not all) the spectra, especially if they were close to metal lines. This includes the He I absorption feature at about 4,121 Å in Panel *c* of Figure 4, which is visible in some of the *Keck/ESI* spectra, and is close to Fe I and O II transitions. For consistency, we did not remove these weaker He I features when present, but we verified that their inclusion had a negligible effect on `cecilia`'s final results, including the best-fit elemental abundances.

[10] https://www.tensorflow.org/?hl=es
[11] https://mit-satori.github.io/.
[12] The Autoencoder has two components: an Encoder and a Decoder. Together, they are used for data compression and dimensionality reduction. The FCNNs are designed to produce a high-resolution synthetic spectrum from 13 stellar labels (i.e. T_{eff}, log g, $\log_{10}$ (H/He), and 10 metal abundances relative to He).

Figure 2. A summary of `cecilia`'s methodology for estimating the main astrophysical properties (or labels) of polluted He-rich WDs from multiple spectroscopic observations. We refer the reader to Badenas-Agusti et al. (2024) for a more comprehensive description of the pipeline.

framework is presented below and in Figure 2. We refer the reader to BA24 for a more detailed discussion of the pipeline.

(i) *Stage 1*: First, `cecilia` invokes its trained FT FCNN2 network[13] to produce a preliminary, high-resolution synthetic spectrum based on user-defined initial guesses for the 13 stellar labels of the WD (i.e. $T_{\rm eff}$, log g, and 11 logarithmic elemental abundances for H, Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni). If any of these labels are unknown (as would likely be the case for a newly discovered polluted WD), `cecilia` assumes chondritic abundance ratios except for H and Ca, for which it adopts their mean abundance ratio from its training set as initial guesses. We note that the assumption of chondritic abundance ratios is often used to approximate the photospheric abundances of elements when, for various reasons, these cannot be derived directly from a WD spectrum (Dufour et al. 2007; Coutu et al. 2019). Although several exceptions have been found —such as water-rich material (Farihi et al. 2013; Raddi et al. 2015; Gentile Fusillo et al. 2017), or differentiated bodies with crustal and core-like compositions (Zuckerman et al. 2011; Jura & Xu 2012b; Harrison et al. 2018)— this assumption is often used given the similarity between the bulk composition of WD pollutants, and those of bulk Earth and the CI chondrites (Allègre et al. 2001; Lodders 2003), especially for rock-forming elements like Ca, Mg, Fe, Si, and O (Zuckerman et al. 2007; Xu et al. 2019; Doyle et al. 2023).

(ii) *Stage 2*: Second, `cecilia` implements several data processing techinques to enable a meaningful comparison between its initial prediction and the observed polluted WD spectrum. For example,

it (i) denormalises its raw ML output; (ii) smears and resamples its denormalised prediction to the the resolving power and wavelength grid of the observed spectrum, respectively; (iii) applies a radial velocity shift (or RV thereafter) to its prediction; and (iv) employs a linear model (or alternatively, an n-th degree polynomial as decided by the user) to correct the spectral "jumps" arising from the training of `cecilia` in independent windows of 200 Å.

(iii) *Stage 3*: Finally, `cecilia` employs a combination of frequentist and Bayesian statistical techniques to optimise its smeared, downsampled, and slope-corrected ML prediction. In particular, it incorporates two fitting routines: an optional and fast non-linear least-squares Levenberg-Marquardt method implemented by the Python `mpfit`[14] library (Moré 1978; Markwardt 2009), and a differential evolution Markov Chain Monte Carlo (MCMC) sampler executed by the Python `edmcmc` package (Vanderburg 2021; Ter Braak 2006). If the user runs `mpfit`, the resulting best-fit parameters are fed into the MCMC as improved initial guesses.

For a given white dwarf, `cecilia` takes, on average, less than a day to produce a complete spectroscopic solution using a single GPU. Moreover, `cecilia` can be scaled to simultaneously process multiple objects if additional GPUs are available, making it a valuable tool in the era of large-scale astronomical surveys. The total execution time is primarily dependent on the computational cost associated to the ML-based spectral interpolation process, rather than on the spectral resolution or S/N of the observations. As discussed in BA24, `cecilia`'s retrieval accuracy can be ≲0.1 dex for up to 10 chemical elements (H, Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni), with Be

---

[13] We note that `cecilia` only employs its FT FCNN2 architecture to generate a spectral prediction. The Autoencoder and FCNN1 are only used during the training stage to improve the overall learning accuracy of our pipeline.

[14] https://github.com/segasai/astrolibpy/blob/master/mpfit/mpfit.py

appearing the hardest element to constrain.[15] This performance is comparable to that of conventional techniques, which usually yield uncertainties of about 0.10 dex-0.20 dex for He-dominated polluted WDs, both in the UV and in the optical, as well as with different types of WD 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).

Beyond `cecilia`'s predictive capabilities, our code differs from conventional WD analysis techniques in two important ways. First, it can quickly fit the spectra of polluted WDs without the need for human supervision. This sets it apart from classical methods, which can be slow and time-intensive due to their reliance on manual and iterative fitting procedures. Second, `cecilia` provides a Bayesian treatment of observational evidence and parameter uncertainties. Through its MCMC, it maximises the log-likelihood function of the free model parameters ($\ln(\mathcal{L})$), while also incorporating the information encapsulated in the spectrum and any "prior" assumptions about the properties of the star. This approach allows for a thorough numerical exploration of the multidimensional parameter space of the stellar labels, yielding full posterior probability distributions, robust statistical errors, and insights into potential correlations between model parameters. These degeneracies are often visualised in scatterplot matrices, or "corner plots," which provide the two-dimensional marginal posterior of a pair of labels alongside their respective one-dimensional histogram distributions (e.g. see Figure 4). Finally, `cecilia` fits the *entire* spectrum simultaneously to determine the 13 stellar labels of the WD. This differs from some classical methods designed to calculate average abundance measurements based on individual, line-by-line fits to visible absorption lines.[16]

4.2 `cecilia`'s Improved Capabilities

In this paper, we present an improved version of `cecilia`'s optimisation procedure (see Figure 2). To begin with, we have modified the MCMC to sample the elemental abundances in linear space. This is different to the pipeline presented in BA24, where logarithmic abundances were used to map out the likelihood distribution of the model parameters. To implement this change, we have updated `cecilia` to: (i) adjust the allowed parameter bounds to linear values (see Table 1 in BA24); (ii) convert the user's initial abundance guesses from log to linear space; and (iii) impose a lower bound of zero on all abundances (representing the complete absence of an element in the atmosphere of the WD). A practical motivation for this switch is that sampling in linear space allows the MCMC to explore very low abundance values and *still* achieve convergence. This is different from our previous log-abundance approach, which could cause the MCMC to diverge towards negative infinity in the case of weak or non-detections. Another advantage of our new approach is that it opens the door to computing statistically robust upper limits on the abundances of the undetected elements. However, we do not attempt to calculate these limits in this work due to computational limitations related to the sparseness of `cecilia`'s training set at very low abundance values.[17]

In addition to sampling the elemental abundances in linear space, we have also implemented two important changes to `cecilia`'s log-likelihood function. First, we have modified our code to allow for a simultaneous fitting of N_{spec} spectra from multiple instrumental facilities —regardless of their intrinsic characteristics (e.g. resolving power, S/N). Second, we have introduced a jitter term A_{jit}^{s} (where s denotes the instrument that acquired the spectrum) to account for unknown sources of noise in the data and in the atmosphere models used to train `cecilia`. The jitter parameter employs the observed scatter in the spectrum to estimate what the uncertainties of the stellar flux should be, lowering or increasing them accordingly if necessary (Ford 2006). Therefore, it regulates the contribution of each spectrum to the final solution and allows `cecilia` to robustly fit multiple spectra simultaneously, even if some datasets are significantly poorer than the rest. Ideally, the jitter term should be close to 1 when the flux errors are well estimated prior to initiating `cecilia`'s optimisation routine. As an additional diagnostic metric, our code also computes the reduced chi-squared ($\chi_{\text{reduced}}^{2}$) for the best-fit solution associated to each spectrum, with the expectation that $\chi_{\text{reduced}}^{2} \approx 1$ when all the observations have contributed similarly to the best-fit solution.

With all the aforementioned changes, `cecilia`'s new log-likelihood MCMC function is given by Eq. 1, where $f_{\text{synth,corr}}^{s,i}$ denotes `cecilia`'s prediction for pixel i from instrument s (after denormalisation, resampling, resolution downgrading, RV-shifting, and slope correction). The terms $f_{\text{obs}}^{s,i}$ and $f_{\text{obs,err}}^{s,i}$ represent the observed stellar flux and its corresponding error, while the priors $\pi_{0,\text{phot}}(T_{\text{eff}}, \log g)$ and $\pi_{0,\text{chondr}}(X_{k})$ are optional contributions to `cecilia`'s spectral model, applied to the parameters T_{eff} and $\log g$, and to the elemental abundances X_{k}, respectively. The photometric priors may be used when T_{eff} and $\log g$ are treated as free model parameters and have reliable external constraints (e.g., from existing photometry).[18] The chondritic priors, in turn, are designed to limit `cecilia`'s exploration of the parameter space and help with MCMC chain convergence.

$$\ln \mathscr{L} = -\frac{1}{2} \left\{ \sum_{s=1}^{N_{\text{spec}}} \sum_{i=1}^{N_{\text{points}}^{(s)}} \left[\left(\frac{f_{\text{obs}}^{s,i} - f_{\text{synth,corr}}^{s,i}}{A_{\text{jit}}^{s} \cdot f_{\text{obs,err}}^{s,i}} \right)^{2} + 2 \ln \left(A_{\text{jit}}^{s} \cdot f_{\text{obs,err}}^{s,i} \right) \right] + \sum_{j=1}^{N_{\text{phot}}} \pi_{0,\text{phot}}(T_{\text{eff}}, \log g) + \sum_{k=1}^{N_{\text{elem}}} \pi_{0,\text{chondr}}(X_{k}) \right\}$$
(1)

In general, the influence of a chondritic prior on an elemental abundance is determined by its inverse variance, $1/\sigma_{\text{prior,chondr}}^{2}$, where $\sigma_{\text{prior,chondr}}$ denotes the standard deviation (or width) of the prior distribution $\pi_{0,\text{chondr}}$. Assuming that both the prior and the likelihood are Gaussian, the posterior uncertainty, σ_{post}, is given by

---

[15] `cecilia`'s performance is discussed in more detail in Section 5 of BA24, based on an analysis of synthetic and real spectroscopic observations.

[16] We note that other conventional (i.e. non-AI-based) methods are also capable of fitting the entire spectrum simultaneously. For instance, Bhattacharjee et al. 2025 recently employed a least-squares minimisation approach to model the spectra of DAZ and DBZ white dwarfs.

[17] We note that our new parametrisation modifies the implicit prior used in our sampling. Indeed, sampling uniformly in logarithmic space (as done in

BA24) corresponds to an implicit prior proportional to $1/x$ in linear space, which favours smaller abundance values. In contrast, our new approach samples uniformly in linear space, which corresponds to an implicit prior that increases exponentially with $\log x$ and favours higher abundances. This distinction is only important when the data provide weak or no constraints on the elemental abundances. Indeed, when the observations are sufficiently informative, the posterior is dominated by the likelihood and the influence of the prior becomes negligible.

[18] The current implementation of `cecilia` does not predict T_{eff} and $\log g$ from photometric observations. Therefore, our code assumes that the user has reasonably good constraints on these two parameters from external photometric fits.

$$\frac{1}{\sigma_{\text{post}}^2} = \frac{1}{\sigma_{\text{obs}}^2} + \frac{1}{\sigma_{\text{prior,chondr}}^2}, \qquad (2)$$

where σ_{obs} represents the observational uncertainty. For elements that are typically easy to detect (e.g., Ca, Mg, Si), we recommend adopting broad chondritic priors (i.e., a large $\sigma_{\text{prior,chondr}}$) so that $1/\sigma_{\text{prior,chondr}}^2$ is very small compared to $1/\sigma_{\text{obs}}^2$ and $\sigma_{\text{post}} \approx \sigma_{\text{obs}}$. This approach minimises the effect of the prior and ensures that the posterior uncertainty is dominated by the observations. In contrast, for elements that are hard to detect, we suggest using narrower chondritic priors (i.e., a small $\sigma_{\text{prior,chondr}}$) to facilitate MCMC convergence; in these cases, if an element is undetected, we recommend not reporting its predicted `cecilia` abundance. In this work, we discuss our choice of prior widths in Section 4.3 and Table 2.

4.3 Estimation of White Dwarf Elemental Abundances

We measured the elemental abundances of the five polluted WDs by performing a joint MCMC fit to their *Keck/ESI* and *SDSS* spectra. Our `cecilia` MCMC model consisted of 17 parameters, namely: the 13 stellar labels underlying `cecilia`'s training set (i.e. T_{eff}, log g, and the logarithmic abundances of H, Ca, Mg, Fe, O, Si, Ti, Be, Cr, Mn, and Ni relative to He); 2 jitter terms ($A_{\text{jit,SDSS}}$, $A_{\text{jit,Keck/ESI}}$); and 2 RV shifts (RV$_{\text{SDSS}}$, RV$_{\text{Keck/ESI}}$).[19] As described in Section 3.3, we decided to freeze T_{eff} and log g to their photometric solution, and only fitted the remaining 15 parameters.

To prepare our initial guesses, we adopted the T_{eff}, log g, H, and Ca results of Coutu et al. (2019; C19), whenever possible.[20] There were two systems, however, lacking C19 measurements: SDSS J0231+2512 and SDSS J1109+1318. For the former, we took as initial guesses the T_{eff} and log g values derived from a fit to *Pan-STARRS* photometry (see the "MWDD He" column in the MWDD), and the H and Ca results from Koester & Kepler (2015). For SDSS J1109+1318, we used the T_{eff} and log g photometric solution of Genest-Beaulieu & Bergeron (2019b), and the H and Ca abundances of Koester & Kepler (2015). Then, for the remaining (and unknown) elemental abundances, we assumed chondritic abundances relative to our initial guess for Ca. Finally, we set the jitter terms to 1 (dimensionless) and the RV shifts to 0 km/s.

After building our spectral model, we executed `cecilia`'s MCMC twice per WD, using the log-likelihood in Eq. 1 and a linear function to minimise the spectral jumps between training windows of 200 Å. Our MCMC hyperparameters consisted of n_{walkers}=50 walkers (or chains), n_{draws}=3000 draws (number of steps per chain), and a 20% "burn-in" phase aimed at removing non-stationary solutions (i.e. $n_{\text{burn}} = 0.2 n_{\text{draws}}$=600 draws).[21] Similarly to BA24, we initialised the positions of the walkers using Gaussian balls with standard deviations of 0.01 dex for the elemental abundances, 1 km/s for the two RV shifts, and 0.01 for the jitter terms. We then imposed priors on

---

[19] Although we would expect the RV of a WD to remain consistent across its spectra, we considered separate RV shifts for each dataset to account for known systematic offsets in our calibrations arising from different observational set-ups, instrumental effects, etc. As explained in the text, we do not report or provide a physical interpretation of our best-fit RV results; instead, we only treat them as nuisance parameters in our model.

[20] In C19, the authors studied 1,023 DBZ/DZ(A) WDs with the same atmosphere code used to generate `cecilia`'s training set; therefore, their results represent the most self-consistent initial guesses for our Bayesian analysis.

[21] The use of multiple walkers (or chains) ensures extensive sampling of the parameter space. Over time, the cumulative distribution of the walkers and draws should approach the final posterior distribution.

Table 2. Adopted prior distributions for `cecilia`'s model parameters (see Section 4.3). The symbols $\mathcal{T}$ and $\mathcal{U}$ denote truncated Gaussian and uniform priors, respectively. Where applicable, the mean μ represents the offset between the chondritic abundance of element X and that of calcium, defined as $\mu = \log(\text{Ca/He}) - \log(X/\text{He})$ based on the chondritic values of Lodders (2003). The standard deviation σ specifies the width of the truncated Gaussian distribution. The bounds of the elemental abundances are given in logarithmic base 10.

Parameter	Prior	Min.	Max.	Mean (μ)	Width (σ)
T_{eff} [K]		\multicolumn{4}{c}{No Prior Used – Fixed Parameter}			
log g [cgs]		\multicolumn{4}{c}{No Prior Used – Fixed Parameter}			
$\log_{10}$ (H/He) [dex]	$\mathcal{T}$	-7.00	-3.00	-1.96	2
$\log_{10}$ (Be/He) [dex]	$\mathcal{T}$	-23.85	-5.61	4.91	0.5
$\log_{10}$ (O/He) [dex]	$\mathcal{T}$	-17.46	1.25	-2.10	2
$\log_{10}$ (Mg/He) [dex]	$\mathcal{T}$	-16.89	-0.17	-1.24	2
$\log_{10}$ (Si/He) [dex]	$\mathcal{T}$	-17.42	-0.51	-1.22	2
$\log_{10}$ (Ca/He) [dex]	$\mathcal{U}$	-12.00	-7.00	–	–
$\log_{10}$ (Ti/He) [dex]	$\mathcal{T}$	-19.35	-2.32	1.40	0.5
$\log_{10}$ (Cr/He) [dex]	$\mathcal{T}$	-19.61	-1.71	0.66	0.5
$\log_{10}$ (Mn/He) [dex]	$\mathcal{T}$	-20.08	-1.65	0.82	0.5
$\log_{10}$ (Fe/He) [dex]	$\mathcal{T}$	-18.20	0.18	-1.16	2
$\log_{10}$ (Ni/He) [dex]	$\mathcal{T}$	-18.96	-1.66	0.10	0.5
RV$_{\text{SDSS}}$ [km/s]	$\mathcal{U}$	-500	500	–	–
RV$_{\text{Keck/ESI}}$ [km/s]	$\mathcal{U}$	-500	500	–	–
$A_{\text{jit,SDSS}}$ [-]	$\mathcal{U}$	0	$+\infty$	–	–
$A_{\text{jit,Keck/ESI}}$ [-]	$\mathcal{U}$	0	$+\infty$	–	–

our free model parameters to limit `cecilia`'s exploration of the parameter space and facilitate MCMC chain convergence. In particular, we applied uniform priors to the Ca abundance, the RV shifts, and the jitter terms, allowing their values to vary within the bounds defined by `cecilia`'s training set. For all other elemental abundances, we adopted truncated Gaussian priors informed by the chondritic abundances of Lodders (2003). To do this, we set the mean of each distribution to the difference between the element's chondritic abundance and that of calcium. We then applied narrow widths of 0.5 dex for elements that are typically hard to detect (Ti, Be, Cr, Mn, Ni), and broader widths of 2 dex to more commonly observed elements (H, Mg, Fe, O, Si). A summary of our priors is presented in Table 2.

For our first MCMC, we fitted all model parameters except T_{eff} and log g, which we froze to our initial guesses. We then recycled `cecilia`'s best-fit results to run a second MCMC with improved values for T_{eff} and log g (see Tables 1-3). To refine these two parameters, we fitted the *SDSS* and *Pan-STARRS* photometry of each star following the procedure of C19, but correcting it for interstellar extinction using the `stilism` reddening models.[22] Our `cecilia` MCMC analysis took between 12.26 hr (SDSS J1109+1318) and 22.68 hr (SDSS J0859+5732) to complete. We then assessed MCMC chain convergence by computing the Gelman-Rubin potential scale reduction factor $\hat{R}$ for each free model parameter and ensuring that it satisfied $\hat{R} < 1.02$ (Gelman & Rubin 1992; Gelman et al. 2004).

Upon concluding our fits, we considered an element to be detected

---

[22] https://stilism.obspm.fr/. For WDs located beyond d$\gtrsim$100 pc, it is important to account for interstellar extinction by de-reddening the photometry with 3D dust maps (Coutu et al. 2019; Genest-Beaulieu & Bergeron 2019a; Gentile Fusillo et al. 2019).

if its elemental abundance had a statistical (i.e. MCMC) uncertainty lower than than or equal to an assumed detectability threshold of $\sigma_{\text{detection}} = 0.10$ dex. We also validated this cut-off empirically by confirming the presence of its absorption feature(s) by eye. Then, for each detected element, we used `cecilia`'s predicted abundance to estimate its mass ratio relative to He ($R_{\text{ratio,Z}}$; Eq. 3), its mass in the stellar atmosphere (M_{Z}; Eq. 4), and its accretion rate ($\dot{M}_{\text{Z}}$; Eq. 5), together with their 1σ errors from their 16th, 50th, and 84th quantiles of their MCMC posterior distributions. More specifically, we first calculated the mass ratio of a given metal Z with

$$R_{\text{ratio,Z}} = n_{\text{Z}} \cdot \frac{u_{\text{Z}}}{u_{\text{He}}}. \tag{3}$$

where u_{Z} represents the atomic mass of the metal, u_{He} is the atomic mass of He (i.e. $u_{\text{He}} = 4$ atomic mass units, or amu), and n_{Z} is the metal elemental abundance in linear space (as measured by `cecilia`). Second, we determined the mass of a given metal (M_{Z}) by invoking the mass of the stellar Convection Zone (M_{cvz}). In He-atmosphere WDs, CVZs tend to be very large (Koester 2009) and can therefore be used as a proxy for the total mass of He in the atmosphere,

$$M_{\text{Z}} = M_{\text{cvz}} \cdot R_{\text{ratio,Z}}. \tag{4}$$

To estimate M_{cvz}, we employed the evolutionary models of Bédard et al. (2020) using our improved photometric solution for T_{eff} and $\log g$.[23] We then inferred the accretion rate of each detected metal ($\dot{M}_{\text{Z}}$) by taking the ratio of its total mass (M_{Z}) to its gravitational sinking timescale (τ_{Z}), where τ_{Z} was obtained from the public diffusion models of the MWDD,[24]

$$\dot{M}_{\text{Z}} = \frac{M_{\text{Z}}}{\tau_{\text{Z}}}. \tag{5}$$

Finally, for each WD in our sample, we computed a lower limit on the total metal accretion rate ($\sum \dot{M}_{\text{Z}}$), along iwth the total mass of metals in the atmosphere ($\sum M_{\text{Z}}$). From the latter, we derived the fractional mass contribution of each element (MF_{Z}) using Eq. 6. For this calculation, we considered the 10 metals fitted by `cecilia`, with the best-fit masses of the undetected elements constrained via our chondritic abundance priors.

$$\text{MF}_{\text{Z}} = \frac{M_{\text{Z}}}{\sum M_{\text{Z}}}. \tag{6}$$

4.4 Estimation of Pollutants' Elemental Abundances

The photospheric abundances of a WD can be used to probe the bulk composition of its pollutant through a set of well-known WD accretion and diffusion equations (Dupuis et al. 1992, 1993a,b; Koester 2009). Assuming a constant accretion rate and the engulfment of a single pollutant,[25] we evaluated these equations for three different

---

[23] The models of Bédard et al. (2020) are publicly available in the MWDD (https://www.montrealwhitedwarfdatabase.org/evolution.html).

[24] Eq. 5 quantifies the rate at which metals diffuse out of the convective layer, and it only corresponds to the true accretion rate under the assumption of steady-state equilibrium, that is, when the influx of metals is balanced by their diffusion in the CVZ. As discussed in Section 5, it is likely that our targets are in the build-up or steady-state phase, which makes Eq. 5 a reasonable approximation.

[25] Previous work has investigated the accretion of multiple bodies (e.g. small asteroids) with different bulk compositions (e.g. Jura 2008; Wyatt et al. 2014; Turner & Wyatt 2019; Trierweiler et al. 2022). However, for simplicity, we assume that our WDs accreted a single large object.

phases. First, there is a "build-up" (or "increasing") phase, when the polluting body has just begun to accrete onto the WD's photosphere. During this stage, the number abundance ratios n_{A} and n_{B} of the pollutant for two different metals A and B are assumed to correspond to the elemental abundances in the atmosphere of the WD,

$$\left(\frac{n_{\text{A}}}{n_{\text{B}}}\right)_{\text{P}} = \left(\frac{n_{\text{A}}}{n_{\text{B}}}\right)_{\text{WD}}, \tag{7}$$

where the subscript "WD" alludes to the observed number abundances in the star —in our case, obtained from `cecilia` (see Table 3)— and "P" refers to the inferred composition of the polluting body. After the increasing phase, there is "steady-state" phase, when the effects of accretion and diffusion are (almost) in balance in the convective layer. During this period, the number abundance ratios of the pollutant are assumed to match those of the star when corrected by their respective diffusion timescales,

$$\left(\frac{n_{\text{A}}}{n_{\text{B}}}\right)_{\text{P}} = \left(\frac{n_{\text{A}}}{n_{\text{B}}} \cdot \frac{\tau_{\text{B}}}{\tau_{\text{A}}}\right)_{\text{WD}}. \tag{8}$$

The last phase of accretion is known as the "declining" (or "decreasing") phase, when the polluting material is no longer replenished in the stellar photosphere and sinks donwards into the interior, hence disappearing from view. During this phase, lightweight elements (i.e. H and He) remain in the outermost layers of the WD, while heavier elements rapidly sink below the photosphere. If t is the time elapsed since the end of accretion, the ratio of heavy elements in the polluting body can be inferred from the stellar abundance ratios after introducing an exponential decay term,

$$\left(\frac{n_{\text{A}}}{n_{\text{B}}}\right)_{\text{P}} = \left(\frac{n_{\text{A}}}{n_{\text{B}}} \cdot \frac{e^{-t/\tau_{\text{B}}}}{e^{-t/\tau_{\text{A}}}}\right)_{\text{WD}}. \tag{9}$$

5 RESULTS

5.1 Bayesian Spectral Modelling with `cecilia`

The results of `cecilia`'s MCMCs are summarised in Table 3, together with their corresponding total uncertainties (σ_{tot}). For each free model parameter, we assumed Gaussian and uncorrelated noise to approximate σ_{tot} as the quadrature sum of its statistical ($\sigma_{\text{stat,MCMC}}$) and systematic error ($\sigma_{\text{sys,ML}}$). The former are directly obtained from the 1σ confidence intervals of our MCMC and encapsulate the errors of the observations. In contrast, the latter arise from inaccuracies in `cecilia`'s ML predictions as well as from inherent imperfections in the WD atmosphere models used to train its networks. To ensure that `cecilia`'s uncertainties would not be greatly underestimated for our abundance parameters, we followed the methodology of BA24 and replaced their total error σ_{tot} by a conservative noise floor of $\sigma_{\text{floor}} = 0.20$ dex whenever $\sigma_{\text{tot}} < \sigma_{\text{floor}}$.

The full spectroscopic solutions of `cecilia`'s optimisation routine are illustrated in Figure 1. In Figure 3 and Figures A1-A4, we provide zoomed-in panels of the *Keck/ESI* best-fit models across different wavelength ranges. To visualise the quality of the model uncertainties, we also show `cecilia`'s predictions when modifying the abundances of the *detected* elements by a factor of $\pm 1\sigma_{\text{tot}}$. We note that `cecilia`'s detections are primarily enabled by the quality of the *Keck/ESI* observations. The *SDSS* spectra are especially helpful for the analysis of the Ca II H&K region between 3,930 Å and 3,970 Å, where the *Keck/ESI* spectrograph has very limited sensitivity. However, they are too noisy elsewhere to contribute significantly to `cecilia`'s detections. Finally, Figure 4 and Figures A5-A8 in the

Figure 3. A selection of spectral windows showing `cecilia`'s best-fit RV-shifted MCMC model (in green) for the median-normalised *Keck/ESI* spectrum of SDSS J0859+5732 (in light blue; in air wavelengths). For reference, we also include `cecilia`'s predictions when modifying the abundances of the *detected* elements by $\pm 1 \sigma_{\mathrm{tot}}$ (red and orange). The green labels show all the detected elements, defined as those with $\sigma_{\mathrm{stat,MCMC}} \leq \sigma_{\mathrm{detection}} = 0.10$ dex and at least one visible absorption line in the spectrum. We note that `cecilia` struggles to model the depth of the Mg line at about 4,481 Å (see panel *d*). This is not the case for other Mg lines, which are fitted reasonably well by our code (e.g. panels *g* and *j*). Such behaviour may be an example of underestimated errors due to the high level of red (i.e. correlated) noise in panel *d*. In Section 6.2, we discuss how `cecilia` can be improved to address this problem.

Table 3. `cecilia`'s best-fit parameters for the five polluted WDs in our sample, obtained from a joint MCMC fit to their *SDSS* and *Keck/ESI* spectra. For clarity, we only report the abundances of the detected elements, as `cecilia`'s results for the remaining chemical species are dominated by our choice of chondritic priors. For each detection, we provide three sources of uncertainty: the statistical errors from `cecilia`'s MCMC ($\sigma_{\mathrm{stat,MCMC}}$; left parenthesis), the systematic errors caused by imperfections in `cecilia`'s ML interpolation and our training models ($\sigma_{\mathrm{sys,ML}}$; center), and the total uncertainties (σ_{tot}; right), with the latter computed as the quadrature sum of $\sigma_{\mathrm{stat,MCMC}}$ and $\sigma_{\mathrm{sys,ML}}$. In all cases, σ_{tot} is consistently below our assumed noise floor ($\sigma_{\mathrm{floor}} = 0.20$ dex), so we follow the methodology of BA24 and enforce $\sigma_{\mathrm{tot}} = \sigma_{\mathrm{floor}}$. In this Table, we denote `cecilia`'s fixed model parameters with a dagger (†). We also exclude our results for the $\mathrm{RV}_{\mathrm{SDSS}}$ and $\mathrm{RV}_{\mathrm{Keck/ESI}}$ terms, as `cecilia` does not account for the stellar barycentric motion and gravitational redshift during its fitting procedure.

Parameter	SDSS J0231+2512	SDSS J0859+5732	SDSS J1109+1318	SDSS J1333+6364	SDSS J2311−0041
$T_{\mathrm{eff}}^{\dagger}$ [K]	12620±503	12677±722	15112±1688	14762±1340	12023±544
$\log g^{\dagger}$ [cgs]	$7.76^{+0.14}_{-0.13}$	$7.95^{+0.16}_{-0.15}$	$8.09^{+0.20}_{-0.19}$	$7.95^{+0.14}_{-0.13}$	$8.07^{+0.10}_{-0.09}$
$\log_{10}$ (H/He) [dex]	–	$-6.60^{+0.08}_{-0.09}(\pm0.04)(\pm0.20)$	$-4.33^{+0.03}_{-0.04}(\pm0.02)(\pm0.20)$	$-5.49^{+0.05}_{-0.05}(\pm0.03)(\pm0.20)$	$-6.26^{+0.07}_{-0.09}(\pm0.04)(\pm0.20)$
$\log_{10}$ (O/He) [dex]	–	$-5.35^{+0.03}_{-0.03}(\pm0.12)(\pm0.20)$	$-5.66^{+0.06}_{-0.06}(\pm0.13)(\pm0.20)$	$-5.85^{+0.08}_{-0.08}(\pm0.13)(\pm0.20)$	$-5.64^{+0.04}_{-0.04}(\pm0.13)(\pm0.20)$
$\log_{10}$ (Mg/He) [dex]	$-7.22^{+0.06}_{-0.07}(\pm0.08)(\pm0.20)$	$-6.24^{+0.02}_{-0.03}(\pm0.07)(\pm0.20)$	$-6.46^{+0.05}_{-0.05}(\pm0.07)(\pm0.20)$	$-6.54^{+0.07}_{-0.09}(\pm0.07)(\pm0.20)$	$-6.57^{+0.02}_{-0.02}(\pm0.07)(\pm0.20)$
$\log_{10}$ (Si/He) [dex]	–	$-6.48^{+0.03}_{-0.03}(\pm0.08)(\pm0.20)$	$-6.72^{+0.06}_{-0.06}(\pm0.09)(\pm0.20)$	–	$-6.78^{+0.07}_{-0.08}(\pm0.09)(\pm0.20)$
$\log_{10}$ (Ca/He) [dex]	$-8.58^{+0.08}_{-0.08}(\pm0.09)(\pm0.20)$	$-7.86^{+0.04}_{-0.04}(\pm0.09)(\pm0.20)$	$-7.34^{+0.07}_{-0.08}(\pm0.09)(\pm0.20)$	$-7.78^{+0.06}_{-0.06}(\pm0.09)(\pm0.20)$	$-7.94^{+0.03}_{-0.03}(\pm0.09)(\pm0.20)$
$\log_{10}$ (Fe/He) [dex]	–	$-6.66^{+0.04}_{-0.05}(\pm0.07)(\pm0.20)$	–	–	$-7.09^{+0.09}_{-0.11}(\pm0.08)(\pm0.20)$
$A_{\mathrm{jit,SDSS}}$ [−]	1.00±0.01	0.97±0.01	1.08±0.01	0.92±0.01	0.98±0.01
$A_{\mathrm{jit,Keck/ESI}}$ [−]	1.25±0.01	1.44±0.01	1.09±0.01	1.08±0.01	1.11±0.01
Detections [no.]	2 total (2 metals)	6 total (5 metals)	5 total (4 metals)	4 total (3 metals)	6 total (5 metals)

Appendix provide the corner plots of `cecilia`'s fits, demonstrating the convergence of our MCMCs and the lack of strong correlations between `cecilia`'s free model parameters.

From `cecilia`'s results, we can confidently identify a total of 2, 6, 5, 4, and 6 elements in the atmospheres of our targets (Table 3). In particular, we detect traces of Ca and Mg in SDSS J0231+2512; H, O, Mg, Si, Ca, and Fe in SDSS J0859+5732; H, O, Mg, Si, and Ca in SDSS J1109+1318; H, O, Mg, and Ca in SDSS J1333+6364; and H, O, Mg, Si, Ca, and Fe in SDSS J2311−0041. As discussed in Section 4.3, we consider these chemical species to be positive detections because their MCMC errors fall below our detectability cut-off (i.e. $\sigma_{\mathrm{stat,MCMC}} < \sigma_{\mathrm{detection}} = 0.10$ dex) and they all have at least one clear absorption line in the spectra. The most polluted WDs in our sample —defined as those with the largest number of detected elements— are SDSS J0859+5732 and SDSS J2311−0041, likely due to a combination of high accretion rates (see Table B1 and below) and relatively low effective temperatures. At lower T_{eff}, the photospheres of WDs tend to be less opaque, which would make it easier for `cecilia` to detect metal pollution.

Given that this paper represents the first application of `cecilia` to the study of polluted WDs with no well-measured abundances, we chose to validate our results by fitting the spectra of our targets with the classical method of Dufour et al. 2012 (DF12). This comparison yielded consistent abundances within 0.2 dex, which is roughly the scatter we would expect due to methodological differences in the two fitting approaches (see Section 4.1). More specifically, DF12 adopt a line-by-line χ^2 minimisation approach, which effectively gives the same weight to each visible absorption feature. This differs from `cecilia`'s approach, which fits the full spectra simultaneously,

assigning different weights to the lines based on their observational characteristics (e.g. line widths and strength, S/N).

5.2 Properties of Detected Metals

For all five WDs, we implemented the methodology described in Section 4.3 to estimate the mass ratio ($R_{\mathrm{ratio,Z}}$; Eq. 3), atmospheric mass (M_{Z}; Eq. 4), and accretion rate ($\dot{M}_{\mathrm{Z}}$; Eq. 5) of `cecilia`'s detected metals. Our results are summarised in Table B1, together with lower limits on the total metal mass ($\sum M_{\mathrm{Z}}$) and total accretion rate ($\sum \dot{M}_{\mathrm{Z}}$) of each system. For reference, we also include the logarithmic metal diffusion timescales ($\log_{10} \tau_{\mathrm{Z}}$) and the masses of the stellar convective layers (M_{cvz}) obtained from the MWDD.

Next, we used the equations presented in Section 4.4 to infer the metal abundances of the WD pollutants. To this end, we focused on the build-up and steady-state phases of accretion, which yielded results more consistent with the engulfment of rocky exoplanetary debris than the decaying phase. For our calculations, we normalised `cecilia`'s best-fit abundances by Mg, that is, we set $n_{\mathrm{B}} = n_{\mathrm{Mg}}$ and $\tau_{\mathrm{B}} = \tau_{\mathrm{Mg}}$ in Eq. 7-8. In Table B1, we provide our normalised abundance measurements ($n_{\mathrm{Z}}/n_{\mathrm{Mg}}$) for `cecilia`'s detected elements. From these results, we also derived their corresponding percent mass fractions (MF$_{\mathrm{Z}}$; Eq. 6). Figure 5 illustrates our estimated abundances and MFs, together with those of bulk Earth (Allègre et al. 2001), the CI chondrites (Alexander 2019a,b), and comet Halley (Jessberger et al. 1988).

Considering the full range of abundance ratios, the compositions of the polluting bodies accreting onto the five WDs are largely consistent with those of CI chondrites to within 1-2σ_{tot}. For SDSS J0859+5732 and SDSS J1109+1318, their Ca/Mg ratios appear slightly lower

Figure 4. MCMC corner plot for SDSS J0859+5732. The off-diagonal plots illustrate the two-dimensional marginalised posterior distributions of the free model parameters, while the histogram panels along the diagonal show their one-dimensional marginalised distributions together with their median value and 1σ confidence interval. This Figure does not show our results for `cecilia`'s undetected elements. It also excludes the RV shifts because `cecilia` does not correct for the barycentric motion and gravitational redshift of the white dwarf.

and higher, respectively, than those of these primitive meteorites, although they are still consistent within $2\sigma_{\rm tot}$. We also find that the O/Mg abundances of the four white dwarfs with oxygen detections (SDSS J0859+5732, SDSS J1109+1318, SDSS J1333+6364, SDSS J2311−0041) are enhanced in comparison to that of rocky, bulk Earth-like material. From these results, we calculated the oxygen excess by assessing how much oxygen binds with the major elements in the accreted material.

The concept of oxygen excess, first introduced by Klein et al. (2010), refers to the number of oxygen atoms that remain unaccounted for after calculating the amount of oxygen needed to form common rock-forming oxides (CaO, SiO_2, FeO, MgO). In this work, we investigated this phenomenon for SDSS J0859+5732 and SDSS J2311−0041, which exhibit the four main rock-forming elements in their spectra (O, Fe, Si, Mg). Assuming steady-state accretion and that iron is in the form of FeO, we followed the methodology of Rogers et al. (2024b) to quantify the oxygen excess and its sig-

nificance. In particular, we employed Monte Carlo techniques to sample plausible oyxgen excess values given our assumed total error on `cecilia`'s predicted abundances. Our analysis indicates that SDSS J0859+5732 and SDSS J2311−0041 have oxygen excesses of $0.64^{+0.13}_{-0.25}$ and $0.68^{+0.12}_{-0.23}$, respectively. Accounting for the asymmetry of the Monte Carlo sampled values, we calculated the fraction of sampled values showing a positive oxygen excess, obtaining 95.5% and 96.4% for SDSS J0859+5732 and SDSS J2311−0041, respectively. Comparing these probabilities to the 2σ significance threshold of Brouwers et al. 2022 (equivalent to 95.45% confidence),[26] we find

---

[26] As discussed in Brouwers et al. (2022), the use of 2σ confidence threshold is a conservative trade-off between minimising false positives (which would become abundant at the 1σ level), and the limitations imposed by errors in the stellar abundances and accretion rates (which would make the 3σ very restrictive). Our methodology follows the same conservative approach, but assumes a steady-state accretion model.

Figure 5. Compositional properties of the WD pollutants during build-up, steady-state, and decaying phase (first, second, and third columns, respectively). The top panels show the Mg-normalised linear abundance ratios of the accreted material relative to those of CI chondrites (red line; Alexander 2019a,b). The bottom panels present the percent metal mass fractions for systems with simultaneous detections of at least the four major rock-forming elements (O, Fe, Si, Mg), with Ca values scaled by a factor of 10 for clarity. In each plot, we only include `cecilia`'s detected elements (excluding Mg, our reference metal), with error bars reflecting a total assumed abundance error of $\sigma_{tot} = \sigma_{floor} = 0.20$ dex (see Table 3). For comparison, we also show the properties of bulk Earth (black '⊕' marker; Allègre et al. 2001), comet Halley (black '◇' marker; Jessberger et al. 1988), and CI chondrites (black '+' marker, bottom panels only). Heavy elements are sorted by increasing condensation temperature T_{cond} (Lodders 2003).

that both targets have statistically significant oxygen excesses. Our results are also similar to those reported by Brouwers et al. (2022) for GALEXJ2339 (0.61, using data from Klein et al. 2021), and WD1232+563 (0.57, using data from Xu et al. 2019), where H_2O is considered the most likely carrier of excess oxygen. It should be noted that the derived H abundances for SDSS J0859+5732 and SDSS J2311−0041 are too low to account for the oxygen excess, so it is likely additional oxygen carriers, such as carbon-bearing species, would be required to explain the observed stellar abundances.

Based on `cecilia`'s metal detections, our work suggests that the total metallic mass content of the WD pollutants ranges from 5.43×10^{21} g (for SDSS J0231+2512) to 1.92×10^{23} g (for SDSS J0859+5732) —in agreement with previous studies of metal pollution (e.g. Zuckerman et al. 2010; Farihi 2016; Xu et al. 2019).

At the lower end, our inferred mass is comparable to that of a small- or medium-sized asteroid in the Solar System. For example, 52 Europa has a mass of approximately 2.4×10^{22} g, slightly above our lower bound. At the upper end, our result is comparable to the accretion of a much larger body, similar to Saturn's icy moon Enceladus (1.08×10^{23} g; Jacobson 2022), or the large asteroids Vesta (2.5×10^{23} g; Russell 2012) and Pallas (2.05×10^{23} g; Carry 2012). Assuming a typical density of 2.2 g/cm^3 for CI chondrites (Britt & Consolmagno 2004), the total accreted masses derived in this work would result in a single pollutant with diameters between approximately 170-550 km. Lastly, in terms of total accretion rate, our results range between $10^7 - 10^9$ g/s, which is consistent with the observed distribution for He-atmosphere polluted WDs with similar effective temperatures (Rogers et al. 2024a).

6 DISCUSSION

6.1 Limitations and Insights of Compositional Analysis

Our geological inferences are inherently limited by the number of detected metals in the WDs and our conservative noise floor of 0.20 dex on the stellar abundances. Another limitation is that these systems are too faint to be observed with current UV facilities —which would typically allow for the detection of more elements, including volatiles such as C, N, or P—, so our conclusions are restricted to the optical spectra considered in this work. Recognising these observational caveats, the five WDs in our sample appear to be accreting rocky extrasolar material with a composition similar to those of CI chondrites. These results should be interpreted with caution, as the absence of major detected elements like Fe, Si, and Al introduces uncertainties in our calculations, especially when assuming chondritic values for some of the primary rock-forming elements. In particular, Fe and Si are critical for the formation of planetary mantles and cores, so their absence in some WDs limits our ability to make reliable assumptions about the geochemistry and internal structure of the accreted debris.

With respect to SDSS J0859+5732 and SDSS J2311−0041, both systems show enhanced oyxgen levels compared to bulk Earth, with significant oxygen excesses ($> 2\sigma$) pointing to the accretion of oxygen-rich exoplanetary material. An important caveat is that the assumed oxidation state of Fe can influence the magnitude of the inferred oxygen excess, as Fe could be in the form of metallic Fe, FeO, or Fe_2O_3. If the Fe were instead fully or partially in metallic form, the inferred oxygen excess would increase; conversely, if all Fe were assumed to be present as Fe_2O_3, the resulting oxygen excesses would be $0.60^{+0.14}_{-0.28}$ for SDSS J0859+5732 and $0.65^{+0.13}_{-0.25}$ and SDSS J2311−0041. In these cases, 94.1% and 95.6% of the Monte Carlo samples lie above zero, indicating a positive oxygen excess. When compared to the 2σ significance threshold of [Brouwers et al. 2022](corresponding to 95.45% confidence), the oxygen excess in SDSS J2311−0041 would remain statistically significant, while the result for SDSS J0859+5732 would fall just below the threshold, offering tentative evidence for an oxygen excess even under this conservative oxidation scenario. Another potential source of uncertainty in our oxygen excess calculations is the phase of accretion, with build-up making the excess more pronounced, and declining making it less significant. The detection of infrared flux excesses would offer independent evidence of ongoing accretion ([Bonsor et al. 2017](link)), but none of our targets have existing *WISE* and *Spitzer* photometry to support this hypothesis. Despite this limitation, the predicted abundances of the polluting bodies resemble those of rocky material from the Solar System assuming build-up or steady-state accretion, which would make it unlikely that these WDs are accreting in the declining phase. Furthermore, SDSS J0859+5732 and SDSS J2311−0041 retain their oxygen excess up to 4.8 and 5 times the Mg sinking timescale into the declining phase, respectively. At these times, the abundance patterns of their accreted material deviate significantly from those of known rocky Solar System bodies (see Figure 5 and Figure A9), which would further support our interpretation that these systems are not in the declining phase and that their oxygen excesses are real. Therefore, our results add two more WDs to the growing smaple of systems with notable oxygen excesses ([Farihi et al. 2011](link), [2013](link); [Raddi et al. 2015](link); [Xu et al. 2017](link); [Hoskin et al. 2020](link); [Klein et al. 2021](link); [Rogers et al. 2024a](link)), demonstrating that oxygen-rich material can survive during post-MS evolution and subsequently accrete onto WDs.

An additional insight from our geological analysis is the tentative evidence for Si depletion relative to bulk Earth and CI chondrites (see Figure 5). For example, our estimated steady-state Mg/Si

ratios are approximately 1.5, 1.7, and 1.7 for SDSS J0859+5732, SDSS J1109+1318, and SDSS J2311−0041, respectively (see Table B1). Although these values have relatively large uncertainties, they appear to exceed Earth's mantle ratio of about 1.3 ([Ringwood 1989](link)). Therefore, if real, they would point to the engulfment of polluting bodies with a mineralogy primarily composed of magnesium silicates —dominated by olivine with lesser amounts of pyroxene—, and relatively depleted in silicon compared to bulk Earth. Together, the high Mg/Si ratios and oxygen excesses observed in SDSS J0859+5732 and SDSS J2311−0041 would imply the accretion of oxygen-rich exoplanetary material with significant amounts of magnesium silicates, such as olivine.

Beyond individual systems, an elevated Mg/Si ratio could have important consequences for extrasolar geochemistry, shaping factors such as water storage capacity, magnetic fields, or tectonic activity. For instance, high Mg/Si ratios would favour the formation of forsterite-rich olivine, which is less dense and has a lower viscosity than pyroxene ([Mackwell 1991](link); [Hansen & Warren 2015](link)). This difference in viscosity would impact the thermal and dynamical evolution of the mantle; in particular, lower-viscosity olivine would lead to more vigorous mantle convection and faster cooling of the core, which could potentially shorten the lifetimes of magnetic fields and plate tectonics ([Spaargaren et al. 2020](link); [Spaargaren et al. 2023](link)) — two critical processes for sustaining long-term planetary habitability ([Stern & Gerya 2024](link)).

6.2 cecilia's Current Performance and Future Work

In this paper, we have presented several upgrades to cecilia, including the implementation of a more complex MCMC log-likelihood function to allow for joint spectroscopic fits with a more careful treatment of noise and spectral lines (see Section 4.2). We have also reconfigured cecilia's MCMC to enable the calculation of upper abundance limits for those elements with unclear or unobservable absorption lines. Although these upgrades have had a positive impact on cecilia's performance, there are still many opportunities to improve our code. For instance, as discussed in BA24, cecilia could be retrained with more WD models featuring synthetic photometric observations, lower and higher abundance ranges, additional heavy elements (e.g. Al, Na), and new regions of the spectrum, such as the UV.

Building upon the recommendations of BA24, we propose several new directions for future work to further enhance cecilia's capabilities. Our suggestions can be divided into two broad categories: data processing, and Bayesian inference techniques. In relation to the former, we would like to improve our methodology for correcting telluric contamination by Earth's atmosphere, particularly redward of about 6,000 Å. In this paper, we cautiously removed all the most important telluric bands from our *Keck/ESI* spectra, even if some of these bands contained potential metal absorption lines (see Section 3.3). To address this limitation, we aim to develop a more selective and targeted approach to removing telluric lines in ground-based optical observations. For example, we could correct an observed WD spectrum using a telluric absorption model generated by packages such as the Fortran line-by-line radiative transfer code LBLRTM ([Clough et al. 2005](link))[27] or its Python telfit implementation[28] ([Gullikson & Dodson-Robinson 2013](link); [Gullikson et al. 2014](link)). The latter uses a non-linear least-squares Levenverg-Marquardt algorithm to

---

[27] http://rtweb.aer.com/lblrtm.html
[28] https://telfit.readthedocs.io/en/latest/

model telluric effects, given a list of molecular line strengths from the HITRAN database and a nighttime atmosphere model containing the pressure, temperature, and abundances of different molecules as a function of height. We note that `telfit` can also employ custom atmosphere profiles at the specific location of the user's observatory.

In addition to developing better data processing techniques, we have also identified multiple opportunities to improve `cecilia`'s optimisation procedure, which we summarise below.

(i) *Use of Gaussian Processes*: Given the inaccuracy of state-of-the-art atmosphere models to account for the problem of neutral He line broadening from atom collisions, we would like to explore the use of Gaussian Process (GP) regression to improve the quality of `cecilia`'s results, especially around poorly-modelled absorption lines. Broadly speaking, GPs are powerful mathematical tools to perform joint non-parametric[29] fits to astronomical observations with both uncorrelated (i.e. *white*) and correlated (i.e. *red*) sources of noise (e.g. instrumental effects or stellar variability). Mathematically, GPs are defined as generalisations of multivariate Gaussian distributions with a covariance matrix (also known as *kernel* function) that encapsulates the underlying stochastic (i.e. random) correlation structures between adjacent datapoints (Rasmussen & Williams 2006). In the context of `cecilia`, we could implement a GP noise model to account for imperfections in the atmosphere models of He-rich polluted WDs. This approach would downweight problematic spectral regions by inflating the systematic errors of poorly fitted absorption lines, hence penalising atmospheric solutions with under- or overestimated abundances. On a practical level, the use of GPs would involve balancing `cecilia`'s model flexibility, uncertainty estimates, and computational speed (Czekala et al. 2015, C15). On the one hand, GPs may be more efficient at capturing underestimated systematics in the WD atmosphere models, potentially leading to more realistic and conservative results. On the other, they may inflate the errors on the retrieved model parameters, while also slowing down `cecilia`'s optimisation routine. These trade-offs may nonetheless be reasonable if they ultimately improve the reliability and accuracy of our code. In the future, we can explore the integration of GPs with publicly available software packages, such as `tinyGP` (Foreman-Mackey et al. 2024) or `starfish` (C15).[30]

(ii) *Exploration of new sampling mechanisms*: The main optimisation method employed by `cecilia` is the differential evolution MCMC algorithm of Ter Braak (2006). As we seek to improve the efficiency of our pipeline, we could also experiment with different inference techniques, such as nested sampling (Skilling 2004; Skilling 2006). Unlike MCMCs, which are designed to sample the posterior distribution directly from the likelihood function and the prior density, nested sampling revolves around estimating the Bayesian evidence of a model. This parameter is very difficult to determine, so nested sampling addresses this problem in a dynamical and iterative way. More specifically, the algorithm draws an ensemble of random "live points" from the prior, removes the point with the lowest likelihood, generates new live points with higher likelihoods, and repeats this process until the Bayesian evidence of a model satisfies a certain threshold. Therefore, unlike MCMCs, which are "memoryless" systems in which the behaviour of two walkers only depends on their previous state, nested sampling methods systematically explore a large volume of the parameter space and gradually compress it based on regions of higher probability. With this strategy, they are more robust to poor initial guesses as well as to multi-modal distributions between different parameters (Ashton et al. 2022). A potential nested sampling implementation is the open-source package `dynesty` (Speagle 2020; Koposov et al. 2023).[31]

(iii) *Bayesian model comparison*: Central to our understanding of metal pollution is our degree of belief in spectroscopic fits of polluted WD spectra, especially when obtained with ML codes that may not often be easily explainable. In the future, we could improve `cecilia` by integrating a Bayesian model comparison framework to assess the relative probability of different ML predictions in the absence or presence of a heavy element. This can be achieved by calculating the so-called "Bayes factor," which is defined as the ratio

$$R \equiv \frac{p_{M_1} \cdot \pi_{M_1}}{p_{M_2} \cdot \pi_{M_2}}, \tag{10}$$

where M_1 and M_2 are two competing models, p_{M_1} and p_{M_2} are their posterior probability distributions given the observed data, and π_{M_1} and π_{M_2} are our *a priori* beliefs on each model. From this expression, R≫1 and B≈0 would indicate a strong preference for M_1 and M_2, respectively. However, more complex interpretations are also possible, such as those of Jeffreys (1998) and Kass & Raftery (1995). In general, Eq. 10 is difficult to solve, so there are several approximate forms of Bayesian model comparison. For example, the so-called Bayesian Information Criterion (BIC) evaluates the accuracy and complexity of different models based on the maximum value of their likelihood function ($\hat{\mathcal{L}}$), the number of points in the observations (n), and their number of free model parameters (k), i.e., BIC $= -2\ln(\hat{\mathcal{L}}) + k\ln(n)$.

(iv) *Improved Treatment of Resolving Power or Spectral Resolution*: Our analysis of *SDSS* observations assumes a constant resolving power of R=2,000 between about 3,800Å to 9,000Å, even though *R* changes significantly by about 67% (from 1,500 to 2,500) across this wide wavelength range. Our simplification does not fully capture the intrinsic properties of the spectra, so future iterations of `cecilia` could adopt a more adaptive approach, either by convolving the data to a constant linear resolution in wavelength space, or by directly supporting a variable resolving power or resolution to better preserve the native characteristics of the observations.

(v) *Automated detection of chemical elements*: In this paper, we consider an element to be detected if (i) its `cecilia` predicted abundance had a statistical error lower than or equal to a conservative detectability threshold of $\sigma_{\text{detection}} = 0.10$ dex, and (ii) if it exhibited at least a clear absorption line in the spectrum. This validation process could be automated in order to streamline all the criteria associated to a robust detection.

7 CONCLUSIONS

In this paper, we have used the ML-pipeline `cecilia` to constrain the physical and chemical properties of five metal polluted He-

---

[29] Parametric models have a fixed number of unknown parameters. This differs from non-parametric models, which can have an arbitrarily large number of dimensions.

[30] These packages can be accessed via https://github.com/dfm/tinygp and https://starfish.readthedocs.io/en/latest/index.html, respectively.

[31] https://dynesty.readthedocs.io/en/latest/dynamic.html

atmosphere WDs. We started by performed a joint fit to their *SDSS* (R=2,000) and *Keck/ESI* (R=4,500) spectra with an MCMC model consisting of 17 parameters (T_{eff}, log g, 11 elemental abundances, and a jitter term and RV shift per spectrum). We then estimated the geological composition of the WD pollutants from `cecilia`'s predicted stellar abundances. The main computational and scientific conclusions of our work are summarised below.

(i) *Upgrades to* `cecilia` *(Section 4.2)*: We have improved `cecilia`'s fitting procedure in two ways. First, we have re-parametrised its MCMC to sample the posterior distribution of the model parameters in linear space. This change will eventually allow us to calculate robust upper abundance limits for those elements with undetected spectral lines. Second, we have modified our log-likelihood function (Eq. 1) to allow for a joint fit of multiple spectra, regardless of their observational characteristics. We have also incorporated a jitter term per spectrum to model unaccounted sources of noise in the observations or in `cecilia`'s training atmosphere models.

(ii) *Atmospheric Analysis (Section 4.3)*: We have identified a total of 2, 6, 5, 4, and 6 chemical elements in the atmospheres of SDSS J0231+2512 (Mg, Ca), SDSS J0859+5732 (H, O, Mg, Si, Ca, Fe), SDSS J1109+1318 (H, O, Mg, Si, Ca), SDSS J1333+6364 (H, O, Mg, Ca), and SDSS J2311−0041 (H, O, Mg, Si, Ca, Fe), respectively. In terms of detected metals, SDSS J0859+5732 and SDSS J2311−0041 are the most heavily polluted in our sample. For all five WDs, `cecilia` constrains the abundances of the detected elements with greater precision than our assumed noise floor of ≤0.20 dex —a performance comparable to that of classical WD fitting methods. We emphasise, however, that `cecilia`'s predictive power is fundamentally restricted by the quality of its training models and by its ML interpolation errors (see BA24 for a detailed discussion of systematics using synthetic data). In other words, although `cecilia` may yield lower statistical uncertainties than conventional techniques, its overall performance cannot exceed that of classical methods because it is inherently limited by the accuracy of its underlying WD atmosphere models.

(iii) *Geological Analysis (Sections 4.4-5.2)*: Our results indicate that the five WDs engulfed rocky extrasolar material with a bulk composition largely consistent with those of CI chondrites (within 1-2σ_{tot}), with the main rock-forming elements (Mg, Fe, Si, and O) accounting for most of their interior composition by mass. Among these systems, SDSS J0859+5732 and SDSS J2311−0041 have statistically significant ($> 2\sigma$) oxygen excesses, which could indicate the accretion of oxygen-rich extrasolar material. These findings are not conclusive given the limited number of detected metals in each WD and the uncertainties in `cecilia`'s stellar abundances. However, even within this margin of error, our analysis is sensitive enough to identify compositional deviations in the WD pollutants. In the coming years, as next-generation optical/UV telescopes become operational, it may be possible to disentangle their full geochemical properties and place them in the broader context of extrasolar compositions.

To conclude, we have demonstrated that `cecilia` can constrain the elemental abundances of He-rich polluted WDs with no *a priori* knowledge of their atmospheric composition and with minimal human supervision. As we venture into the era of massively multiplexed observational surveys (e.g. DESI, SDSS-V, WEAVE, 4MOST), `cecilia` can be used to rapidly analyse large volumes of data within a practical amount of time, therefore offering a solution to the human-in-the-loop problem of conventional WD characterisation methods. In doing so, `cecilia` aims to lay the foundations for a methodological shift towards population-wide studies of metal pollution —a "Big Data" revolution with the potential to offer new insights into the properties of extrasolar worlds.

DATA AVAILABILITY

The *SDSS* and *Keck/ESI* spectra of our targets can be downloaded from the *SDSS* DR18 and the *Keck* online databases. The spectra of the standard stars described in Section 3 can be obtained from the STScI archive.

REFERENCES

Abadi M., et al., 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, https://www.tensorflow.org/

Abadi M., et al., 2016, in Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation. OSDI'16. USENIX Association, USA, pp 265–283

Alexander C. M. O., 2019a, Geochimica Cosmochimica Acta, 254, 246

Alexander C. M. O., 2019b, Geochimica Cosmochimica Acta, 254, 277

Allègre C., Manhès G., Lewin É., 2001, Earth and Planetary Science Letters, 185, 49

Almeida A., et al., 2023, ApJS, 267, 44

Ashton G., et al., 2022, Nature Reviews Methods Primers, 2, 39

Badenas-Agusti M., Viaña J., Vanderburg A., Blouin S., Dufour P., Xu S., Sha L., 2024, MNRAS, 529, 1688

Becklin E. E., Farihi J., Jura M., Song I., Weinberger A. J., Zuckerman B., 2005, The Astrophysical Journal, 632, L119

Bédard A., Bergeron P., Brassard P., Fontaine G., 2020, ApJ, 901, 93

Bédard A., Bergeron P., Brassard P., 2022, ApJ, 930, 8

Bergeron P., et al., 2011, ApJ, 737, 28

Bhattacharjee S., et al., 2025, arXiv e-prints, p. arXiv:2502.05502

Bigelow B. C., Nelson J. E., 1998, in D'Odorico S., ed., Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 3355, Optical Astronomical Instrumentation. pp 164–174, doi:10.1117/12.316829

Blouin S., Dufour P., Allard N. F., 2018a, ApJ, 863, 184

Blouin S., Dufour P., Allard N. F., Kilic M., 2018b, ApJ, 867, 161

Bohlin R. C., Gordon K. D., Tremblay P. E., 2014, PASP, 126, 711

Bonsor A., Farihi J., Wyatt M. C., van Lieshout R., 2017, MNRAS, 468, 154

Britt D. T., Consolmagno G. J., 2004, in Mackwell S., Stansbery E., eds, Lunar and Planetary Science Conference. Lunar and Planetary Science Conference. p. 2108

Brouwers M. G., Buchan A. M., Bonsor A., Malamud U., Lynch E., Rogers L., Koester D., 2022, Monthly Notices of the Royal Astronomical Society, 519, 2663

Buton C., et al., 2013, A&A, 549, A8

Camisassa M. E., Althaus L. G., Rohrmann R. D., García-Berro E., Torres S., Córsico A. H., Wachlin F. C., 2017, ApJ, 839, 11

Campbell B., Walker G. A. H., Yang S., 1988, ApJ, 331, 902

Carry B., 2012, Planetary and Space Science, 73, 98

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

Clough S. A., Shephard M. W., Mlawer E. J., Delamere J. S., Iacono M. J., Cady-Pereira K., Boukabara S., Brown P. D., 2005, J. Quant. Spectrosc. Radiative Transfer, 91, 233

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

Cukanovaite E., Tremblay P.-E., Bergeron P., Freytag B., Ludwig H.-G., Steffen M., 2020, Monthly Notices of the Royal Astronomical Society, 501, 5274

Czekala I., Andrews S. M., Mandel K. S., Hogg D. W., Green G. M., 2015, ApJ, 812, 128

DESI Collaboration et al., 2016a, arXiv e-prints, p. arXiv:1611.00036

DESI Collaboration et al., 2016b, arXiv e-prints, p. arXiv: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), doi:10.1117/12.2055132

Dorn C., Khan A., Heng K., Connolly J. A. D., Alibert Y., Benz W., Tackley P., 2015, A&A, 577, A83

Doyle A. E., Young E. D., Klein B., Zuckerman B., Schlichting H. E., 2019, Science, 366, 356

Doyle A. E., et al., 2023, ApJ, 950, 93

Dressing C. D., et al., 2015, The Astrophysical Journal, 800, 135

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

Dupuis J., Fontaine G., Pelletier C., Wesemael F., 1992, ApJS, 82, 505

Dupuis J., Fontaine G., Pelletier C., Wesemael F., 1993a, ApJS, 84, 73

Dupuis J., Fontaine G., Wesemael F., 1993b, ApJS, 87, 345

Farihi J., 2016, New Astron. Rev., 71, 9

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

Fontaine G., Michaud G., 1979, ApJ, 231, 826

Ford E. B., 2006, ApJ, 642, 505

Foreman-Mackey D., 2016, The Journal of Open Source Software, 1, 24

Foreman-Mackey D., et al., 2024, dfm/tinygp: The tiniest of Gaussian Process libraries, doi:10.5281/zenodo.10463641, https://doi.org/10.5281/zenodo.10463641

Fortney J. J., Marley M. S., Barnes J. W., 2007, The Astrophysical Journal, 659, 1661

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

Gänsicke B. T., Koester D., Farihi J., et al. 2012, MNRAS, 424, 333

Gelman A., Rubin D. B., 1992, Statistical Science, 7, 457

Gelman A., Carlin J. B., Stern H. S., Rubin D. B., 2004, Bayesian Data Analysis. Chapman & Hall/CRC, Boca Raton, FL

Genest-Beaulieu C., Bergeron P., 2019a, ApJ, 871, 169

Genest-Beaulieu C., Bergeron P., 2019b, ApJ, 882, 106

Gentile Fusillo N. P., Gänsicke B. T., Farihi J., Koester D., Schreiber M. R., Pala A. F., 2017, MNRAS, 468, 971

Gentile Fusillo N. P., et al., 2019, MNRAS, 482, 4570

Gullikson K., Dodson-Robinson S., 2013, AJ, 145, 3

Gullikson K., Dodson-Robinson S., Kraus A., 2014, AJ, 148, 53

Gunn J. E., et al., 2006, AJ, 131, 2332

Hansen L. N., Warren J. M., 2015, Journal of Geophysical Research: Solid Earth, 120, 2717

Harrison J. H. D., Bonsor A., Madhusudhan N., 2018, Monthly Notices of the Royal Astronomical Society, 479, 3814

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

Jacobson R. A., 2022, The Astronomical Journal, 164, 199

Jeffreys H., 1998, Theory of Probability. International series of monographs on physics, Clarendon Press, https://books.google.com/books?id=_PuRmAEACAAJ

Jessberger E. K., Christoforidis A., Kissel J., 1988, Nature, 332, 691

Jontof-Hutter D., 2019, Annual Review of Earth and Planetary Sciences, 47, 141

Jura M., 2006, ApJ, 653, 613

Jura M., 2008, AJ, 135, 1785

Jura M., Xu S., 2012a, AJ, 143, 6

Jura M., Xu S., 2012b, The Astronomical Journal, 145, 30

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

Kass R. E., Raftery A. E., 1995, Journal of the American Statistical Association, 90, 773

Klein B., Jura M., Koester D., Zuckerman B., Melis C., 2010, The Astrophysical Journal, 709, 950

Klein B. L., Doyle A. E., Zuckerman B., Dufour P., Blouin S., Melis C., Weinberger A. J., Young E. D., 2021, ApJ, 914, 61

Koester D., 2009, A&A, 498, 517

Koester D., 2010, Mem. Soc. Astron. Italiana, 81, 921

Koester D., Kepler S. O., 2015, A&A, 583, A86

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

Koposov S., et al., 2023, joshspeagle/dynesty: v2.1.3, doi:10.5281/zenodo.8408702, https://doi.org/10.5281/zenodo.8408702

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.

Kupka F., Piskunov N., Ryabchikova T. A., Stempels H. C., Weiss W. W., 1999, A&AS, 138, 119

Lodders K., 2003, ApJ, 591, 1220

Luger R., Lustig-Yaeger J., Agol E., 2017, ApJ, 851, 94

Mackwell S. J., 1991, Geophysical Research Letters, 18, 2027

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), doi:10.48550/arXiv.0902.2850

Martin D. C., et al., 2005, ApJ, 619, L1

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

Oliphant T. E., 2015, Guide to NumPy, 2nd edn. CreateSpace Independent Publishing Platform, North Charleston, SC, USA

Paquette C., Pelletier C., Fontaine G., Michaud G., 1986, ApJS, 61, 177

Pelletier C., Fontaine G., Wesemael F., Michaud G., Wegner G., 1986, ApJ, 307, 242

Penoyre Z., Belokurov V., Wyn Evans N., Everall A., Koposov S. E., 2020, Monthly Notices of the Royal Astronomical Society, 495, 321

Piskunov N. E., Kupka F., Ryabchikova T. A., Weiss W. W., Jeffery C. S., 1995, A&AS, 112, 525

Price-Whelan A. M., et al., 2018, AJ, 156, 123

Putirka K. D., Xu S., 2021, Nature Communications, 12, 6168

Raddi R., Gänsicke B. T., Koester D., Farihi J., Hermes J. J., Scaringi S., Breedt E., Girven J., 2015, MNRAS, 450, 2083

Rasmussen C. E., Williams C. K. I., 2006, Gaussian Processes for Machine Learning

Rebassa-Mansergas A., Nebot Gomez-Moran A., Schreiber M. R., Girven J., Gansicke B. T., 2011, Monthly Notices of the Royal Astronomical Society, 413, 1121

Ringwood A., 1989, Earth and Planetary Science Letters, 95, 1

Rogers L. A., Seager S., 2010, ApJ, 712, 974

Rogers L. K., et al., 2024a, MNRAS, 527, 6038

Rogers L. K., et al., 2024b, MNRAS, 532, 3866

Russell C. T. e. a., 2012, Science, 336, 684

Ryabchikova T., Piskunov N., Kurucz R. L., Stempels H. C., Heiter U., Pakhomov Y., Barklem P. S., 2015, Phys. Scr., 90, 054005

Saumon 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

Sheinis A. I., Bolte M., Epps H. W., Kibrick R. I., Miller J. S., Radovan M. V., Bigelow B. C., Sutin B. M., 2002, PASP, 114, 851

Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V., eds, American Institute of Physics Conference Series Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. pp 395–405, doi:10.1063/1.1835238

Skilling J., 2006, Bayesian Analysis, 1, 833

Spaargaren R. J., Ballmer M. D., Bower D. J., Dorn C., Tackley P. J., 2020, A&A, 643, A44

Spaargaren R. J., Wang H. S., Mojzsis S. J., Ballmer M. D., Tackley P. J., 2023, The Astrophysical Journal, 948, 53

Speagle J. S., 2020, Monthly Notices of the Royal Astronomical Society, 493, 3132

Stern R. J., Gerya T. V., 2024, Scientific Reports, 14, 8552

Swan A., Farihi J., Koester D., Hollands M., Parsons S., Cauley P. W., Redfield S., Gänsicke B. T., 2019, Monthly Notices of the Royal Astronomical Society, 490, 202

Ter Braak C. J. F., 2006, Statistics and Computing, 16, 239

Trierweiler I. L., Doyle A. E., Melis C., Walsh K. J., Young E. D., 2022, ApJ, 936, 30

Turner S. G. D., Wyatt M. C., 2019, Monthly Notices of the Royal Astronomical Society, 491, 4672

Vanderburg A., 2021, avanderburg/edmcmc: v1.0.0, doi:10.5281/zenodo.5599854, https://doi.org/10.5281/zenodo.5599854

Veras D., 2021, Planetary Systems Around White Dwarfs, doi:10.1093/acrefore/9780190647926.013.238, https://doi.org/10.1093/acrefore/9780190647926.013.238

Vincent O., Bergeron P., Dufour P., 2023, Monthly Notices of the Royal Astronomical Society, 521, 760

Virtanen P., et al., 2020, Nature Methods,

Weidemann V., 1960, ApJ, 131, 638

Weidemann V., Koester D., 1983, A&A, 121, 77

Williams J. T., Gänsicke B. T., Swan A., O'Brien M. W., Izquierdo P., Cutolo A. M., Cunningham T., 2024, A&A, 691, A352

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

Wyatt M. C., Farihi J., Pringle J. E., Bonsor A., 2014, Monthly Notices of the Royal Astronomical Society, 439, 3371

Xu S., Jura M., Klein B., Koester D., Zuckerman B., 2013, ApJ, 766, 132

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

Xu S., Rogers L. K., Blouin S., 2024, Reviews in Mineralogy and Geochemistry, 90, 171

York D. G., et al., 2000, AJ, 120, 1579

Zeng L., Sasselov D., 2013, Publications of the Astronomical Society of the Pacific, 125, 227

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

Zuckerman B., Koester D., Dufour P., Melis C., Klein B., Jura M., 2011, ApJ, 739, 101

van Maanen A., 1917, Publications of the Astronomical Society of the Pacific, 29, 258

APPENDIX A: FIGURES

In Figures A1-A4, we present `cecilia`'s RV-shifted MCMC solutions (in green) for the median-normalised *Keck/ESI* spectra of

SDSS J0231+2512, SDSS J1109+1318, SDSS J1333+6364, and SDSS J2311−0041 (in blue; air wavelengths). These figures also illustrate `cecilia`'s predictions when modifying the best-fit abundances of the detected elements by $\pm 1\sigma_{tot}$ (in red and orange). For reference, we use green labels to denote `cecilia`'s detections, i.e. those elements with $\sigma_{stat,MCMC} \leq \sigma_{detection} = 0.10$ dex and at least one visible absorption line in the spectrum. In Figures A5-A8, we show the corner plots of `cecilia`'s MCMC solutions, excluding the undetected or tentative elements, as well as the the RV shifts of each spectrum. Lastly, Figure A9 shows the continuous decay evolution of metal mass fractions for SDSS J0859+5732 and SDSS J2311−0041, focusing on `cecilia`'s detected metals.

APPENDIX B: TABLES

In Table B1, we provide the main compositional properties of our targets during the build-up and steady-state phases of accretion. In Table B2, we compare `cecilia`'s best-fit elemental abundances to those obtained with the classical, line-by-line fitting method of DF12 (see Section 5.1 for a brief summary of the DF12 approach).

This paper has been typeset from a TEX/LATEX file prepared by the author.

Figure A1. `cecilia`'s best-fit MCMC model (in green) for the median-normalised *Keck/ESI* spectrum of SDSS J0231+2512, with its corresponding $\pm 1\sigma_{\rm tot}$ models (in red and orange).

Figure A2. `cecilia`'s best-fit MCMC model (in green) for the median-normalised *Keck/ESI* spectrum of SDSS J1109+1318, with its corresponding $\pm 1\sigma_{\rm tot}$ models (in red and orange).

Figure A3. `cecilia`'s best-fit MCMC model (in green) for the median-normalised *Keck/ESI* spectrum of SDSS J1333+6364, with its corresponding $\pm 1\sigma_{tot}$ models (in red and orange).

Figure A4. cecilia's best-fit MCMC model (in green) for the median-normalised *Keck/ESI* spectrum of SDSS J2311−0041, with its corresponding ±1σ_tot models (in red and orange).

SDSS J0231 + 2512

Figure A5. MCMC corner plot for SDSS J0231+2512.

SDSS J1109 + 1318

Figure A6. MCMC corner plot for SDSS J1109+1318.

SDSS J1333 + 6364

Figure A7. MCMC corner plot for SDSS J1333+6364.

SDSS J2311 − 0041

Figure A8. MCMC corner plot for SDSS J2311−0041.

Table B1. Compositional properties of the five polluted WDs and their accreted material during build-up and steady-state (see Sections 4.3-4.4). The metal mass fractions are only reported for those systems which exhibit the four main-rock forming elements at the same time (Mg, Si, Fe, and O). In this Table, we only show `cecilia`'s detected metals, sorted by increasing condensation temperature T_{cond} (Lodders 2003).

| Property	SDSS J0231+2512	SDSS J0859+5732	SDSS J1109+1318	SDSS J1333+6364	SDSS J2311−0041
			Oxygen (T_{cond}=182 K)		
$\log_{10}(\tau_O)$ [yrs]		6.35	5.74	6.07	6.16
M_O [10^{20} g]		$1303.56^{+770.35}_{-485.76}$	$156.85^{+96.50}_{-60.10}$	$172.14^{+110.23}_{-67.28}$	$461.60^{+275.90}_{-173.57}$
$\dot{M}_O$ [10^8 g/s; Steady]		$18.38^{+10.86}_{-6.85}$	$9.15^{+5.63}_{-3.51}$	$4.66^{+2.98}_{-1.82}$	$10.15^{+6.07}_{-3.82}$
n_O/n_{Mg} [Observed; Build-up]		$8.15^{+4.86}_{-3.05}$	$6.98^{+4.43}_{-2.71}$	$5.06^{+3.53}_{-2.07}$	$8.76^{+5.27}_{-3.31}$
n_O/n_{Mg} [Steady]	Not Detected	$8.27^{+4.92}_{-3.10}$	$7.02^{+4.46}_{-2.73}$	$5.06^{+3.53}_{-2.07}$	$8.99^{+5.41}_{-3.40}$
$MR_{O,Mg}$ [Observed; Build-up]		$5.36^{+4.97}_{-2.59}$	$4.58^{+4.41}_{-2.24}$	$3.33^{+3.37}_{-1.68}$	$5.75^{+5.39}_{-2.78}$
$MR_{O,Mg}$ [Steady]		$5.43^{+5.04}_{-2.62}$	$4.61^{+4.44}_{-2.26}$	$3.33^{+3.37}_{-1.68}$	$5.90^{+5.53}_{-2.85}$
MF_O [Observed; Build-up]		$52.42^{+12.77}_{-13.28}$	–	–	$62.47^{+11.33}_{-13.05}$
MF_O [Steady]		$45.65^{+13.28}_{-12.88}$	–	–	$56.72^{+12.12}_{-13.36}$
			Iron (T_{cond}=1357 K)		
$\log_{10}(\tau_{Fe})$ [yrs]		6.16			5.96
M_{Fe} [10^{20} g]		$205.07^{+123.90}_{-76.84}$			$58.19^{+38.74}_{-23.56}$
$\dot{M}_{Fe}$ [10^8 g/s; Steady]		$4.51^{+2.72}_{-1.69}$			$2.01^{+1.34}_{-0.81}$
n_{Fe}/n_{Mg} [Observed; Build-up]		$0.37^{+0.22}_{-0.14}$			$0.32^{+0.21}_{-0.13}$
n_{Fe}/n_{Mg} [Steady]	Not Detected	$0.58^{+0.35}_{-0.22}$	Not Detected	Not Detected	$0.51^{+0.34}_{-0.21}$
$MR_{Fe,Mg}$ [Observed; Build-up]		$0.84^{+0.78}_{-0.41}$			$0.72^{+0.72}_{-0.36}$
$MR_{Fe,Mg}$ [Steady]		$1.33^{+1.24}_{-0.64}$			$1.16^{+1.16}_{-0.59}$
MF_{Fe} [Observed; Build-up]		$8.07^{+5.22}_{-3.28}$			$7.75^{+5.57}_{-3.43}$
MF_{Fe} [Steady]		$10.98^{+6.69}_{-4.34}$			$11.07^{+7.36}_{-4.77}$
			Magnesium (T_{cond}=1397 K)		
$\log_{10}(\tau_{Mg})$ [yrs]	6.72	6.36	5.74	6.07	6.17
M_{Mg} [10^{20} g]	$50.19^{+31.39}_{-19.37}$	$243.31^{+143.83}_{-90.05}$	$34.23^{+20.68}_{-12.91}$	$51.70^{+32.94}_{-20.14}$	$80.17^{+47.46}_{-29.67}$
$\dot{M}_{Mg}$ [10^8 g/s; Steady]	$0.30^{+0.19}_{-0.12}$	$3.38^{+2.00}_{-1.25}$	$1.98^{+1.20}_{-0.75}$	$1.40^{+0.89}_{-0.54}$	$1.72^{+1.02}_{-0.64}$
n_{Mg}/n_{Mg} [Observed; Build-up]	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$
n_{Mg}/n_{Mg} [Steady]	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$	$1.00^{+0.59}_{-0.37}$
$MR_{Mg,Mg}$ [Observed; Build-up]	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$
$MR_{Mg,Mg}$ [Steady]	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$	$1.00^{+0.00}_{-0.00}$
MF_{Mg} [Observed; Build-up]	–	$9.60^{+5.98}_{-3.87}$	–	–	$10.74^{+6.79}_{-4.43}$
MF_{Mg} [Steady]	–	$8.22^{+5.17}_{-3.29}$	–	–	$9.45^{+5.96}_{-3.85}$
			Silicon (T_{cond}=1529 K)		
$\log_{10}(\tau_{Si})$ [yrs]		6.36	5.72		6.18
M_{Si} [10^{20} g]		$162.48^{+96.50}_{-60.31}$	$19.33^{+11.95}_{-7.35}$		$47.44^{+29.85}_{-18.27}$
$\dot{M}_{Si}$ [10^8 g/s; Steady]		$2.24^{+1.33}_{-0.83}$	$1.17^{+0.72}_{-0.44}$		$1.00^{+0.63}_{-0.39}$
n_{Si}/n_{Mg} [Observed; Build-up]		$0.58^{+0.35}_{-0.21}$	$0.49^{+0.31}_{-0.19}$		$0.51^{+0.32}_{-0.20}$
n_{Si}/n_{Mg} [Steady]	Not Detected	$0.57^{+0.34}_{-0.21}$	$0.51^{+0.32}_{-0.20}$	Not Detected	$0.50^{+0.32}_{-0.20}$
$MR_{Si,Mg}$ [Observed; Build-up]		$0.67^{+0.61}_{-0.32}$	$0.57^{+0.54}_{-0.28}$		$0.59^{+0.56}_{-0.29}$
$MR_{Si,Mg}$ [Steady]		$0.66^{+0.61}_{-0.32}$	$0.59^{+0.56}_{-0.29}$		$0.58^{+0.55}_{-0.29}$
MF_{Si} [Observed; Build-up]		$6.39^{+4.14}_{-2.62}$	–		$6.32^{+4.44}_{-2.73}$
MF_{Si} [Steady]		$5.42^{+3.53}_{-2.21}$	–		$5.48^{+3.82}_{-2.34}$
			Calcium (T_{cond}=1659 K)		
$\log_{10}(\tau_{Ca})$ [yrs]	6.55	6.18	5.55	5.89	5.98
M_{Ca} [10^{20} g]	$4.15^{+2.65}_{-1.62}$	$9.44^{+5.71}_{-3.51}$	$6.60^{+4.23}_{-2.55}$	$6.84^{+4.26}_{-2.60}$	$5.62^{+3.36}_{-2.09}$
$\dot{M}_{Ca}$ [10^8 g/s; Steady]	$0.04^{+0.02}_{-0.01}$	$0.20^{+0.12}_{-0.07}$	$0.59^{+0.38}_{-0.23}$	$0.28^{+0.17}_{-0.11}$	$0.19^{+0.11}_{-0.07}$
n_{Ca}/n_{Mg} [Observed; Build-up]	$0.05^{+0.03}_{-0.02}$	$0.02^{+0.01}_{-0.01}$	$0.12^{+0.08}_{-0.05}$	$0.08^{+0.05}_{-0.03}$	$0.04^{+0.03}_{-0.02}$
n_{Ca}/n_{Mg} [Steady]	$0.07^{+0.05}_{-0.03}$	$0.04^{+0.02}_{-0.01}$	$0.18^{+0.12}_{-0.07}$	$0.12^{+0.08}_{-0.05}$	$0.07^{+0.04}_{-0.02}$
$MR_{Ca,Mg}$ [Observed; Build-up]	$0.08^{+0.08}_{-0.04}$	$0.04^{+0.04}_{-0.02}$	$0.19^{+0.19}_{-0.09}$	$0.13^{+0.13}_{-0.07}$	$0.07^{+0.07}_{-0.03}$
$MR_{Ca,Mg}$ [Steady]	$0.12^{+0.12}_{-0.06}$	$0.06^{+0.06}_{-0.03}$	$0.30^{+0.29}_{-0.15}$	$0.20^{+0.20}_{-0.10}$	$0.11^{+0.10}_{-0.05}$
MF_{Ca} [Observed; Build-up]	–	$0.37^{+0.26}_{-0.15}$	–	–	$0.74^{+0.55}_{-0.32}$
MF_{Ca} [Steady]	–	$0.48^{+0.34}_{-0.20}$	–	–	$1.02^{+0.73}_{-0.43}$
M_{cvz} [10^{-6} $M_\odot$]	7.08	3.51	0.78	1.57	2.43
$\sum M_Z$ [g]	$> 5.43 \times 10^{21}$	$> 1.92 \times 10^{23}$	$> 2.17 \times 10^{22}$	$> 2.31 \times 10^{22}$	$> 6.53 \times 10^{22}$
$\sum \dot{M}_Z$ [g/s]	$> 3.40 \times 10^7$	$> 2.87 \times 10^9$	$> 1.29 \times 10^9$	$> 6.33 \times 10^8$	$> 1.51 \times 10^9$

Figure A9. Mass fractions of `cecilia`'s detected metals for SDSS J0859+5732 (left) and SDSS J2311−0041 (right), normalised to those of CI chondrites (black line; Alexander 2019a,b). The dashed vertical lines at $4.8\tau_{\mathrm{Mg}}$ and $5\tau_{\mathrm{Mg}}$ indicate the point at which these systems cease to show evidence of oxygen excess, respectively (see Section 6.1).

Table B2. Comparison between the elemental abundances obtained with the classical, line-by-line fitting method of DF12 (left) and those predicted by `cecilia`'s ML-based optimisation routine. The parentheses in the DF12 columns denote the number of absorption lines used to constrain each abundance measurement. The uncertainties of `cecilia`'s results represent the total error (σ_{tot}), calculated as the quadrature sum of the statistical ($\sigma_{\mathrm{stat,MCMC}}$) and systematic error sources ($\sigma_{\mathrm{sys,ML}}$). As discussed in Section 5, `cecilia`'s estimated abundances are consistent with those of DF12 within 0.2 dex.

Property	SDSS J0231+2512		SDSS J0859+5732		SDSS J1109+1318		SDSS J1333+6364		SDSS J2311−0041	
	DF12	cecilia	DF12	cecilia	DF12	cecilia	DF12	cecilia	DF12	cecilia
$\log_{10}$ (H/He)	-6.26 (1)	–	–	-6.60±0.20	-4.39 (1)	-4.33±0.20	-5.40 (1)	-5.49±0.20	-6.07 (1)	-6.26±0.20
$\log_{10}$ (O/He)	–	–	-5.37±0.10 (2)	-5.35±0.20	-5.35±0.19 (2)	-5.66±0.20	-5.49±0.24 (2)	-5.85±0.20	-5.51±0.15 (2)	-5.64±0.20
$\log_{10}$ (Mg/He)	-7.35 (1)	-7.22±0.20	-6.20±0.07 (3)	-6.24±0.20	-6.57 (1)	-6.46±0.20	-6.48 (1)	-6.54±0.20	-6.60 (1)	-6.57±0.20
$\log_{10}$ (Si/He)	-6.90 (1)	–	-6.30±0.10 (3)	-6.48±0.20	-6.44±0.11 (3)	-6.72±0.20	-6.91±0.16 (3)	–	-6.69±0.15 (3)	-6.78±0.20
$\log_{10}$ (Ca/He)	-8.89±0.10 (2)	-8.58±0.20	-8.00±0.17 (3)	-7.86±0.20	-7.58±0.17 (3)	-7.34±0.20	-7.50±0.20 (3)	-7.78±0.20	-7.97±0.04 (3)	-7.94±0.20
$\log_{10}$ (Fe/He)	–	–	-6.62±0.13 (3)	-6.66±0.20	–	–	–	–	-6.77±0.18 (2)	-7.09±0.20