A Data-Driven M Dwarf Model and Detailed Abundances for ∼17,000 M Dwarfs in SDSS-V

Aida Behmard,[1,2] Melissa K. Ness,[3] Andrew R. Casey,[4,5,6] Ruth Angus,[2] Katia Cunha,[7,8] Diogo Souto,[9] Yuxi(Lucy) Lu,[10,11] and Jennifer A. Johnson[12]

[1]*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA; abehmard@flatironinstitute.org**
[2]*American Museum of Natural History, 200 Central Park West, Manhattan, NY 10024, USA*
[3]*Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2601, Australia*
[4]*School of Physics and Astronomy, Monash University, Melbourne, VIC 3800, Australia*
[5]*ARC Centre of Excellence for All Sky Astrophysics in Three Dimensions (ASTRO-3D), Australia*
[6]*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA*
[7]*University of Arizona, Tucson, AZ 85719, USA*
[8]*Observatório Nacional, São Cristóvão, Rio de Janeiro, Brazil*
[9]*Departamento de Física, Universidade Federal de Sergipe, Av. Marcelo Deda Chagas, S/N Cep 49.107-230, São Cristóvão, SE, Brazil*
[10]*Department of Astronomy, The Ohio State University, Columbus, 140 W 18th Ave, OH 43210, USA*
[11]*Center for Cosmology and Astroparticle Physics (CCAPP), The Ohio State University, 191 W. Woodruff Ave., Columbus, OH 43210, USA*
[12]*Department of Astronomy and Center for Cosmology and AstroParticle Physics, Ohio State University, Columbus, OH 43210, USA*

ABSTRACT

The cool temperatures of M dwarf atmospheres enable complex molecular chemistry, making robust characterization of M dwarf compositions a long-standing challenge. Recent modifications to spectral synthesis pipelines have enabled more accurate modeling of M dwarf atmospheres, but these methods are too slow for characterizing more than a handful of stars at a time. Data-driven methods such as *The Cannon* are viable alternatives, and can harness the information content of many M dwarfs from large spectroscopic surveys. Here, we train *The Cannon* on M dwarfs with FGK binary companions from the Sloan Digital Sky Survey-V/Milky Way Mapper (SDSS-V/MWM), with spectra from the Apache Point Observatory Galactic Evolution Experiment (APOGEE). The FGK-M pairs are assumed to be chemically homogeneous and span $-0.56 < $ [Fe/H] < 0.31 dex. The resulting model is capable of inferring M dwarf $T_{\rm eff}$ and elemental abundances for Fe, Mg, Al, Si, C, N, O, Ca, Ti, Cr, and Ni with median uncertainties of 13 K and $0.018-0.029$ dex, respectively. We test the model by verifying that it reproduces reported abundance values of M dwarfs in open clusters and benchmark M dwarf datasets, as well as expected metallicity trends from stellar evolution. We apply the model to 16,590 M dwarfs in SDSS-V/MWM and provide their detailed abundances in our accompanying catalog.

Keywords: stars: abundances

1. INTRODUCTION

M dwarfs are the most common stars, comprising ∼70% of the Galactic stellar population (Miller & Scalo 1979; Bochanski et al. 2010; Henry et al. 2018). Their low masses translate to slow hydrogen fusion rates at their cores, resulting in long main-sequence lifetimes that exceed the age of the universe (e.g., Woolf & Wallerstein 2020). Because M dwarfs are often old, their chemical compositions encode nucleosynthetic processes and interstellar medium enrichment from early gener-

ations of higher mass stars. M dwarfs are thus fossil records of Galactic chemical evolution (Bochanski et al. 2010; Woolf & West 2012). Additionally, M dwarfs are ideal for exoplanet detection and characterization. Their small sizes and low masses lead to strong transit and radial velocity signals, even for small, cool planets that approach the Earth-like regime (Nutzman & Charbonneau 2008; Trifonov et al. 2018). This makes M dwarfs especially attractive targets for exoplanet surveys, e.g., the Transiting Exoplanet Survey Satellite (TESS) and PLAnetary Transits and Oscillations of stars (PLATO) missions (Rauer et al. 2024), and for planetary atmosphere investigations with JWST (e.g., Clampin 2008;

Muirhead et al. 2018). Planet properties gain valuable context with knowledge of host star chemistry, which reflect the compositions of planet building block material from protoplanetary disks. Thus, robust methods for measuring M dwarf chemistry will revolutionize our understanding of both planet formation and the assembly history of our galaxy.

Constraining M dwarf chemical compositions is notoriously difficult. Traditional spectroscopic methods that rely on physical stellar atmosphere models are optimized for solar-like stars with temperatures above ~4500 K (e.g., Brewer & Fischer 2018; Jönsson et al. 2020; Hayes et al. 2022). M dwarfs are cooler, so their atmospheres can harbor molecules (e.g., TiO, VO, MgH, CaH, FeH, H_2O, CO) that create dense clusters of molecular lines in the optical and near-infrared regions of spectra (e.g., Allard et al. 1997; Rojas-Ayala et al. 2012). Traditional stellar models lack the necessary line lists/opacity information to reproduce molecular features (e.g., Mann et al. 2013b), which severely hampers their ability to model M dwarf atmospheres. A few studies have made modifications to line lists and model fitting methodologies that more accurately capture M dwarf chemistry (e.g., Souto et al. 2022; Melo et al. 2024; Hejazi et al. 2024). However, these customized M dwarf modeling pipelines are slow, and have only been used to characterize small (~20 stars; Souto et al. 2022) M dwarf samples to date.

This has spurred development of empirical stellar characterization techniques that do not rely on physical stellar models, such as the *The Cannon* (Ness et al. 2015; Casey et al. 2016). A data-driven framework, *The Cannon* uses a training set of spectra from stars with well-determined "labels" (e.g., stellar parameters and/or elemental abundances) to construct a predictive model of the flux at every pixel in the wavelength range of the spectra. While the training set labels originate from other datasets that may utilize physical stellar models, *The Cannon* is able to construct a model built purely from the labeled training spectra. In that sense, *The Cannon* does not utilize physical stellar models; it only requires a training set of spectra with high quality labels that span the label parameter space of the stars we seek to characterize. The resulting data-driven model can then infer the parameter/abundance labels of other stellar samples given their spectra. This makes *The Cannon* a valuable tool for characterizing cool stars that current stellar models struggle with, e.g., M dwarfs (Behmard et al. 2019; Birky et al. 2020; Galgano et al. 2020; Rains et al. 2024). Data-driven approaches are also computationally inexpensive compared to methods that rely on stellar atmosphere modeling, so *The Cannon* is well-

suited to characterizing stellar samples from large spectroscopic surveys.

Here, we use *The Cannon* to infer a wide set of elemental abundances for ~17,000 M dwarfs with Milky Way Mapper (MWM) data from SDSS-V, the current phase of the Sloan Digital Sky Survey. We outline *The Cannon* in Section 2, and describe the MWM data and how we processed it for *The Cannon* in Section 3. We apply *The Cannon* to a small (~20 stars) benchmark M dwarf sample and evaluate its performance in Section 4. In Section 5, we construct a larger M dwarf training set drawn from SDSS-V/MWM and assess its improved performance by testing against the aforementioned benchmark sample and M dwarfs from the Hyades open cluster. We apply the SDSS-V/MWM training set to a sample of ~17,000 SDSS-V/MWM M dwarfs in Section 6, and report their abundances for Fe, Mg, Al, Si, C, N, O, Ca, Ti, Cr, and Ni.

2. *THE CANNON*

The Cannon is a regression model that operates under two key assumptions: that stellar spectra with identical labels look identical at every pixel, and that the flux at every pixel in a spectrum changes continuously as a function of the stellar labels. *The Cannon* can infer labels for stellar samples given their spectra via a two-step process: a "training" step in which the spectra and labels of stars that compose the training set are used to construct a predictive model of the flux at every pixel in the wavelength range, and a "test" step in which the model is applied to spectra of other stars in order to infer their labels. In principle, the labels can be any physical parameters that typically parameterize stellar atmosphere models (e.g., T_{eff}, $\log g$, [Fe/H], etc.), or empirical labels that serve as proxies for physical parameters (e.g., spectral types, colors, magnitudes, etc.) (Birky et al. 2020). Whatever the training set labels are, they must be high-quality as the inferred labels will only be as accurate as the training labels, and only precise if the training labels are measured consistently across the training set stars. The training set labels must also span the parameter space of the true labels that we seek to infer for the test set, because *The Cannon* does not extrapolate well outside the training set parameter space.

In this study, we use *The Cannon 2*, the second implementation of *The Cannon* developed by Casey et al. (2016). Hereafter, we will refer to *The Cannon 2* simply as *The Cannon*. This version allows for more complex flux models than the original, which aids inference of more labels, e.g., a large set of abundances beyond [Fe/H]. It has also been used to successfully infer M

dwarf labels in previous studies (e.g., Behmard et al. 2019).

2.1. Training Step

The Cannon flux model for a star n at wavelength pixel j is expressed as

$$f_{jn} = v(\ell_n) \cdot \theta_j + e_{jn}, \tag{1}$$

where θ_j is a vector containing the set of flux model coefficients at each pixel j, and $v(\ell_n)$ is a function of the label list ℓ_n that is unique for each spectrum n. The function $v(\ell_n)$ is referred to as the "vectorizer" which can accommodate functions beyond simple polynomial expansions of the label list ℓ_n (e.g., sums of sines and cosines). The noise term e_{jn} can be taken as drawn from a Gaussian with zero mean and variance $\sigma_{jn}^2 + s_j^2$, where σ_{jn}^2 is the flux uncertainty reported for the training set spectra, and s_j^2 is the intrinsic scatter of the model at each pixel j. This intrinsic scatter is analogous to the expected deviation of the model from the spectrum at pixel j. We can relate the flux model to a single-pixel log-likelihood function:

$$\ln p(f_{jn}|\theta_j, l_n, s_j^2) = -\frac{[f_{jn} - v(l_n) \cdot \theta_j]^2}{\sigma_{jn}^2 + s_j^2} \\ -\ln(\sigma_{jn}^2 + s_j^2) + \Lambda_j Q(\theta_j), \tag{2}$$

where Λ_j is a regularization parameter and $Q(\theta_j)$ is a regularizing function that encourages the flux model coefficients θ_j to approach zero, which combats overfitting. This is potentially useful for inferring label sets that include many elemental abundances because we expect that only a small number of abundances will affect the flux at each pixel in the wavelength range.

In the training step, each log-likelihood is maximized to derive the best-fit model coefficients θ_j and intrinsic scatter s_j^2 at each wavelength pixel j:

$$\theta_j, s_j^2 = \underset{\text{argmax}}{\theta_j, s_j} \left[\sum_{n=0}^{N-1} \ln p(f_{jn}|\theta_j, l_n, s_j^2) \right]. \tag{3}$$

2.2. Test Step

In the test step, we use the optimized model coefficients and scatter (θ_j, s_j^2) in our set of single-pixel log-likelihood functions to infer the label list. We do this by maximizing the log-likelihood functions, or equivalently, minimizing the negative log-likelihood functions:

$$l_n = \underset{l_n}{\text{argmin}} \left[\sum_{j=0}^{J-1} -\ln p(f_{jn}|\theta_j, l_n, s_j^2) \right] \\ = \underset{l_n}{\text{argmin}} \left[\sum_{j=0}^{J-1} -\frac{[f_{jn} - v(l_n) \cdot \theta]^2}{\sigma_{jn}^2 + s_j^2} \right], \tag{4}$$

where we can consider $-\ln p(f_{jn}|\theta_j, l_n, s_j^2)$ (hereafter χ^2) a goodness-of-fit metric that assesses how well the flux model fits the actual spectra of each test set star.
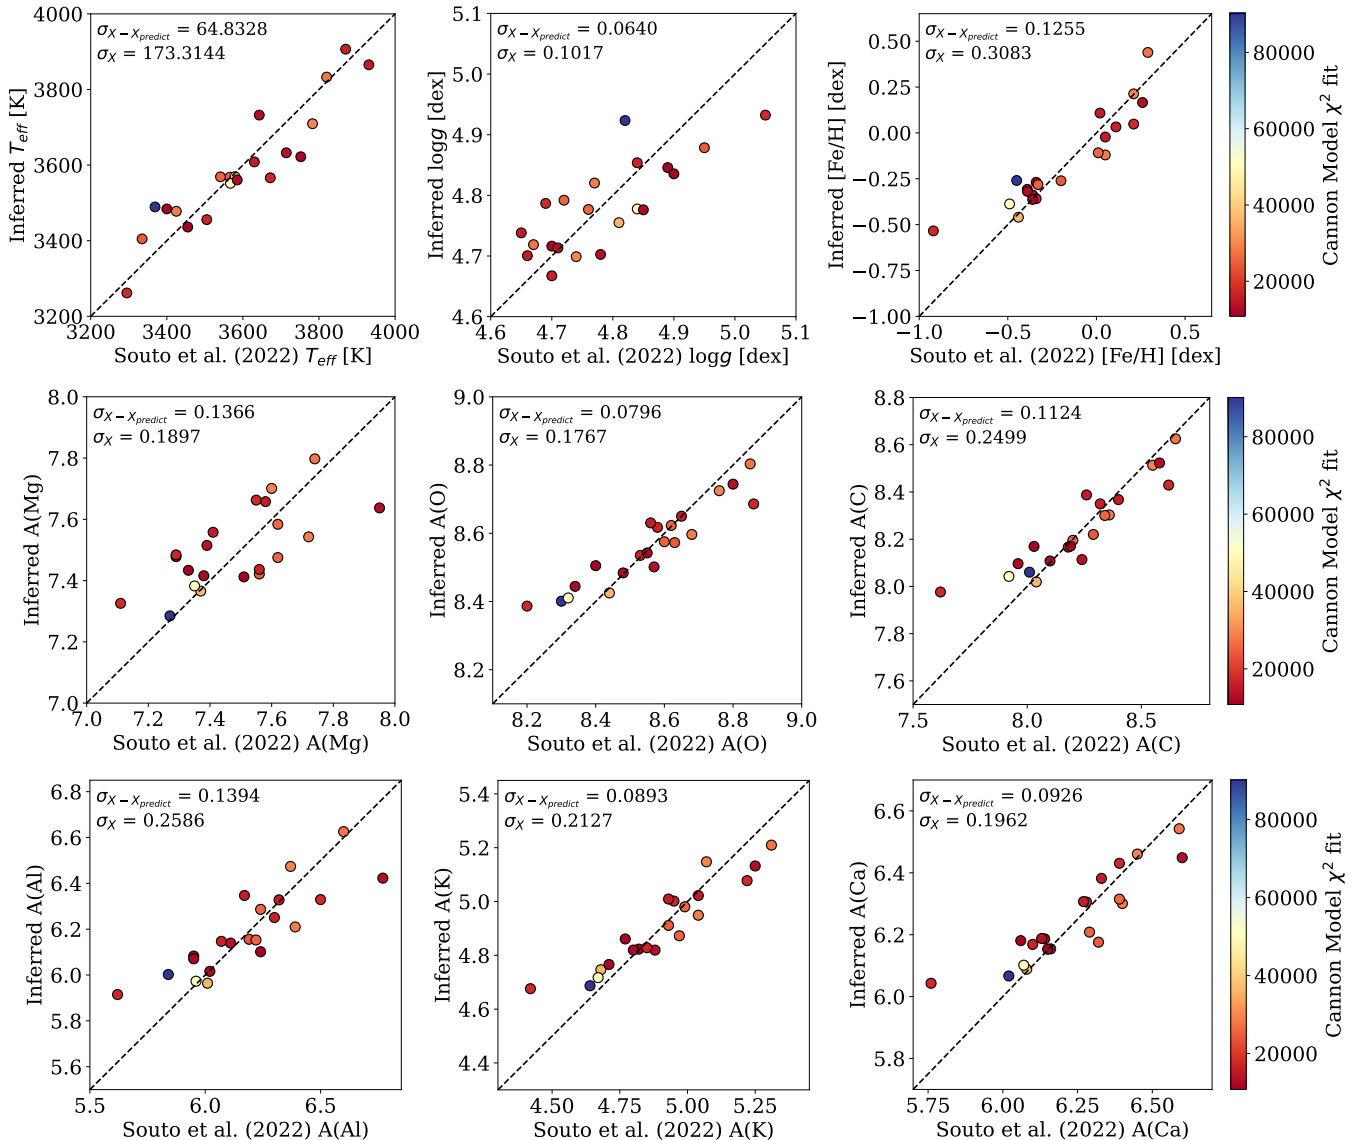
Figure 1. 1-to-1 plots of our inferred versus the reported labels for the Souto et al. (2022) sample from our LOOCV scheme with *The Cannon*. The points are colored by the χ^2 of the flux model fit. We report the rms scatter between the inferred and Souto et al. (2022) labels in the top left of each plot, and the scatter of the Souto et al. (2022)-reported values for that label below. The former values are smaller than the latter for every label, indicating that *The Cannon* recovers the labels with rms scatter well within the intrinsic label scatter.

3. SDSS-V/MWM DATA

The SDSS-V/MWM survey has an H-band ($1.51-1.7$ μm) component that employs the Apache Point Observatory Galactic Evolution Experiment (APOGEE) spectrographs, through which it conducts a high resolution ($R \sim 22,500$) survey that currently provides over 1 million stellar spectra through ongoing observations (Majewski et al. 2017; Wilson et al. 2019; Almeida et al. 2023). While the APOGEE survey is primarily designed for targeting bright objects, namely red giants, it has surveyed $\sim$50,000 M dwarfs to date. The APOGEE Stellar Parameter and Chemical Abundances Pipeline

(ASPCAP) provides abundances for a wide set of elements at typical precisions of <0.1 dex (García Pérez et al. 2016). However, main-sequence stars with <4500 K temperatures are known to have problematic ASPCAP abundances, making ASPCAP unreliable for M-dwarfs. This motivates our work on constructing a data-driven model for inferring detailed M dwarf abundances.

APOGEE provides pseudo-continuum-normalized, rest-frame-shifted, co-added spectra across all observed epochs. However, we opt to use the non-pseudo-continuum-normalized spectra because *The Cannon* requires input spectra to be normalized via a linear operation; this is not satisfied by the APOGEE pseudo-

continuum-normalization procedure that makes use of high-order polynomials (Jönsson et al. 2020). Instead, we carried out continuum normalization via error-weighted, broad Gaussian smoothing with

$$\bar{f}(\lambda_0) = \frac{\sum_j (f_j \sigma_j^{-2} w_j(\lambda_0))}{\sum_j (\sigma_j^{-2} w_j(\lambda_0))} \,, \qquad (5)$$

where f_j is the flux at pixel j of the wavelength range, σ_j is the uncertainty at pixel j, and the weight $w_j(\lambda_0)$ is drawn from a Gaussian:

$$w_j(\lambda_0) = e^{-\frac{(\lambda_0 - \lambda_j)^2}{L^2}} \,, \qquad (6)$$

where L is chosen to be 10 Å, which is slightly larger than typical absorption features in the APOGEE spectra. Gaussian smoothing is often used to continuum normalize spectra in preparation for *The Cannon*, with L adjusted according to the spectral resolution (e.g., Ho et al. 2017; Behmard et al. 2019; Rampalli et al. 2021).

The APOGEE spectra signal-to-noise (SNR) ratios are typically high, with average SNR levels for M dwarfs of approximately 100/pix. However APOGEE M dwarf spectra are often reported with SNR $\gtrsim 200$/pix. These high SNR targets have underestimated flux uncertainties that artificially inflate flux model spectral fit χ^2 values from *The Cannon*. To address this, we set all <0.005 flux uncertainty values to 0.005; typical normalized flux uncertainties should scale as ~ 1/SNR, and APOGEE targets have maximum effective SNR levels of ~ 200/pix. Thus, flux uncertainty values smaller than $1/200 = 0.005$ are unrealistic.

4. SMALL TRAINING SET LIMITATIONS

Previous studies demonstrate that *The Cannon* can infer detailed elemental abundances for red giants (Casey et al. 2016) and FGK dwarfs (e.g., Rice & Brewer 2020; Rampalli et al. 2024; Angelo et al. 2024), but it has never been used to infer abundances beyond [Fe/H] or occasionally [Ti/Fe] for M dwarfs (Behmard et al. 2019; Birky et al. 2020; Rains et al. 2024). This is largely because we lack robust M dwarf training sets with well-determined abundances for other elements. The largest M dwarf dataset with a wide set of abundance measurements is from Souto et al. (2022). This dataset consists of 21 M dwarfs observed during SDSS-IV/APOGEE-1 (Blanton et al. 2017), with spectra collected by the APOGEE-N instrument (Gunn et al. 2006; Wilson et al. 2019) and reduced as part of the 16th SDSS data release (DR16; Nidever et al. 2015; Ahumada et al. 2020; Jönsson et al. 2020). In short, Souto et al. (2022) employed LTE MARCS model atmospheres (Gustafsson

et al. 2008) and the TurboSpectrum spectral synthesis code (Alvarez & Plez 1998; Plez 2012), along with a custom line list based on the APOGEE line list from DR17 (e.g., Smith et al. 2013; Hasselquist et al. 2016; Cunha et al. 2017). This enabled them to measure T_{eff}, log g, Fe, C, O, Mg, Al, K, and Ca for their 21 M dwarfs.

We find that the small Souto et al. (2022) sample is an effective training set for constructing a model with *The Cannon* that reproduces the reported abundances to high precision. We demonstrate this through a leave-one-out cross-validation (LOOCV) scheme in which we train *The Cannon* on all 21 M dwarfs but one, and then infer the detailed abundances of the removed M dwarf with the resultant model trained on $N-1$. We loop through the entire Souto et al. (2022) sample with this procedure to infer abundances for each M dwarf. For this test, we use SDSS-V/MWM spectra for the Souto et al. (2022) sample rather than the SDSS-IV/APOGEE DR16 spectra used in the original study. The SNR levels range from $60-570$/pix.

The models are functions of the parameter and abundance labels, listed in the label vector below:

$$\ell_n = [1, \ T_{\text{eff}}, \ \log g, \ [\text{Fe/H}], \ A(\text{Mg}), \ A(\text{O}), \ A(\text{C}), \\ A(\text{Al}), \ A(\text{K}), \ A(\text{Ca})] \qquad (7)$$

We test models that are both linear and quadratic functions of the labels. The linear models are significantly faster to train (which is unsurprising given the large set of abundance labels), and exhibit slightly better performance in recovering the reported abundances. This unsurprising because the small Souto et al. (2022) sample does not span a large metallicity range, and many elemental abundances are correlated with each other. Thus, quadratic label terms likely introduce excessive/unneeded model complexity. Consequently we employ linear models for this LOOCV test.

The LOOCV results are shown in Figure 1. We find that our LOOCV implementation reproduces the detailed M dwarf abundances reported in Souto et al. (2022) to precisions of $0.09-0.14$ dex. These precisions are affected by the uncertainties on the Souto et al. (2022) abundances, which are reported as $0.02-0.15$ dex and constitute uncertainties on our model labels. To calculate these uncertainties, Souto et al. (2022) propagated the uncertainties on their adopted atmospheric parameters (T_{eff}, log g, [M/H], C/O) and pseudo-continuum normalization procedure (Souto et al. 2017, 2022). While our LOOCV abundance precisions are affected by these reported abundance uncertainties, it is difficult to disentangle their contribution from other error sources in the APOGEE data reduction process, and

the inherent scatter of our model. We reserve a longer discussion of uncertainties on our inferred abundances for later in the manuscript (see Section 6).

The peak flux model spectral fit χ^2 of all the Souto et al. (2022) M dwarfs is ∼12,000, which translates to a reduced χ^2 of ∼1.6 considering that there are ∼7400 pixels in the APOGEE wavelength range (which can be considered the number of degrees of freedom). An excellent flux model fit would yield a reduced χ^2 of approximately 1, but our peak χ^2 is slightly higher because our flux model fits are imperfect, which is unsurprising given the complexity of M dwarf spectra. Still, a reduced χ^2 of ∼1.6 indicates a good flux model fit with *The Cannon* (e.g., Birky et al. 2020; Rampalli et al. 2024).

There is one M dwarf (SDSS ID = 80419035) in the sample with an anomalously large flux model spectral fit χ^2 value of ∼90,000. We removed this star and re-ran the LOOCV scheme, and found only marginal improvement to the inferred abundance precisions. This star is a fast rotator with $v \sin i = 13.5 \pm 2.0$ km s^{-1} (Souto et al. 2020). Upon visual inspection of its spectrum, we see no obvious rotational broadening, but rapid rotation can induce magnetic activity in late type M dwarfs that affect the flux in other ways (e.g., Suárez Mascareño et al. 2016). For example, through flares, starspots, and even radii inflation due to magnetic inhibition of convection. All these factors likely contribute to the relatively poor flux fit our model achieves for this star. In general however, this scheme successfully recovers the correct M dwarf abundances to high precisions. This demonstrates, for the first time, that *The Cannon* is capable of inferring precise M dwarf abundances for a wide set of elements.

5. FGK-M TRAINING SET

We use SDSS-V/MWM to construct a larger training set of 79 M dwarfs (compared to 21 M dwarfs in the Souto et al. (2022) sample) with detailed abundance measurements. Because ASPCAP elemental abundances are unreliable for M dwarfs, we cannot simply use the ASPCAP abundances as labels for *The Cannon*. Instead, we tag M dwarfs with the ASPCAP abundances of FGK dwarf binary companions. Unlike M dwarfs, FGK stars are well-modeled by ASPCAP because physical stellar models and synthetic spectra are reliable in the solar-like regime. Tagging M dwarfs with the abundances of solar-like binary companions is considered a "gold standard" method (e.g., Mann et al. 2013a, 2014; Newton et al. 2014; Maldonado et al. 2020; Duque-Arribas et al. 2024). It is based on the assump-

tion that binary companions share a parent molecular cloud and thus formed from the same material, making them approximately chemically homogeneous at birth (e.g., De Silva et al. 2007, 2009; Bland-Hawthorn et al. 2010). However, there is a caveat to this assumption; solar-like stars are affected by diffusion processes (i.e., gravitational settling and radiative levitation) that alter their surface abundances over time (Dotter et al. 2017; Souto et al. 2019). Conversely, M dwarfs are relatively immune from diffusion effects because they have deeper convective envelopes, and diffusion efficiency decreases with depth (e.g., Liu et al. 2019; Moedas et al. 2022; Wanderley et al. 2023). This may produce abundance variations of 0.01−0.12 dex between FGK and M dwarf companions assuming solar age (Choi et al. 2016). We can potentially correct for diffusion processes by reporting abundances in the form of [X/Fe] rather than [X/H] because diffusion effects are roughly similar across different elements. This was done by Souto et al. (2022) to enable comparison of their M dwarf abundances to FGK companion abundances. We leave this as an option in our analysis if diffusion appears to be an issue.

Table 1. M Dwarf Training Set Properties

Gaia DR3 ID	SDSS ID	T_{eff} K	[Fe/H] dex	[Mg/H] dex	[Al/H] dex	[Si/H] dex	[C/H] dex	...
4665710629633988736	91725384	3910	-0.10 ± 0.01	0.01 ± 0.01	0.01 ± 0.03	-0.09 ± 0.02	-0.07 ± 0.02	
4819175927157254784	92916465	3704	0.10 ± 0.01	0.12 ± 0.02	0.12 ± 0.03	0.07 ± 0.02	0.05 ± 0.03	
4763618910272182400	92524557	3396	0.11 ± 0.01	0.14 ± 0.01	0.21 ± 0.03	0.14 ± 0.02	0.14 ± 0.02	
831399531274288512	57258526	3543	0.18 ± 0.01	0.14 ± 0.02	0.25 ± 0.03	0.21 ± 0.02	0.03 ± 0.03	
839567326416583808	57301576	3382	0.06 ± 0.01	0.12 ± 0.02	0.05 ± 0.03	0.02 ± 0.02	0.08 ± 0.02	
3922083599776523520	80399914	3637	-0.10 ± 0.01	-0.10 ± 0.02	-0.15 ± 0.03	-0.13 ± 0.02	-0.13 ± 0.03	
1534086421765263360	61843665	3654	-0.29 ± 0.01	-0.03 ± 0.01	-0.06 ± 0.03	-0.15 ± 0.02	-0.09 ± 0.02	
2581806619466249728	70532624	3729	-0.26 ± 0.01	-0.15 ± 0.01	-0.20 ± 0.03	-0.20 ± 0.02	-0.31 ± 0.03	
315289980082496256	116701844	3551	-0.08 ± 0.01	-0.14 ± 0.01	-0.17 ± 0.03	-0.15 ± 0.02	-0.15 ± 0.02	
3637535866023121536	78881207	3735	0.08 ± 0.01	0.12 ± 0.01	0.13 ± 0.03	0.10 ± 0.02	0.09 ± 0.02	
935566511271187072	58005724	3419	0.22 ± 0.01	0.26 ± 0.01	0.33 ± 0.03	0.25 ± 0.02	0.27 ± 0.02	
678053878665346816	56426757	3337	0.01 ± 0.01	-0.04 ± 0.01	-0.10 ± 0.03	-0.03 ± 0.02	0.02 ± 0.02	
702082159096764288	56575737	3516	0.03 ± 0.01	0.02 ± 0.01	0.00 ± 0.03	0.03 ± 0.02	-0.01 ± 0.02	
					...			

NOTE—This table lists the properties of the M dwarfs in our FGK-M training set drawn from SDSS-V/MWM. The properties consist of all labels used to construct our flux models with *The Cannon* (photometric temperatures and abundances for all elements of interest). The abundances and their errors are the reported ASPCAP abundances of the FGK companions. We only list a subset of the abundances in this table, but the full set is provided in the downloadable version.

(This table is available in its entirety in machine-readable form.)

To identify FGK-M binary systems in SDSS-V/MWM, we cross-match all SDSS-V stars with AS-PCAP results with the El-Badry et al. (2021) binary catalog. Because the ASPCAP log g values are unreliable for cool M dwarfs, we select M dwarf secondaries in binary systems according to M dwarf type-color sequence relations from Pecaut & Mamajek (2013):

$$G - RP > 0.92$$
$$G + 5\log(\varpi/100) > 8.16 \tag{8}$$

where G and RP are photometric passbands and ϖ is the parallax, all from *Gaia* DR3 (Gaia Collaboration et al. 2023). To ensure the primaries in the remaining binary systems are FGK dwarfs with reliable abundances, we apply the following ASPCAP parameter and abundance cuts:

$$\log g > 4 \text{ dex}$$
$$T_{\text{eff}} = 4500 - 6500 \text{ K}$$
$$[X/H]_{\text{err}} < 0.2 \text{ dex} \tag{9}$$
$$\text{Flags x\_h\_flags not set}$$

The T_{eff} cut delineates the range within which AS-PCAP elemental abundances are reliable. Finally, we apply a cut of SNR > 50 pix^{-1} to the M dwarfs to ensure that their spectra are sufficiently high quality for

training robust flux models with *The Cannon*. This resulted in a final training set of 79 FGK-M binaries with SNR = 51−440/pix. The FGK companions span −0.56 < [Fe/H] < 0.31 dex, and the abundance uncertainties across all elements we use as training labels range from 0.01−0.14 dex.

We provide the M dwarf training set labels in Table 1. These consist of the M dwarf T_{eff} values, and elemental abundances from the FGK companions. While the AS-PCAP T_{eff} values are not bad, they are still derived from extrapolations to models that are not designed to accommodate M dwarfs. For this reason, it is hard to say how truly reliable they are across the entire cool temperature range (3000−4000 K) spanned by our *Cannon* model. A handful of our training set M dwarfs also lack ASPCAP T_{eff}. Instead, we use an empirical color-temperature relation (Curtis et al. 2020) and *Gaia* DR3 photometry to calculate new T_{eff} values as training set labels. We check the color extinction values using the Bayestar19 3D dust map (Green et al. 2019) implemented in the dustmaps Python package (Green 2018). The extinction coefficients are taken from Danielski et al. (2018). >95% of our training set M dwarfs have E(BP−RP) < 0.03, corresponding to potential T_{eff} shifts due to reddening of <20 K. Two of our M dwarfs have higher extinction values of ∼0.06 and ∼0.16, corresponding to potential T_{eff} shifts of ∼40 K and ∼140 K, but both these stars have reported ASPCAP T_{eff} that agree to within 60 K
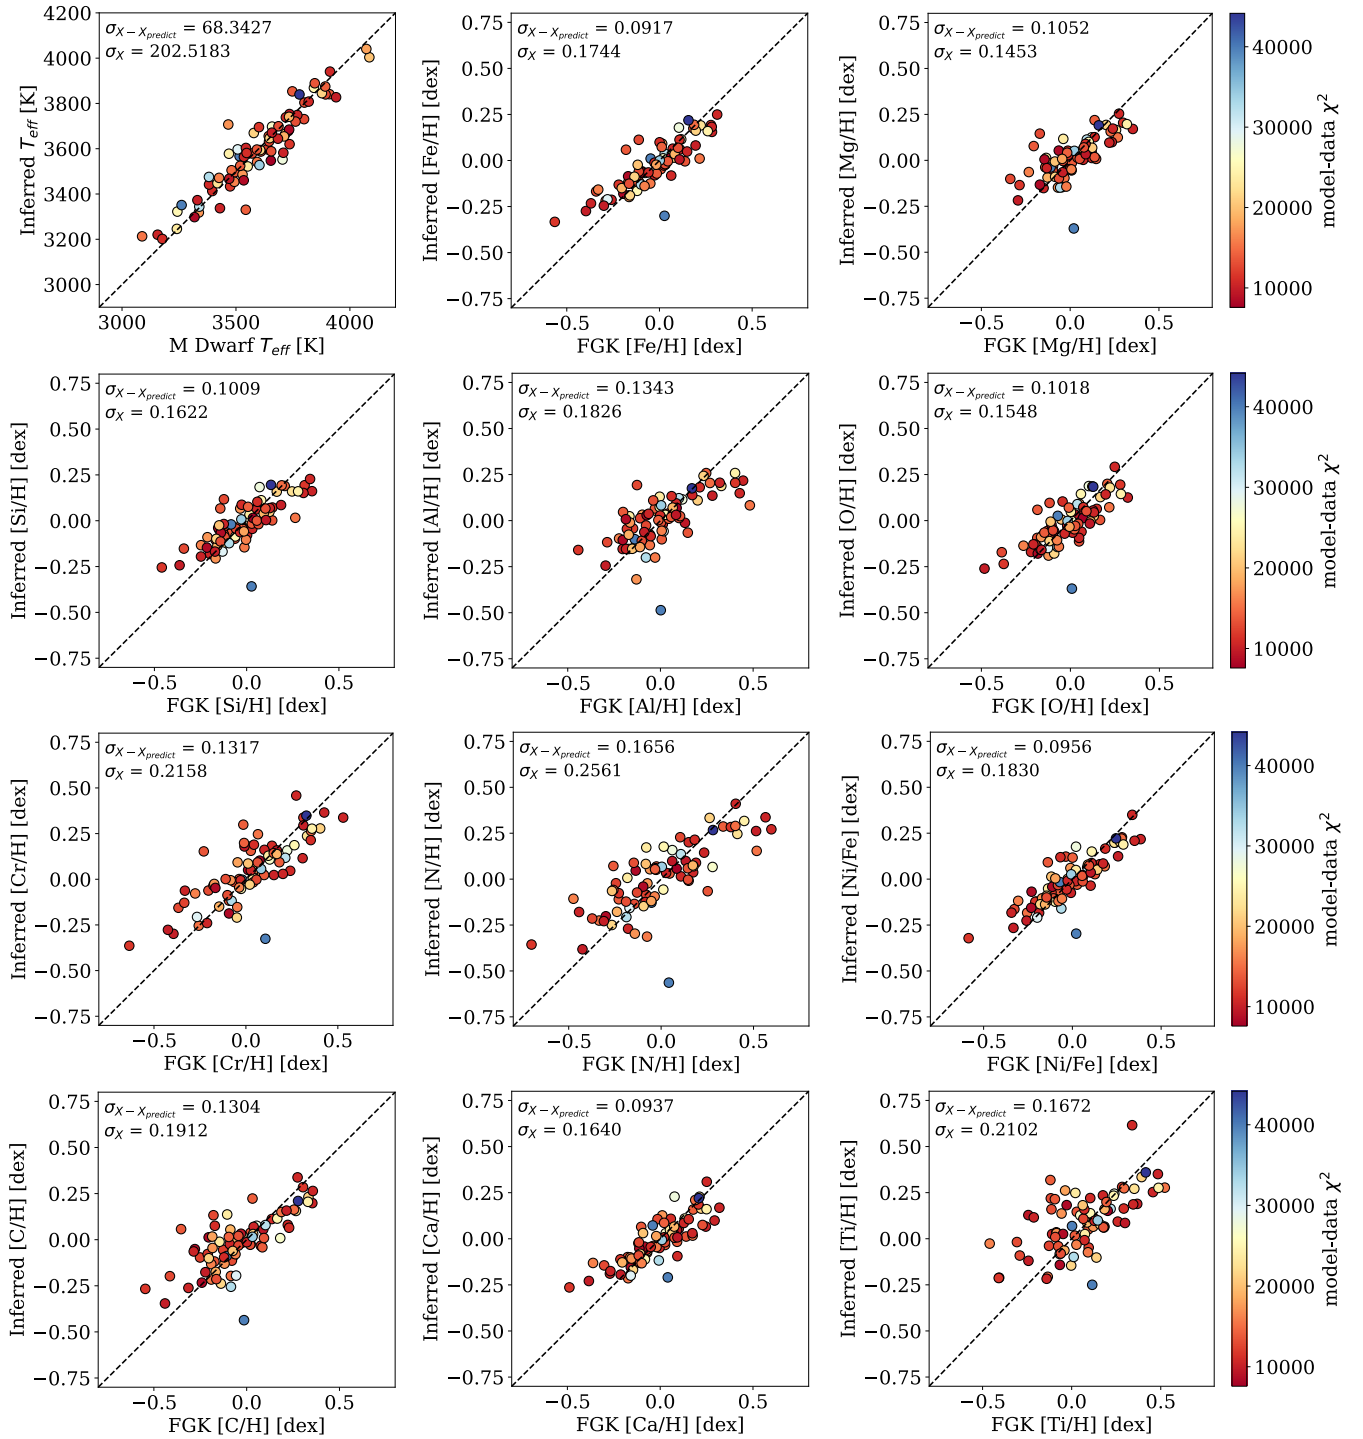
Figure 2. 1-to-1 plots of our inferred versus the reported labels for our FGK-M training set after applying LOOCV with *The Cannon*. The points are colored by the χ^2 of the flux model spectral fit. We report the rms scatter between the inferred M dwarf and FGK companion labels in the top left of each plot, and the scatter of the FGK companion labels below. The former values are smaller than the latter for every label, indicating that *The Cannon* recovers the labels using the M dwarf spectra with rms scatter that are well within the intrinsic FGK label scatter.

of their photometric $T_{\rm eff}$ from the Curtis et al. (2020) relation. We conclude that the photometric $T_{\rm eff}$ values of our training set M dwarfs are not affected by reddening.

After training our model on this FGK-M binary sample, we provide the flux model coefficients θ_j at each pixel j for each label in Figures A.1, A.2, and A.3. We compare the abundance label coefficients with the APOGEE line list (Smith et al. 2021), and additional M dwarf lines identified in Souto et al. (2022). We find that the coefficient amplitudes are often large at the locations of strong absorption features. This indicates that our flux models include reliable abundance information.

5.1. *Leave-One-Out Cross-Validation*

To assess the validity of the FGK-M abundance labels, we apply LOOCV cross-validation. Our label vector again corresponds to a linear model, justified by the narrow metallicity range of our FGK-M training set. It includes a large set of abundances:

$$\ell_n = [1,\ T_{\rm eff},\ [Fe/H],\ [Mg/H],\ [Al/H],\ [Si/H],$$
$$[C/H],\ [N/H],\ [O/H], [Ca/H],\ [Ti/H],\quad (10)$$
$$[V/H],\ [Cr/H],\ [Ni/H]]$$

The elemental abundance labels are a combination of those considered most reliable for dwarfs from the ASPCAP pipeline (Fe, C, Mg, Si, Ni), and those involved in common molecular species in M dwarf atmospheres (O, N, Ti, V, Ca, Cr) (e.g., Rajpurohit et al. 2018). We also include Al because it is well-represented in bulk Earth composition (McDonough 2003), and is therefore common in rocky planet-forming material. Including this wide set of abundances in our model results in optimal performance; subsets of these abundances do not recover the known training set labels as well. This indicates that each of these elements contain spectral information content, and contribute significantly to M dwarf chemical compositions. The M dwarf $T_{\rm eff}$ label values are derived from the Curtis et al. (2020) color-temperature relation as outlined earlier. We do not include log g as a label in our model because M dwarfs evolve slowly, so their physical properties do not change much over the age of the universe after they reach the zero-age main sequence. Because of this, the information contained in M dwarf log g and metallicity can be considered redundant (Birky et al. 2020).

We illustrate our LOOCV results in Figure 2. The χ^2 values for all M dwarfs in our training set peak at $\sim$10,000 (reduced χ^2 of $\sim$1.4), indicating that they are well-fit by the flux models from *The Cannon*. There is one noticeable outlier star in the 1-to-1 abundance plots with a relatively high χ^2. It does not have an unusual SNR, FGK companion separation, or reddening level according BP−RP color. It also seems to occupy a well-populated area of abundance parameter space (solar-like) according to the abundances of its FGK companion. Still, it is possible that this star is not well-represented by the N-1 training set stars in terms of its parameters and/or chemistry. It also does not appear to be a fast rotator based on visual inspection of the flux model fits, but it may still be rapidly rotating as in the case of SDSS ID = 80419035 in the Souto et al. (2022) sample. Thus, its relatively poor flux fit may be due to rapid rotation and/or magnetic activity.

In terms of precision, LOOCV recovers the photometric M dwarf $T_{\rm eff}$ to 68 K, and the FGK companion abundance labels to 0.09−0.17 dex in rms scatter. These values are smaller than the scatter of the FGK labels themselves (second row of values in top left corners of each Figure 2 panel), indicating that our inferred labels are high precision, and are not just reproducing the input label scatter. The only exception is vanadium, which has a high rms scatter of 0.33 dex and does not exhibit a convincing 1-to-1 trend between the inferred and FGK abundances. For this reason we do not include V in Figure 2. This suggests that the ASPCAP V abundances for the FGK dwarf companions are inaccurate, or are not mapped well to the spectra of their M dwarf companions. The former may be true considering that V is measured from a single, weak line in APOGEE spectra (e.g., Grilo et al. 2024), and the latter may also be true because hyperfine structure affects vanadium spectral features of cool versus solar-like stars differently (Shan et al. 2021). Nevertheless, including V as a label in our model improves the precision of our other inferred abundances, indicating that the V abundances still contain useful information.

We also ran our LOOCV scheme with FGK abundance labels in the form of [X/Fe] rather than [X/H], but find that our results are worse in terms of inferred abundance precisions and 1-to-1 trends. We conclude that there is no noticeable advantage to using [X/Fe] in the interest of mitigating differential diffusion effects on the surface abundances of FGK and M dwarf companions. It is possible that any advantage is wiped out by compounding the uncertainties on Fe and X, where X is any other element of interest. Additionally, X and Fe will be strongly correlated if they share nucleosynthetic channels, which will cause [X/Fe] to lack much of the information inherent in X. In such cases, abundances in the form [X/Fe] will not be informative labels for training *The Cannon*. We conclude that [X/H] abundance labels are best for our purposes.
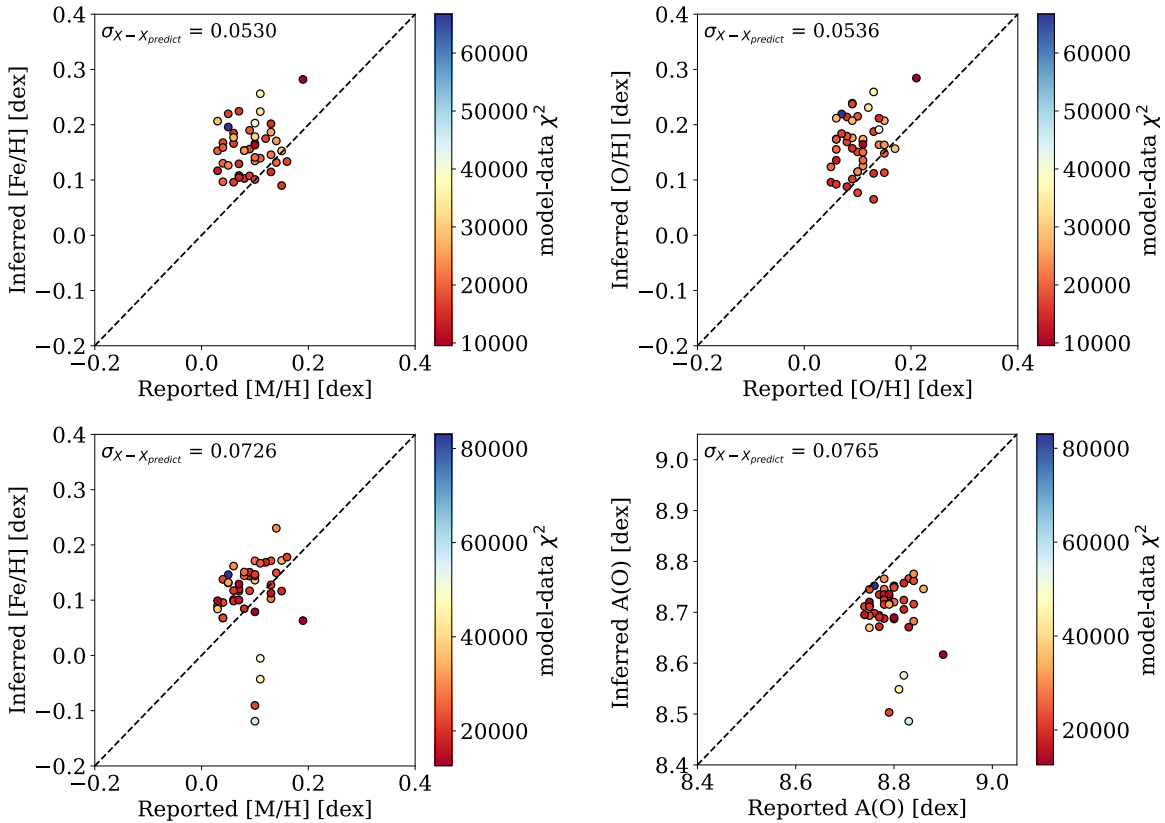
Figure 3. 1-to-1 plots of our inferred abundances vs. the reported abundances of M dwarfs from the Hyades cluster using either our FGK-M sample (top row) or the Souto et al. (2022) sample (bottom row) as the training set. The points are colored by the χ^2 of the flux model spectral fit. We report the rms scatter between the inferred and reported abundance values from Wanderley et al. (2023) in the top left of each plot, and the scatter of the reported values for that label below. The rms scatter from training on our FGK-M sample ($\sim$0.05 dex) are lower than those from training on the Souto et al. (2022) sample ($\sim$0.07−0.08 dex).

As in the case of LOOCV on the Souto et al. (2022) sample, our abundance precisions from LOOCV on our FGK-M sample are affected by the label uncertainties. In this case, these could be considered the reported uncertainties on the FGK companion ASPCAP abundances. However, it is not clear how reliable the ASP-CAP abundance uncertainties are (discussion with A. R. Casey). In our final sample, we use M dwarfs with repeat APOGEE observations to more robustly determine uncertainties on our inferred abundances (see Section 6).

5.2. *Validation with Hyades and Souto et al. (2022) Datasets*

To examine how much better our FGK-M sample performs as a training set compared to the Souto et al. (2022) sample, we tested how well each reproduces reported [M/H] and A(O) measurements for M dwarfs in the Hyades open cluster (Wanderley et al. 2023). We use the same implementation of *The Cannon* as in our FGK-M LOOCV test, namely a linear model in $T_{\rm eff}$ and all considered elemental abundances in Equation 10. Our results are shown in Figure 3. Our FGK-M sample (Figure 3, upper row) does a noticeably better job at reproducing the Hyades abundances from Wanderley et al. (2023) compared to the Souto et al. (2022) sample (Figure 3, lower row). In the latter case there are a handful of M dwarfs with very discrepant inferred versus reported abundances. This is likely because these few Hyades M dwarfs fall into a sparsely populated region of the abundance parameter space spanned by the Souto et al. (2022) sample, which is small and poorly populated in general because the Souto et al. (2022) sample
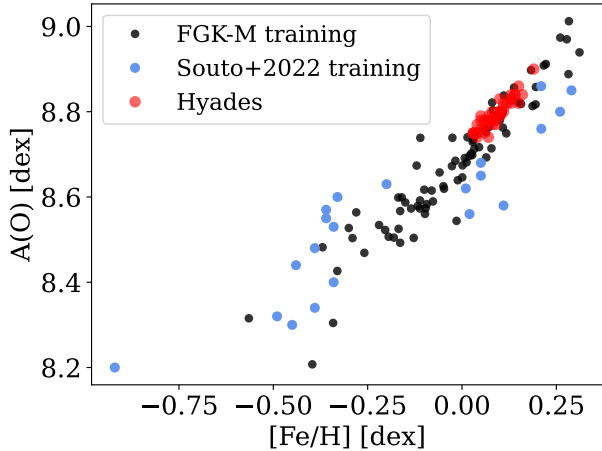
Figure 4. The A(O) and [M/H]/[Fe/H] parameter space spanned by the Hyades M dwarfs from Wanderley et al. (2023) (red), the training set of 21 Souto et al. (2022) M dwarfs (blue), and our FGK-M training set (black). Because the Souto et al. (2022) sample is small, its abundance parameter space is small and sparse, and does not cover the Hyades sample. The region of [M/H]/[Fe/H] ≈ 0.1 dex and A(O) ≈ 8.8 dex is particularly sparse, and likely responsible for the handful of large outliers in inferred vs. reported abundances in the lower panels of Figure 3.

is small (Figure 4, blue points). In contrast, training on our FGK-M sample does not result in large discrepancies between inferred versus reported Hyades M dwarf abundances. Our FGK-M sample performs better because it is bigger than the Souto et al. (2022) sample, and thus spans a larger and more well-populated abundance parameter space (Figure 4, black points). Our inferred abundances agree with the reported Hyades abundances to within ∼0.05 dex (as opposed to ∼0.07−0.08 dex with the Souto et al. (2022) training set). For reference, the scatter of the metallicities and oxygen abundances reported in Wanderley et al. (2023) is 0.03 dex. Compared to these reported abundances, our inferred abundances also adhere more convincingly to the 1-to-1 trend. We also do not see any significant inferred abundance trends with $T_{\rm eff}$, which is reassuring. The flux model spectral fit χ^2 values are also lower compared to those from training on Souto et al. (2022), which peak at ∼16,000 as opposed to ∼23,000 when we trained on the Souto et al. (2022) sample. Figure 5 illustrates the improved flux model fit for an example Hyades M dwarf from the Wanderley et al. (2023) sample, whose abundances are not well-inferred using the Souto et al. (2022) training set. The left panel exhibits the model fit from Souto et al. (2022), and the right panel exhibits the model fit from our larger FGK-M sample. The latter fit is a clear improvement, especially around prominent OH lines used to infer the oxygen abundance [O/H] or A(O).

We use the Hyades sample to test adding regularization to our model. A regularization parameter Λ is built into *The Cannon* and can be assigned different strengths, which encourage model coefficients to take on zero values. This results in simpler models that are less prone to overfitting. We test regularization parameter values ranging from $\Lambda = 10^2$ to 10^5, and find that including regularization always results in less precise label predictions. For this reason we do not include regularization in our model. This is somewhat surprising as Casey et al. (2016) found that including regularization resulted in better inference of a wide set of red giant elemental abundances. However, it is possible that many different abundances contribute to the flux at each wavelength point in M dwarf spectra, making overfitting less of an issue. This would make sense given the complexity of M dwarf spectra (e.g., large number of molecular lines) due to the low $T_{\rm eff}$ of M dwarf atmospheres.

We also assess how well our FGK-M training set recovers the reported abundances of the Souto et al. (2022) sample. Our training label set overlaps with the Souto et al. (2022) abundances for Fe, Mg, Al, C, O, and Ca. We again use the same implementation of *The Cannon* as in our LOOCV test and train on all labels in Equation 10. Our inferred versus reported abundances are presented in Figure 6. We do not show our 1-to-1 plot for $T_{\rm eff}$, but the inferred and reported values agree to within 48 K and exhibit a strong 1-to-1 trend. The abundances agree to within 0.1−0.17 dex, with good adherence to the 1-to-1 trends. However, our model does struggle to reproduce the low metallicities of Souto et al. (2022) stars outside the range of our training set ([Fe/H] < −0.56 dex), namely the lowest metallicity Souto et al. (2022) M dwarf ([Fe/H] = −0.92 dex). We would ideally exclude such low metallicity sources from our analysis of SDSS-V/MWM M dwarfs, but we cannot cut on the AS-PCAP [Fe/H] values because they are unreliable. The lowest metallicity Souto et al. (2022) star also does not have a high model fit χ^2 value, or noticeably poor inferences for other labels (e.g., $T_{\rm eff}$) that could compensate to produce a good model fit. We conclude that there is no clear way to identify M dwarfs with metallicities outside the range of our training set parameter space beforehand. However, we do not expect many SDSS-V/MWM M dwarfs to have metallicities outside −0.56 < [Fe/H] < 0.31 dex. Because M dwarfs are faint, those observed will be mostly within the solar neighborhood, and the majority of solar neighborhood main sequence dwarfs are within this metallicity threshold (e.g., Bensby et al. 2014). Still, we highlight this to caution the reader, and encourage users of our M dwarf model to inspect the flux model fits in cases where the M dwarf metallicity is

Figure 5. Plots of flux model fits to the spectrum of an example M dwarf from the Hyades cluster (SDSS ID = 77382206). The left panel displays the fit resulting from training on the Souto et al. (2022) sample, and the right panel displays the fit from training on the FGK-M sample. In both panels the flux model fit is in red, and the real spectrum is in black. Prominent OH lines used for calculating [O/H] or A(O) are marked by the dashed gray lines. It is apparent that the model fit resulting from the FGK-M training set is a better fit to the M dwarf spectrum (right panel).

suspected to be outside this range. Overall though, we reproduce the reported Souto et al. (2022) abundances and Hyades [M/H] and A(O) abundances from Wanderley et al. (2023) well, demonstrating that our M dwarf model trained on our sample of FGK-M binaries is robust.

6. SDSS-V/MWM TEST SAMPLE

We then use our FGK-M training set to infer detailed abundances for the largest sample of M dwarfs we can get from SDSS-V/MWM. Beginning with the entire SDSS-V/MWM catalog of ∼1 million stars, we select test set M dwarfs according to the type-color sequence relations from Pecaut & Mamajek (2013). We then remove sources with negative *Gaia* DR3 parallaxes, leaving us with ∼48,000 stars. Next, we implement additional cuts to select M dwarfs with high quality spectra, minimal binary contamination, and properties contained within the training set parameter space. We outline these cuts below:

1. We apply a cut of SNR $\geq$ 50 to ensure that the test set spectra are sufficiently high quality for inferring high fidelity labels with *The Cannon*.

2. To remove contamination from binaries, we use *Gaia* DR3 quantities to only retain M dwarfs that are well-fit by single star astrometric solutions. Specifically, we made cuts on the *Gaia* Renormalised Unit Weight Error of RUWE < 1.4, and set `non_single_star` = 0.

3. We remove M dwarfs outside of the training set $T_{\rm eff}$ boundaries with 3088 K < $T_{\rm eff}$ < 4085 K. As

mentioned in Section 5.1, we calculate $T_{\rm eff}$ values from an empirical color-temperature relation (Curtis et al. 2020), and use these photometric $T_{\rm eff}$ values for this cut.

4. We calculate log g values using an empirical relation based on *Gaia* DR3 astrometry and K-band magnitudes (Mann et al. 2019), and remove sources with log g outside of 4−5.5 dex, the typical range that M dwarfs span (e.g., Casagrande et al. 2008). Approximately 2% of our sample lack K-band magnitudes, so we do not cut on log g for these sources.

5. After running our *Cannon* model on the remaining M dwarfs in our test set following these cuts, we remove those with flux model spectral fit χ^2 values >100,000 (Birky et al. 2020).

6. We also remove stars that lack sensible Hessian matrices from fits with our *Cannon* model. This indicates that the log-likelihood space is flat around the critical values set by the best-inferred labels according our model, and as a consequence the model cannot identify the best direction to move in for optimizing the log-likelihood. In other words, the model fails to achieve a sensible fit.

Following this set of cuts, we are left with ∼17,000 M dwarfs. Their flux model spectral fit χ^2 values peak at ∼12,500 (reduced χ^2 of ∼1.7), indicating that they are well-fit by the flux models. We show example flux model fits to four randomly selected M dwarfs from our test set that span a wide range of temperatures and metallicities in Figure A.4.
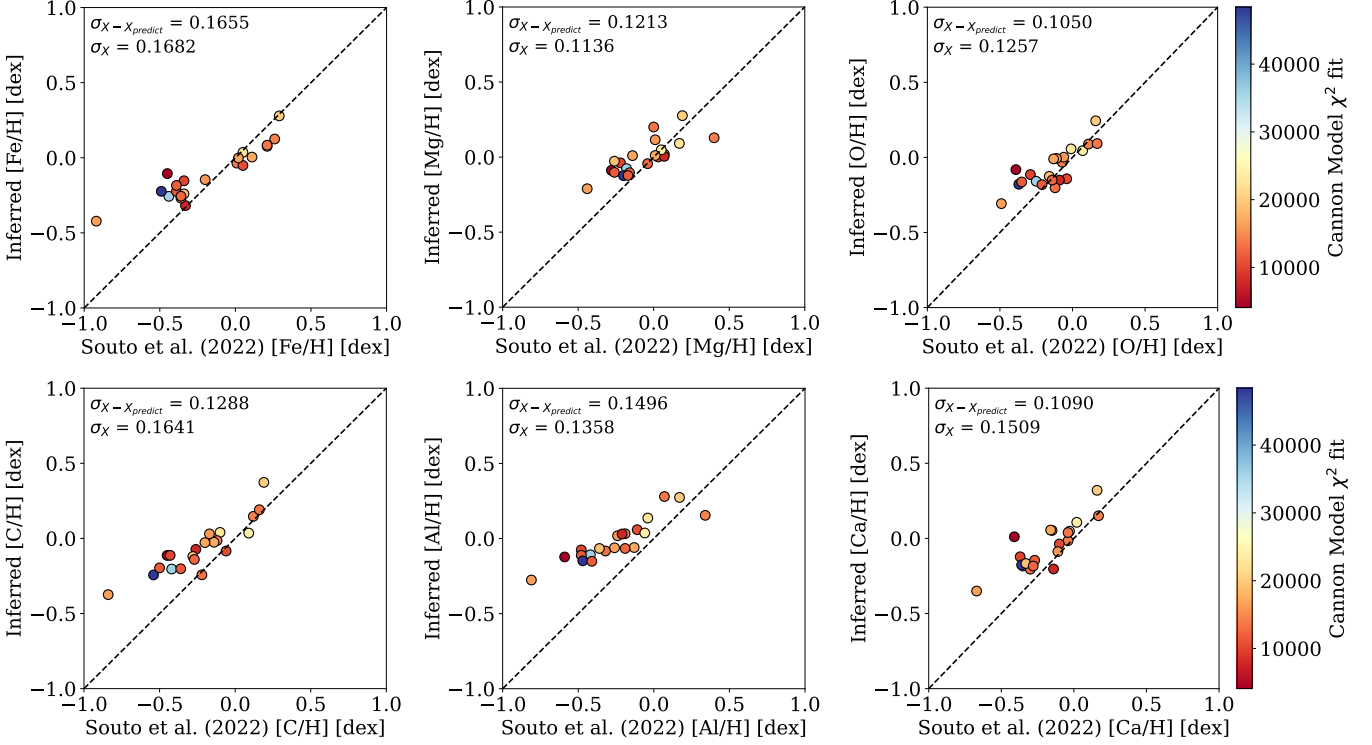
Figure 6. 1-to-1 plots of the inferred versus reported abundances for M dwarfs from Souto et al. (2022) using our FGK-M training set. The points are colored by the χ^2 of the flux model spectral fits. We report the rms scatter between the inferred and reported abundance values from Wanderley et al. (2023) in the top left of each plot, and the scatter of the reported values for that label below.


Figure 7. $T_{\rm eff}$ inferred from *The Cannon* vs. $T_{\rm eff}$ calculated from the Curtis et al. (2020) relation for all ~17,000 M dwarfs in our test set. The points are colored by the χ^2 of the flux model spectral fits. The two dashed lines on either side of the 1-to-1 line mark the 2-σ boundaries in $T_{\rm eff}$ agreement.

We also use our flux model fits to explicitly show that our inferred abundances contain unique chemical information for each element. To do this, we selected four M dwarfs (SDSS ID = 79024069 overlaps with the set in Figure A.4), of which two have very different inferred [Fe/H] vs. [Si/H], and the other two have very different

inferred [O/H] vs. [Si/H]. We compare our flux models of these four M dwarfs with alternative flux models generated from substituting [Si/H] with [Fe/H] or [O/H], depending on which is more discrepant. It can be seen that our chosen flux models (red) outperform the alternative flux models (blue) in all cases, as shown by their fits to prominent Si and OH spectral features (Figure A.5).

As another sanity check, we also compare the $T_{\rm eff}$ values we infer from our model with $T_{\rm eff}$ calculated from the empirical color-temperature relation derived in Curtis et al. (2020). The photometric $T_{\rm eff}$ calculated from the empirical relation are not affected by reddening. More specifically, we find that ~90% of the M dwarfs have color extinction values E(BP−RP) < 0.06, which corresponds to a maximum possible temperature shift of ~60 K. Thus, our test set M dwarf Curtis et al. (2020) relation $T_{\rm eff}$ values are not strongly affected by reddening, and can be used for this comparison test. These $T_{\rm eff}$ values and our inferred $T_{\rm eff}$ values exhibit a clear 1-to-1 relation, with 89% agreeing to within 100 K. In terms of significance, $T_{\rm eff}$ agreement to within 2-σ is at ~193 K, and includes the bulk of sources near the 1-to-1 line (Figure 7). The reason that some test set M dwarfs have large discrepancies between their $T_{\rm eff}$ values from

Figure 8. Kiel diagrams of the M dwarfs in our test set (log g from Mann et al. (2019) versus the inferred $T_{\rm eff}$ from our *Cannon* model), colored by their [Fe/H] from the ASPCAP pipeline (not those of any potential FGK companions) (left), and inferred [Fe/H] from our *Cannon* model (right).

the Curtis et al. (2020) relation versus from our model is likely because they fall outside the $T_{\rm eff}$ and abundance parameter space spanned by our training set, and are being labeled with inaccurate $T_{\rm eff}$ and abundance values that allow our model to achieve a good flux fit. We advise users of our catalog to treat such M dwarfs with caution, and include a flag (temp_agree) on those with Curtis et al. (2020) versus model $T_{\rm eff}$ values discrepant beyond 2-σ.

Using our $T_{\rm eff}$ inferred from *The Cannon* and log g calculated from the Mann et al. (2019) relation, we plot a Kiel diagram for our test set M dwarfs. The M dwarfs are colored by our inferred [Fe/H], which exhibit the evolutionary tracks we expect as dictated by varying stellar metallicities (e.g., Hejazi et al. 2022) (Figure 8, right panel). Our inferred [Fe/H] values also appear more reasonable than the ASPCAP [Fe/H] values, which are unrealistically metal-poor (>80% have ASPCAP [Fe/H] < 0 dex) (Figure 8, left panel). This indicates that our M dwarf abundances inferred from *The Cannon* are reliable, and an improvement on those from the current SDSS-V ASPCAP pipeline.

In order to quantify the inferred $T_{\rm eff}$ and abundance label uncertainties of the ~17,000 M dwarfs in our test set, we estimate the scatter between labels from repeat APOGEE observations of the same stars. We identify ~500 stars in our sample with two visit spectra from APOGEE, denoted A and B, and define a quantity $Z_{A,B}$ for each star and label:

$$Z_{A,B} = \frac{\ell_A - \ell_B}{\sqrt{\sigma_A^2 + \sigma_B^2 + 2(\sigma_{inflate}^2)}}, \tag{11}$$

where ℓ_A and ℓ_B are the inferred values for a particular label ($T_{\rm eff}$ or an elemental abundance) using spectrum A or B, and σ_A and σ_B are the associated scatter derived from our *Cannon* model covariance matrices. We fit for $\sigma_{inflate}$, which can be thought of as the additional scatter we need to account for the difference in inferred labels from different APOGEE visit spectra A and B. The $\sigma_{inflate}$ term is multiplied by two to account for the two visit spectra.

We fit for $\sigma_{inflate}$ factors so that the 13 distributions of $Z_{A,B}$ (corresponding to our 13 labels) from our sample of ~500 stars are normal distributions centered at 0, with standard deviations of 1. The resultant $\sigma_{inflate}$ values range from $0.016-0.025$ dex. To derive the final $T_{\rm eff}$ and abundance label uncertainties for each star, we add these $\sigma_{inflate}$ factors in quadrature with each star's scatter for that label from their individual covariance matrix. The median values of the resultant $T_{\rm eff}$ and abundance errors across all elements is 13 K and $0.018-0.029$ dex, respectively. We provide these uncertainties and all other properties of our final catalog containing 16,590 M dwarfs in Table 2.

7. CONCLUSIONS

Using *The Cannon*, we constructed a data-driven model for inferring M dwarf abundances across a wide set of elements. We apply our model to 16,590 M dwarfs in SDSS-V/MWM, and provide a catalog of their inferred $T_{\rm eff}$ and abundances. We anticipate that this catalog will be invaluable for star and planet formation investigations with SDSS-V data.

We note that our M dwarf model parameter space spans $-0.56 <$ [Fe/H] < 0.31 dex, and is not suitable for inferring abundances for M dwarfs outside this metallicity range. We do not expect that many M dwarfs

Table 2. M Dwarf Test Set Properties

Gaia DR3 ID	SDSS ID	$T_{\rm eff}$ K	[Fe/H] dex	[Mg/H] dex	[Al/H] dex	χ^2	TEFF_AGREE	...
421503353090268416	54381371	3815 ± 13.2	0.00 ± 0.02	0.03 ± 0.02	0.05 ± 0.02	8754	True	
421535342004177920	54381716	3830 ± 13.2	-0.06 ± 0.02	-0.05 ± 0.02	-0.10 ± 0.02	12584	True	
421917594095590272	54386478	3382 ± 13.2	0.04 ± 0.02	0.14 ± 0.02	0.17 ± 0.02	16320	True	
422026445746556928	54387723	3603 ± 13.2	0.16 ± 0.02	0.16 ± 0.02	0.21 ± 0.02	19318	True	
422148079220822912	54389138	3792 ± 13.2	-0.06 ± 0.02	-0.04 ± 0.02	-0.05 ± 0.02	59377	True	
422204253090333824	54389783	3280 ± 13.2	-0.03 ± 0.02	0.03 ± 0.02	0.03 ± 0.02	23258	True	
422228171769143552	54390069	3578 ± 13.2	0.04 ± 0.02	0.10 ± 0.03	0.11 ± 0.03	8691	True	
422491023770086528	54392812	3755 ± 13.2	-0.05 ± 0.02	-0.01 ± 0.02	-0.03 ± 0.02	59154	True	
422673886296537600	54394735	3674 ± 13.2	-0.19 ± 0.02	-0.10 ± 0.02	-0.08 ± 0.02	11333	True	
422794420257500160	54396157	3505 ± 13.2	0.08 ± 0.02	0.13 ± 0.02	0.17 ± 0.02	20607	True	
422813352474333952	54396372	3815 ± 13.2	-0.01 ± 0.02	0.03 ± 0.02	0.02 ± 0.02	21023	True	
423209863856063232	54400818	3861 ± 13.2	0.13 ± 0.02	0.13 ± 0.02	0.08 ± 0.02	16279	True	
423840051515323904	54408223	3816 ± 13.2	0.22 ± 0.02	0.20 ± 0.02	0.17 ± 0.02	26926	True	
424859165658264192	54421755	3755 ± 13.2	0.03 ± 0.02	0.12 ± 0.02	0.14 ± 0.02	11221	True	
424986674646826624	54423703	3345 ± 13.2	0.09 ± 0.02	0.12 ± 0.02	0.15 ± 0.02	10208	True	
425420638133142400	54429533	4119 ± 13.2	0.00 ± 0.02	0.00 ± 0.02	-0.24 ± 0.02	16614	False	
425851195719880704	54434698	3769 ± 13.2	0.05 ± 0.02	0.08 ± 0.02	0.11 ± 0.02	15270	True	
426520355931927936	54444470	3716 ± 13.2	-0.07 ± 0.02	-0.04 ± 0.02	-0.02 ± 0.02	8546	True	
426617864561108736	54446189	3803 ± 13.2	0.05 ± 0.02	0.03 ± 0.02	-0.00 ± 0.02	25069	True	
426622198194555648	54446259	3855 ± 13.2	0.07 ± 0.02	0.08 ± 0.02	0.05 ± 0.02	10447	True	
427277335326484096	54456095	3380 ± 13.2	0.17 ± 0.02	0.22 ± 0.02	0.29 ± 0.02	15818	True	
427880245651962880	54464028	3904 ± 13.2	0.02 ± 0.02	0.05 ± 0.02	0.02 ± 0.02	17876	True	
428015073265343872	54465640	3610 ± 13.2	-0.17 ± 0.02	-0.08 ± 0.02	-0.08 ± 0.02	17593	True	
428216833650287232	54467942	3509 ± 13.2	-0.01 ± 0.02	-0.00 ± 0.03	0.00 ± 0.03	11967	False	
428262703901618432	54468635	3713 ± 13.2	0.03 ± 0.02	0.05 ± 0.02	0.03 ± 0.02	19496	True	
428342933890361216	54469672	3944 ± 13.2	-0.09 ± 0.02	-0.01 ± 0.02	-0.06 ± 0.02	10746	True	
428347778613655424	54469716	3544 ± 13.2	0.05 ± 0.02	0.06 ± 0.02	0.09 ± 0.02	37802	True	
428489409455550336	54471368	3614 ± 13.2	0.02 ± 0.02	0.02 ± 0.02	0.04 ± 0.02	8269	True	
428538818759899648	54472086	3841 ± 13.2	0.13 ± 0.02	0.13 ± 0.02	0.10 ± 0.02	10299	True	
428548645644559104	54472241	3187 ± 13.2	0.10 ± 0.02	0.14 ± 0.02	0.17 ± 0.02	33632	True	

...

NOTE—This table lists the properties of the 16,590 M dwarfs in our SDSS-V/MWM test set. The properties consist of all labels inferred from our implementation of *The Cannon* ($T_{\rm eff}$, abundances for all elements of interest, and flux model spectral fit χ^2 values). The errors on the $T_{\rm eff}$ and abundance labels are the scatter in labels from resampling from flux errors 10 times for each M dwarf. We also include a TEFF_AGREE) flag describing whether the photometric and *Cannon*-inferred $T_{\rm eff}$ agree to within 2-σ. We only list a subset of the abundances in this table, but the full set is provided in the downloadable version.

(This table is available in its entirety in machine-readable form.)

in SDSS-V/MWM will be more metal-poor than -0.56 dex, or more metal-rich than 0.31 dex given that M dwarfs are often quite old. Still, users should be confident that their M dwarfs of interest fall within this metallicity range before using our catalog. In the future it may be possible to expand our model parameter space with a larger, more diverse FGK-M training set, but this will require identification of more FGK-M binaries in SDSS-V/MWM. For now, the metallicity range spanned by our catalog is about as wide as any other ex-

isting method for inferring M dwarf abundances in large stellar samples (e.g., Birky et al. 2020).

Our validation tests with the Wanderley et al. (2023) Hyades M dwarf abundances and the Souto et al. (2022) M dwarf sample indicate that our M dwarf model is robust. The inferred metallicities for our final sample of 16,590 M dwarfs also exhibit the evolutionary tracks we expect according to the stellar metallicities. As more M dwarfs are observed through SDSS-V/MWM, our model can be used to infer their detailed abundances as their spectra become available. Additionally, our model

may be compatible with M dwarf spectra collected with other instruments/surveys, but this may require additional processing, e.g., spectra from an instrument with higher resolution compared to APOGEE must be artificially degraded to match the SDSS-V/MWM training set. The spectra must also have a wavelength range encompassed by APOGEE, i.e., H-band. Expanding our model to work with M dwarf spectra from other instruments could be quite fruitful, but we leave these investigations for future studies.

REFERENCES

Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, ApJS, 249, 3, doi: 10.3847/1538-4365/ab929e

Allard, F., Hauschildt, P. H., Alexander, D. R., & Starrfield, S. 1997, ARA&A, 35, 137, doi: 10.1146/annurev.astro.35.1.137

Almeida, A., Anderson, S. F., Argudo-Fernández, M., et al. 2023, ApJS, 267, 44, doi: 10.3847/1538-4365/acda98

Alvarez, R., & Plez, B. 1998, A&A, 330, 1109, doi: 10.48550/arXiv.astro-ph/9710157

Angelo, I., Bedell, M., Petigura, E., & Ness, M. 2024, arXiv e-prints, arXiv:2407.19016, doi: 10.48550/arXiv.2407.19016

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Behmard, A., Petigura, E. A., & Howard, A. W. 2019, ApJ, 876, 68, doi: 10.3847/1538-4357/ab14e0

Bensby, T., Feltzing, S., & Oey, M. S. 2014, A&A, 562, A71, doi: 10.1051/0004-6361/201322631

Birky, J., Hogg, D. W., Mann, A. W., & Burgasser, A. 2020, ApJ, 892, 31, doi: 10.3847/1538-4357/ab7004

Bland-Hawthorn, J., Krumholz, M. R., & Freeman, K. 2010, ApJ, 713, 166, doi: 10.1088/0004-637X/713/1/166

Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, AJ, 154, 28, doi: 10.3847/1538-3881/aa7567

Bochanski, J. J., Hawley, S. L., Covey, K. R., et al. 2010, AJ, 139, 2679, doi: 10.1088/0004-6256/139/6/2679

Brewer, J. M., & Fischer, D. A. 2018, ApJS, 237, 38, doi: 10.3847/1538-4365/aad501

Casagrande, L., Flynn, C., & Bessell, M. 2008, MNRAS, 389, 585, doi: 10.1111/j.1365-2966.2008.13573.x

Casey, A. R., Hogg, D. W., Ness, M., et al. 2016, ArXiv e-prints. https://arxiv.org/abs/1603.03040

Choi, J., Dotter, A., Conroy, C., et al. 2016, ApJ, 823, 102, doi: 10.3847/0004-637X/823/2/102

Clampin, M. 2008, Advances in Space Research, 41, 1983, doi: 10.1016/j.asr.2008.01.010

Cunha, K., Smith, V. V., Hasselquist, S., et al. 2017, ApJ, 844, 145, doi: 10.3847/1538-4357/aa7beb

Curtis, J. L., Agüeros, M. A., Matt, S. P., et al. 2020, ApJ, 904, 140, doi: 10.3847/1538-4357/abbf58

Danielski, C., Babusiaux, C., Ruiz-Dern, L., Sartoretti, P., & Arenou, F. 2018, A&A, 614, A19, doi: 10.1051/0004-6361/201732327

De Silva, G. M., Freeman, K. C., Asplund, M., et al. 2007, AJ, 133, 1161, doi: 10.1086/511182

De Silva, G. M., Freeman, K. C., & Bland-Hawthorn, J. 2009, PASA, 26, 11, doi: 10.1071/AS08019

Dotter, A., Conroy, C., Cargile, P., & Asplund, M. 2017, ApJ, 840, 99, doi: 10.3847/1538-4357/aa6d10

Duque-Arribas, C., Tabernero, H. M., Montes, D., & Caballero, J. A. 2024, MNRAS, 528, 3028, doi: 10.1093/mnras/stae076

El-Badry, K., Rix, H.-W., & Heintz, T. M. 2021, MNRAS, 506, 2269, doi: 10.1093/mnras/stab323

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1, doi: 10.1051/0004-6361/202243940

Galgano, B., Stassun, K., & Rojas-Ayala, B. 2020, The Astronomical Journal, 159, 193, doi: 10.3847/1538-3881/ab7f37

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, AJ, 151, 144, doi: 10.3847/0004-6256/151/6/144

Green, G. M. 2018, The Journal of Open Source Software, 3, 695, doi: 10.21105/joss.00695

Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, ApJ, 887, 93, doi: 10.3847/1538-4357/ab5362

Grilo, V., Souto, D., Cunha, K., et al. 2024, MNRAS, doi: 10.1093/mnras/stae2209

Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, AJ, 131, 2332, doi: 10.1086/500975

Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, A&A, 486, 951, doi: 10.1051/0004-6361:200809724

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2

Hasselquist, S., Shetrone, M., Cunha, K., et al. 2016, ApJ, 833, 81, doi: 10.3847/1538-4357/833/1/81

Hayes, C. R., Masseron, T., Sobeck, J., et al. 2022, ApJS, 262, 34, doi: 10.3847/1538-4365/ac839f

Hejazi, N., Lépine, S., & Nordlander, T. 2022, ApJ, 927, 122, doi: 10.3847/1538-4357/ac4e16

Hejazi, N., Crossfield, I. J. M., Souto, D., et al. 2024, ApJ, 973, 31, doi: 10.3847/1538-4357/ad61dc

Henry, T. J., Jao, W.-C., Winters, J. G., et al. 2018, AJ, 155, 265, doi: 10.3847/1538-3881/aac262

Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, ApJ, 836, 5, doi: 10.3847/1538-4357/836/1/5

Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90, doi: 10.1109/MCSE.2007.55

Jönsson, H., Holtzman, J. A., Allende Prieto, C., et al. 2020, AJ, 160, 120, doi: 10.3847/1538-3881/aba592

Liu, F., Asplund, M., Yong, D., et al. 2019, A&A, 627, A117, doi: 10.1051/0004-6361/201935306

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94, doi: 10.3847/1538-3881/aa784d

Maldonado, J., Micela, G., Baratella, M., et al. 2020, A&A, 644, A68, doi: 10.1051/0004-6361/202039478

Mann, A. W., Brewer, J. M., Gaidos, E., Lépine, S., & Hilton, E. J. 2013a, AJ, 145, 52, doi: 10.1088/0004-6256/145/2/52

Mann, A. W., Deacon, N. R., Gaidos, E., et al. 2014, AJ, 147, 160, doi: 10.1088/0004-6256/147/6/160

Mann, A. W., Gaidos, E., & Ansdell, M. 2013b, ApJ, 779, 188, doi: 10.1088/0004-637X/779/2/188

Mann, A. W., Dupuy, T., Kraus, A. L., et al. 2019, ApJ, 871, 63, doi: 10.3847/1538-4357/aaf3bc

McDonough, W. F. 2003, Treatise on Geochemistry, 2, 568, doi: 10.1016/B0-08-043751-6/02015-6

Melo, E., Souto, D., Cunha, K., et al. 2024, ApJ, 973, 90, doi: 10.3847/1538-4357/ad5004

Miller, G. E., & Scalo, J. M. 1979, ApJS, 41, 513, doi: 10.1086/190629

Moedas, N., Deal, M., Bossini, D., & Campilho, B. 2022, A&A, 666, A43, doi: 10.1051/0004-6361/202243210

Muirhead, P. S., Dressing, C. D., Mann, A. W., et al. 2018, AJ, 155, 180, doi: 10.3847/1538-3881/aab710

Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015, ApJ, 808, 16, doi: 10.1088/0004-637X/808/1/16

Newton, E. R., Charbonneau, D., Irwin, J., et al. 2014, AJ, 147, 20, doi: 10.1088/0004-6256/147/1/20

Nidever, D. L., Holtzman, J. A., Allende Prieto, C., et al. 2015, AJ, 150, 173, doi: 10.1088/0004-6256/150/6/173

Nutzman, P., & Charbonneau, D. 2008, PASP, 120, 317, doi: 10.1086/533420

Pecaut, M. J., & Mamajek, E. E. 2013, ApJS, 208, 9, doi: 10.1088/0067-0049/208/1/9

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825

Plez, B. 2012, Turbospectrum: Code for spectral synthesis, Astrophysics Source Code Library, record ascl:1205.004

Rains, A. D., Nordlander, T., Monty, S., et al. 2024, MNRAS, 529, 3171, doi: 10.1093/mnras/stae560

Rajpurohit, A. S., Allard, F., Rajpurohit, S., et al. 2018, A&A, 620, A180, doi: 10.1051/0004-6361/201833500

Rampalli, R., Ness, M., & Wylie, S. 2021, ApJ, 921, 78, doi: 10.3847/1538-4357/ac1ac8

Rampalli, R., Ness, M. K., Edwards, G. H., Newton, E. R., & Bedell, M. 2024, ApJ, 965, 176, doi: 10.3847/1538-4357/ad303e

Rauer, H., Aerts, C., Cabrera, J., et al. 2024, arXiv e-prints, arXiv:2406.05447, doi: 10.48550/arXiv.2406.05447

Rice, M., & Brewer, J. M. 2020, ApJ, 898, 119, doi: 10.3847/1538-4357/ab9f96

Rojas-Ayala, B., Covey, K. R., Muirhead, P. S., & Lloyd, J. P. 2012, ApJ, 748, 93, doi: 10.1088/0004-637X/748/2/93

Shan, Y., Reiners, A., Fabbian, D., et al. 2021, A&A, 654, A118, doi: 10.1051/0004-6361/202141530

Smith, V. V., Cunha, K., Shetrone, M. D., et al. 2013, ApJ, 765, 16, doi: 10.1088/0004-637X/765/1/16

Smith, V. V., Bizyaev, D., Cunha, K., et al. 2021, AJ, 161, 254, doi: 10.3847/1538-3881/abefdc

Souto, D., Cunha, K., García-Hernández, D. A., et al. 2017, ApJ, 835, 239, doi: 10.3847/1538-4357/835/2/239

Souto, D., Allende Prieto, C., Cunha, K., et al. 2019, ApJ, 874, 97, doi: 10.3847/1538-4357/ab0b43

Souto, D., Cunha, K., Smith, V. V., et al. 2020, ApJ, 890, 133, doi: 10.3847/1538-4357/ab6d07

—. 2022, ApJ, 927, 123, doi: 10.3847/1538-4357/ac4891

Suárez Mascareño, A., Rebolo, R., & González Hernández, J. I. 2016, A&A, 595, A12, doi: 10.1051/0004-6361/201628586

Trifonov, T., Kürster, M., Zechmeister, M., et al. 2018, A&A, 609, A117, doi: 10.1051/0004-6361/201731442

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: 10.1038/s41592-019-0686-2

Wanderley, F., Cunha, K., Souto, D., et al. 2023, ApJ, 951, 90, doi: 10.3847/1538-4357/acd4bd

Wes McKinney. 2010, in Proceedings of the 9th Python in Science Conference, ed. Stéfan van der Walt & Jarrod Millman, 56 – 61, doi: 10.25080/Majora-92bf1922-00a

Wilson, J. C., Hearty, F. R., Skrutskie, M. F., et al. 2019, PASP, 131, 055001, doi: 10.1088/1538-3873/ab0075

Woolf, V. M., & Wallerstein, G. 2020, MNRAS, 494, 2718, doi: 10.1093/mnras/staa878

Woolf, V. M., & West, A. A. 2012, MNRAS, 422, 1489, doi: 10.1111/j.1365-2966.2012.20722.x

APPENDIX

Figure A.1. The scatter and normalized flux coefficients for the first four labels of our *Cannon* model trained on FGK-M binaries. The wavelength pixels with large coefficient amplitudes correspond to locations in the spectra that contain the most information for that particular label. We highlight a few prominent absorption features whose locations match high amplitude pixels. We also highlight two wavelength pixels with poor sky subtraction in gray, which should be disregarded (these wavelength pixels are assigned high intrinsic model scatter s_j^2, so they do not affect the results). For ease of viewing, we only show the wavelength range covered by the first APOGEE chip.

**Figure A.2.** The scatter and normalized flux coefficients for the next four labels of our *Cannon* model.

Figure A.3. The scatter and normalized flux coefficients for the final five labels of our *Cannon* model.

Figure A.4. Plots of flux model fits from our *Cannon* model to four randomly selected M dwarfs from our test set that span a wide range of temperatures and metallicities. We show the wavelength range of the middle APOGEE chip. The model fits are in red, and the real spectra are in black. The SDSS IDs of each M dwarf are provided in the top right corner of each panel.
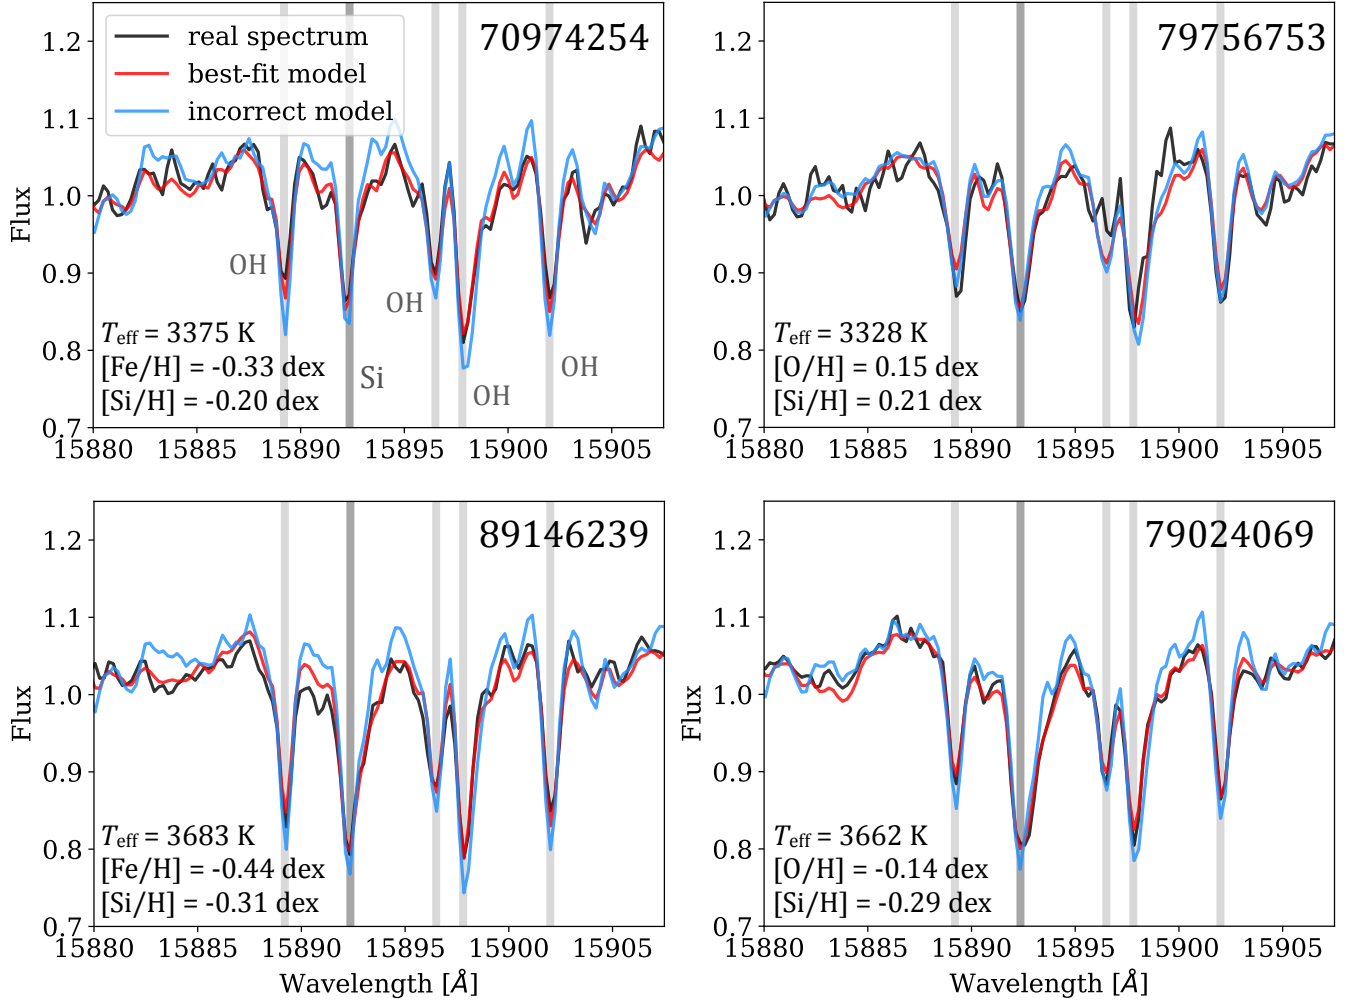
Figure A.5. Plots of our flux model fits to four M dwarfs in our test set (SDSS IDs provided in the top right corner of each panel), of which two have very different inferred [Si/H] and [Fe/H] values (left column), and the remaining two have very different interred [Si/H] and [O/H] values (right column). The spectra (black) are compared to our chosen flux models (red), and alternative flux models where we substitute [Si/H] with either [Fe/H] or [O/H] depending on which is most discrepant. Our chosen models outperform these alternative models for all M dwarfs, especially in the region we highlight which contains prominent Si (dark gray line), and OH (lighter gray lines) absorption features.