

Measuring the 8621 Å Diffuse Interstellar Band in Gaia DR3 RVS Spectra: Obtaining a Clean Catalog by Marginalizing over Stellar Types

Andrew K. Saydjari^{1,2} , Ana Sofía M. Uzsoy² , Catherine Zucker^{2,3,5} , J. E. G. Peek^{3,4} , and Douglas P. Finkbeiner^{1,2} 

¹ Department of Physics, Harvard University, 17 Oxford St., Cambridge, MA 02138, USA; andrew.saydjari@cfa.harvard.edu

² Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA

³ Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218, USA

⁴ Department of Physics & Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

Received 2022 December 7; revised 2023 April 26; accepted 2023 April 28; published 2023 August 31

Abstract

Diffuse interstellar bands (DIBs) are broad absorption features associated with interstellar dust and can serve as chemical and kinematic tracers. Conventional measurements of DIBs in stellar spectra are complicated by residuals between observations and best-fit stellar models. To overcome this, we simultaneously model the spectrum as a combination of stellar, dust, and residual components, with full posteriors on the joint distribution of the components. This decomposition is obtained by modeling each component as a draw from a high-dimensional Gaussian distribution in the data space (the observed spectrum)—a method we call “Marginalized Analytic Data-space Gaussian Inference for Component Separation” (MADGICS). We use a data-driven prior for the stellar component, which avoids missing stellar features not well modeled by synthetic spectra. This technique provides statistically rigorous uncertainties and detection thresholds, which are required to work in the low signal-to-noise regime that is commonplace for dusty lines of sight. We reprocess all public Gaia DR3 RVS spectra and present an improved 8621 Å DIB catalog, free of detectable stellar line contamination. We constrain the rest-frame wavelength to 8623.14 ± 0.087 Å (vacuum), find no significant evidence for DIBs in the Local Bubble from the 1/6th of RVS spectra that are public, and show unprecedented correlation with kinematic substructure in Galactic CO maps. We validate the catalog, its reported uncertainties, and biases using synthetic injection tests. We believe MADGICS provides a viable path forward for large-scale spectral line measurements in the presence of complex spectral contamination.

Unified Astronomy Thesaurus concepts: [Diffuse interstellar bands \(379\)](#); [Astronomy data reduction \(1861\)](#); [Catalogs \(205\)](#)

Supporting material: [interactive figure](#)

1. Introduction

Diffuse interstellar bands (DIBs) are broad absorption features associated with interstellar dust and can be used as chemical tracers, though few DIB carriers have been conclusively identified (Campbell et al. 2015). Despite this, the moments of a simple Gaussian model of the DIB line shape contain a wealth of information. Comparisons of the zeroth moment (the integral under the curve, often called the equivalent width, EW_{DIB}) with dust extinction carries information about variations in the interstellar medium (ISM) environment which may enhance or deplete DIB carriers (Lan et al. 2015). The first moment (the detection wavelength, λ_{DIB}), can be converted into a velocity ($v_{\text{DIB}}^{\text{LSR}}$) given a known rest-frame wavelength and used as a kinematic tracer of the ISM. Previous work used the 15273 Å near-infrared (NIR) DIB detected with APOGEE (Majewski et al. 2017) to map the ISM velocity field and measure large-scale Galactic rotation (Zasowski et al. 2015) and spiral substructure (Tchernyshyov et al. 2018). The second moment (the width of the Gaussian, σ_{DIB}), contains information about the velocity distribution of

dust along the line of sight, the instrumental line spread function, and the intrinsic lifetime of the DIB carrier, useful for carrier identification (Campbell et al. 2015). While coarse estimates of σ_{DIB} have been used in the past, careful reporting of single-detection σ_{DIB} and corresponding uncertainties remains a challenge.

Early identification of DIBs focused on taking high S/N (signal-to-noise ratio), and sometimes high-resolution, spectra of a handful of stars along lines of sight with high extinction (van Loon et al. 2013). Features that correlated with dust extinction and were not present in low-extinction reference stars were identified as DIBs. This often limited those studies to hot stars with relatively few stellar features in order to minimize confusion between the DIBs and stellar features. More recently, larger population studies in the Sloan Digital Sky Survey (SDSS, York et al. 2000) and the Radial Velocity Experiment (RAVE, Steinmetz et al. 2006) have utilized stacking in order to improve S/N enough for statistically significant detections (Kos & Zwitter 2013; Lan et al. 2015). While stacking can enable the detection of fainter DIBs, the information contained in the moments of the Gaussian profile of the DIBs is significantly reduced in that they are averaged over all of the combined spectra.

Now, sufficient data quality and quantities are available to enable large-scale single-detection DIB catalogs. In producing a catalog, we produce a low-dimensional representation of the detected DIBs and their properties, generally by listing their zeroth, first, and second moments along with the stars toward

⁵ Hubble Fellow.

which they were detected. This has been done in the NIR for the 15273 Å DIB in APOGEE (Zasowski et al. 2015), which searched over ~ 100 k spectra, and in the optical for the 8621 Å DIB in Gaia RVS spectra (Schultheis et al. 2023), which searched over ~ 5.5 M spectra. Yet, statistical methods for modeling and detecting DIBs have not kept pace with the increase in data quantity and quality.

DIB catalogs are still made by obtaining a best-fit stellar spectrum, dividing the observed spectrum by the stellar model, and fitting a Gaussian peak to the residuals centered near the DIB wavelength (Zasowski et al. 2015; Schultheis et al. 2023). However, stellar line lists and models often have significant residuals relative to observations, which can cause confusion when fitting or finding DIB features in the residuals. Further, such fits are susceptible to biases resulting from variability in the continuum normalization and are often limited to the high stellar S/N regime, despite the fact that dusty lines of sight are more likely to have low stellar S/N due to dust extinction.

In this work, we introduce a new method for detecting and modeling DIBs that jointly models the stellar and DIB components of the spectra, marginalizing over stellar types and uncertainties in stellar lines. This method is termed “Marginalized Analytic Data-space Gaussian Inference for Component Separation” (MADGICS). We reprocess all public Gaia DR3 RVS spectra with MADGICS and demonstrate reduced stellar line contamination in the new MADGICS DIB catalog presented herein. We believe that the speed and stability of the method, even at low stellar S/N, establishes MADGICS as an important statistical tool in the era of large spectroscopic (DIB) surveys. We anticipate the applicability of this method also extends to atomic interstellar absorption (Welty & Hobbs 2001) and emission (Kollmeier et al. 2017) line surveys.

We begin by comparing the 8621 Å MADGICS DIB catalog to the Gaia DIB catalog and presenting astrophysical applications and validations of the catalog (Section 3). Then, we describe the method, catalog construction, and further validation tests (Sections 4–5).

2. Data

The primary data used in this work are the Gaia DR3 RVS spectra, which were obtained from the Gaia archive (Gaia Collaboration et al. 2016; Vallenari et al. 2023). These are optical spectra spanning 8460–8700 Å which are shifted into the stellar rest frame, averaged over all visits, and interpolated onto a common linear (0.1 Å spacing) wavelength grid (2401 bins). The wavelength calibration of the spectra is reported in vacuum.⁶ Both the flux and flux uncertainty per wavelength bin are provided by Gaia. See Section 5.2 of Vallenari et al. (2023) for more details.

The primary Gaia analysis of these spectra falls under the General Stellar Parameteriser-spectroscopy (GSP-Spec) module, which is a purely spectroscopic analysis to obtain chemophysical stellar parameters from combined RVS spectra (Recio-Blanco et al. 2022). Within GSP-Spec, DIB detections and parameter estimates were made predominantly by dividing each RVS spectrum by the best-matching synthetic spectrum

⁶ Throughout this work, we work exclusively in vacuum wavelength, except in referring to the catalog and DIB by the common 8621 Å nomenclature reflective of the historical identification of DIBs by wavelength in air.

(using the GSP-Spec atmospheric parameters) and performing a least-squares fit of a Gaussian function (Schultheis et al. 2023).

Of the 5,591,594 RVS spectra processed by GSP-Spec, only 999,645 are publicly available. Based on the full 5,591,594 spectra, the Gaia DIB catalog has entries for 476,532 spectra, of which 50,787 have public spectra (Table 1). The Gaia DIB team strongly suggests restricting this catalog to a “high-quality” (HQ) sample, which includes 137,536 spectra, of which 9763 are public (Schultheis et al. 2023). While the Gaia DIB catalog entries are available for all stars—even those without public spectra—we will restrict our comparisons with the Gaia DIB catalog to stars with public spectra, unless otherwise specified. The Gaia DIB catalog, and other Gaia-derived properties of the stars or spectral processing were obtained from the DR3 “source” and “astrophysical parameters” tables.

The cuts defining the Gaia HQ sample are summarized in Table 1. We divide these cuts into two classes: cuts on (1) S/N and (2) the overall model goodness of fit for the spectrum. The entire Gaia 8621 Å DIB catalog is limited to stellar S/N > 70 . A further S/N cut on the DIB detection is imposed by requiring $\text{EW}_{\text{DIB}}/\sigma(\text{EW}_{\text{DIB}}) > 2.86$. The efficacy of cuts imposed this way of course depends strongly on the quality of the $\sigma(\text{EW}_{\text{DIB}})$ estimate.

Additional cuts (Schultheis et al. 2023) are imposed to remove low-quality detections, often as a result of poor stellar modeling. The GSP-Spec flags 1–13, which indicate the goodness of fit for the stellar modeling, are required to be < 2 (reasonably good). Cuts on the DIB peak fit parameters limit λ_{DIB} and σ_{DIB} to reasonable ranges, 8620–8626 Å and 0.6–3.2 Å respectively. The amplitude of the DIB peak, α_{DIB} , is required to be larger than a residual threshold R_i , which has a constant threshold of at least 0.15. The residual threshold R_i depends on σ_{DIB} and is $\max(R_A, R_B, 0.15)$ for $0.6 < \sigma_{\text{DIB}} < 1.2$ and $\max(R_B, 0.15)$ for $1.2 < \sigma_{\text{DIB}} < 3.2$. Here R_A denotes the standard deviation of the “global” data-model residuals 8605–8640 Å and R_B denotes the standard deviation of the “local” residuals “within the DIB profile,” neither of which are publicly available (Schultheis et al. 2023).

The selection function of public RVS spectra remains forthcoming (G. Seabroke et al. 2022, in preparation), but is nonuniform across the sky and generally concentrated within 1–2 kpc of the Sun (see Appendix A). There is also significant heterogeneity in the normalization of the RVS spectra (see Appendix B) as a result of a discrete choice in the Gaia pipeline to either pseudocontinuum normalize or rescale the spectra by a constant, as in the case of cool or low S/N stars. (Vallenari et al. 2023).

3. Astrophysical Catalog Validation

3.1. Stellar Frame

We first compare the MADGICS and Gaia HQ 8621 Å DIB catalog in terms of the first and second moments of the detected peak—the center wavelength in the stellar frame λ_{DIB} and its width σ_{DIB} (Figure 1). The advantage of plotting the catalog in the stellar frame is that it readily identifies the impact of stellar features on the DIB identification. In both panels, we provide guidelines corresponding to the:

1. putative DIB rest-frame wavelength $\lambda_{\text{DIB}_{\text{rest}}}$ (black),
2. calibrated stellar lines in the GSP-Spec line list (turquoise, Contursi et al. 2021a, 2021b),

3. and stellar lines not included or not well calibrated in GSP-Spec, but coincident with features of stellar origin (see also Section 4.2; magenta).

For the lines not well modeled by GSP-Spec, we do not claim or validate the assignment of the observed stellar features as a result of these lines. We simply plot the wavelength and identification of the nearest transition with a large contribution to the solar spectrum from the Kurucz archive⁷ (see Section 4.2.1 and Appendix G). We focus on demonstrating the impact of stellar lines on DIB catalogs rather than an identification of the stellar lines themselves.

In the right panel, the Gaia HQ catalog is shown within the central gray box, which denotes the $(\lambda_{\text{DIB}}, \sigma_{\text{DIB}})$ cuts imposed in defining the Gaia HQ sample (see Table 1). Outside of that box, we show DIB detections from the broader Gaia DIB catalog with only the stellar S/N, DIB S/N, and GSP-Spec flag cuts imposed. Even inside the HQ catalog limits (gray box), the influence of stellar features is apparent. The largest density of detections peaks at the edge of the σ_{DIB} cut (0.6 Å) and is centered on the 8625.11 Å Fe I line, which is not in the GSP-Spec line list (see Appendix G). This peak in density extends to larger σ_{DIB} until ~ 1.5 Å when it broadens and shifts slowly to shorter wavelengths. While not clearly associated with a specific stellar line above $\sigma_{\text{DIB}} = 1.5$ Å the correlation between σ_{DIB} and λ_{DIB} suggests a non-negligible impact of the stellar lines on the properties of the DIB detection. There is also a second peak at low σ_{DIB} , which appears to straddle the 8620.51 Å Fe I]⁸ 8620.79 Å Ti I, 8621.49 Å Fe I] and 8621.73 Å CN lines.

These correlations continue and become more prominent outside the parameter cuts (gray box) included in the definition of the HQ sample. An additional peak in the detection density appears near the wavelength of the 8617.67 Å Fe I line, which is also not in the GSP-Spec line list (see Appendix G). The broad 8620–8622 Å peak in density splits at lower σ_{DIB} into two modes centered on either the 8620.51 Å Fe I] and 8620.79 Å Ti I lines or the 8621.49 Å Fe I] and 8621.73 Å CN lines. Peaks in density near other GSP-Spec stellar lines appear to a lesser extent.

To contextualize the impact of these wavelength pileups on kinematic DIB studies, the right y-axis shows the radial velocity corresponding to the wavelength shift relative to the DIB center wavelength. This is not indicative of the true DIB radial velocity, which must account for the stellar radial velocity (see Section 3.4.1).

This comparison with the less restrictive sample suggests that the Gaia goodness of fit cuts on $(\alpha_{\text{DIB}}, \lambda_{\text{DIB}}, \sigma_{\text{DIB}})$ are not fully effective at removing the influence of stellar lines on the DIB catalog, even within the HQ sample. This list of stellar features that are either poorly modeled by or missing from GSP-Spec and impact the Gaia 8621 Å DIB catalog is not meant to be exhaustive, but merely illustrative that detections of unmodeled stellar features dominate the catalog, with a small contribution from even the carefully calibrated stellar features known in GSP-Spec as well (see Appendix G). A comparison of both the MADGICS and Gaia catalogs removing their respective goodness of fit cuts is provided in Appendix C. In general, we aim to (1) impose goodness of fit cuts that eliminate any influence of stellar lines and (2) use methods that are maximally robust to stellar mismodeling so that real DIBs

Table 1
8621 Å DIB Catalogs

Cuts	MADGICS	Gaia HQ	Gaia Full
S/N Cuts			
Stellar S/N	>15	>70	>70
DIB S/N	>3.8	>2.86	
Overall Spectrum Goodness of Fit Cuts			
$\chi^2_{\text{tot}}/\text{dof}$	0.71 – 1.41		
GSP-Spec Flags		<2	
α_{DIB}		$\alpha_{\text{DIB}} > R_i$	
λ_{DIB}		8620 – 8626 Å	
σ_{DIB}		0.6 – 3.2 Å	
Number of Spectra with Detections			
Public (999,645)	7789	9763	50,787
Total (5,591,594)		137,536	476,532

can be detected and modeled well in the largest number of spectra possible. These quality cuts for the MADGICS DIB catalog are detailed in Table 1.

In the left panel, the MADGICS catalog morphology is dominated by a Gaussian centered at $(\lambda_{\text{DIB}}, \sigma_{\text{DIB}}) = (8623.37 \text{ Å}, 1.9 \text{ Å})$. It has a slightly heavy tail toward larger σ_{DIB} , which can be understood in terms of a low DIB S/N detection bias (see Section 5.3). The λ_{DIB} dispersion is 1.3 Å. Assuming all of the λ_{DIB} dispersion is the result of the relative star-dust velocity dispersion, this λ_{DIB} dispersion corresponds to a velocity dispersion of 46 km s^{-1} . This velocity dispersion is consistent with a value of 49 km s^{-1} , which is obtained by adding in quadrature the dispersion of radial velocities for the stars in the sample measured by Gaia (35 km s^{-1}) and the typical velocity dispersion of CO gas ($\sim 35 \text{ km s}^{-1}$, as a proxy for the dust velocity; Dame et al. 2001). We explicitly note there is no variation in the density of detections or the σ_{DIB} as the λ_{DIB} approaches any of the stellar features indicated.

Thus, the MADGICS catalog delivers a distribution of DIB detections, free of detectable contamination by stellar lines, with a catalog defined only by simple S/N-based detection thresholds and a single χ^2 cut representing the overall goodness of fit. This constitutes a clear improvement over conventional methods of least-squares fitting DIB line shape functions to the data divided by a best-fit synthetic stellar model.

3.2. Local Bubble

While DIBs are known to be correlated with dust extinction, their carriers are suspected to be similar to the only clearly identified DIB carrier, the buckminsterfullerene cation (C_{60}^+ ; Edwards & Leach 1993; Campbell et al. 2015; Walker et al. 2015). Thus, DIBs should be detected when large molecules are present, not dust grains themselves, and the correlation is likely due to coupled formation mechanisms. A recent body of literature has claimed detections of DIBs within the Local Bubble, despite the low dust extinction therein (Bailey et al. 2016; Farhang et al. 2019; Schultheis et al. 2023). The Gaia HQ sample supports these claims by reporting confident detections, albeit of low EW, within the Local Bubble.⁹

⁷ <http://kurucz.harvard.edu/>

⁸ The bracket is a spectroscopic notation indicating that the transition is an intercombination line, meaning it is semiforbidden.

⁹ A large number of such detections are included in the public subset of Gaia RVS spectra.

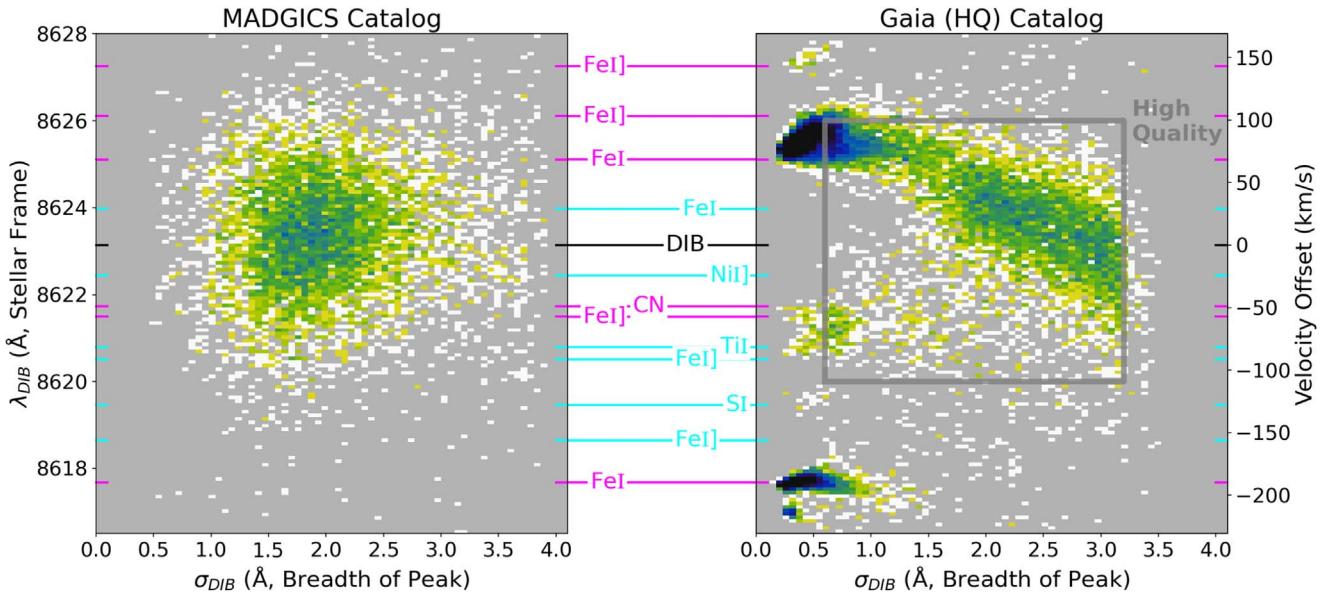


Figure 1. The density of detections in the MADGICS (this work, left) and Gaia (Schultheis et al. 2023, right) 8621 Å DIB catalog as a function of the DIB width (σ_{DIB}) and its central wavelength (λ_{DIB}). The color scale indicates density, where darker colors indicate more detections, and is a common logarithmic stretch for both catalogs. The (λ_{DIB} , σ_{DIB}) cuts imposed in defining the Gaia HQ sample are shown as a gray box, inside of which the HQ sample is shown. Outside of that box, the Gaia catalog is shown with only the stellar S/N, DIB S/N, and GSP-Spec flag cuts imposed. Guidelines for the central wavelength of the DIB and stellar lines are labeled and color coded between the plots. Stellar lines in magenta indicate lines not included or not well calibrated in the GSP-Spec line list, and the Gaia DIB detections preferentially cluster at those wavelengths, especially at low σ_{DIB} . Because the MADGICS catalog inference marginalizes over a data-driven stellar model, it is free from detectable contamination from these stellar lines and provides a more robust characterization of the DIB properties.

Within the MADGICS DIB catalog, there are no confident DIB detections (DIB S/N > 5) within the Local Bubble (Figure 2). All confident DIB detections fall outside the bubble and cluster along lines of sight passing through known local molecular clouds such as Taurus, Perseus, and Ophiuchus (Zucker et al. 2022). In projection, three significant detections appear to fall inside the bubble, but do not when examined in 3D. The 3D distribution can be visualized with the interactive version of Figure 2.¹⁰ This is an important demonstration of the necessity of marginalizing over possible stellar models when reporting the confidence of DIB detections. Using a single fixed stellar model per star can easily lead to stellar residuals that can confuse conventional DIB fitting approaches because there is confidently a feature in the stellar residuals, but that feature is of stellar origin. The MADGICS DIB S/N is instead based upon the $\Delta\chi^2$ resulting from adding a DIB component to the component separation, using an empirical model to marginalize over stellar features (see Section 4).

Because the MADGICS stellar model is empirical (see Section 4.2), one might worry that, if they exist, low EW DIBs uncorrelated with dust are baked into the stellar model. While possible, we show in Section 4.2 no features resembling the DIB in the stellar component prior. Further, upon loosening the cut on training stars used in constructing the stellar prior (allowing dustier stars into the training sample), we find only slightly reduced sensitivity to DIB detection (see Appendix D). Thus, using a more careful definition of DIB S/N, we find no evidence for DIB detections within the Local Bubble in public Gaia DR3 RVS spectra (Schultheis et al. 2023), which calls for careful reconsideration of related previous claims. We look forward to the future release of more and higher stellar S/N spectra to help resolve this question.

3.3. Moment-0

The moment-0 (EW_{DIB}) map of DIB detections (Figure 3) closely resembles a noisy version of 3D dust maps obtained using other techniques (Green et al. 2019). Figure 3 also clearly features detections predominantly along lines of sight with the highest extinction. This makes sense given the known positive correlation between DIBs and dust and the low to moderate S/N of most of the stellar spectra; confident DIB detections (high DIB S/N) require large DIB signals and, by correlation, large dust column densities.

Studies of DIBs often focus on the linearity of the correlation of the EW_{DIB} with dust column density, extinction A_V , or reddening $E(B-V)$ (Kos & Zwitter 2013; Schultheis et al. 2023). However, several DIBs are known to “flatten out” or even “turn over,” having a negative correlation of EW_{DIB} with dust extinction past a critical point (Lan et al. 2015). This could be, for example, because the DIB carrier freezes out of the gas phase onto grains, chemical reactions deplete the DIB carrier, or there are several dust clouds along the line of sight, some of which do not host DIBs.

Previous work using purely statistical correlations without imposing detection cuts has found linear correlations with near-zero intercepts (Zasowski et al. 2015). However, the creation of a catalog necessarily imposes a detection bias that more significantly impacts DIBs at low reddening and can thus lead to a positive intercept of EW_{DIB} with A_V .

In Figure 4, we show a weighted linear least-squares fit of EW_{DIB} versus the extinction to the background star, as measured by the Bayestar19 3D dust map (Green et al. 2019). The weights are simply the inverse variance of the EW_{DIB} and the fit is performed in the range $0 < A_V < 1$. We find a slope of $106 \pm 33 \text{ m}\text{\AA} \text{ mag}^{-1}$ and an intercept of $17 \pm 26 \text{ m}\text{\AA}$.

Because Bayestar19 has significant scatter around the line of sight fits and we do not know what fraction of the reddening to the star is associated with the DIB, there is a large, but hard to

¹⁰ https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/DIB_Local_Bubble_SNR_3D.html

quantify, uncertainty in the extinction axis. Thus, we derive our error bars on the slope and intercept from independent fits to the DIB detections in each quadrant in Galactic longitude. Further, in all fits and figures, we restrict to lines of sight where the Bayestar19 extinction prediction is unchanged within $\pm 1\sigma$ of the distance to the background star according to the distance uncertainties in Bailer-Jones et al. (2021). While this selection could introduce a bias with respect to Galactic latitude and the DIB environment, we leave the investigation of these more subtle effects to future work.

For reference, the 15273 Å DIB which has previously been used for successful ISM mapping has a reported slope of $102 \pm 1 \text{ mÅ mag}^{-1}$ versus A_V (Zasowski et al. 2015). To directly compare the slope to Schultheis et al. (2023), which measured the same 8621 Å DIB, we convert A_V Bayestar19 to $E(B - V)$ to find $290 \pm 91 \text{ mÅ mag}^{-1}$ versus $E(B - V)$. While Schultheis et al. (2023) report in mag mÅ^{-1} , a simple inversion of the slope yields 455 mÅ mag^{-1} versus $E(B - V)$, which is 1.8σ larger than the value found here. Careful control of EW_{DIB} biases and uncertainties and using those uncertainties in the fit is important to avoid slight overestimation of the slope.

3.4. Moment-1

3.4.1. Kinematics

To use the DIB to study kinematics (Figure 5), we must first convert the λ_{DIB} measured in the stellar frame to $v_{\text{DIB}}^{\text{LSR}}$. We remove the wavelength shift caused by reporting spectra in the stellar frame using the barycentric radial velocity v_{star} of the star reported by Gaia.¹¹ At the same time, we remove the wavelength shift caused by the motion of the Sun relative to the LSR, given the unit direction vector to the source \hat{d} . We use the LSR convention where $V_{\text{Sun}} = (U, V, W) = (10.6, 10.7, 7.6) \text{ km s}^{-1}$ (Reid et al. 2019).¹² Then, using the rest-frame wavelength λ_{rest} and the speed of light c , we convert the “LSR” wavelength λ_{LSR} to a velocity.

$$\lambda_{\text{LSR}} = \lambda_{\text{DIB}}(1 + v_{\text{star}}/c + V_{\text{Sun}}/c \cdot \hat{d}) \quad (1)$$

$$v_{\text{DIB}}^{\text{LSR}} = \left(1 - \frac{\lambda_{\text{rest}}}{\lambda_{\text{LSR}}}\right)c \quad (2)$$

The reported $v_{\text{DIB}}^{\text{LSR}}$ uncertainty is obtained by propagating the λ_{DIB} uncertainty only and does not include uncertainty contributions from v_{star} or overall zero-point uncertainty related to V_{Sun} .

It is essential to measure the rest-frame (intrinsic) wavelength of the DIB to enable the use of the DIB in studies of Galactic kinematics. Additionally, precise measurement of the rest-frame wavelength aids in identifying its carrier in comparison to theoretical chemical models. Here we use the common method of measuring the wavelength intercept of the DIB detections toward the Galactic anticenter (Munari et al. 2008; Zasowski et al. 2015; Zhao et al. 2021; Schultheis et al. 2023). Assuming a purely tangential, azimuthally symmetric Galactic rotation, DIB carriers rotating with the galaxy would exhibit no shift in the λ_{LSR} relative to the rest-frame wavelength toward the Galactic anticenter.

¹¹ We use a positive sign convention for radial velocity moving away from the observer.

¹² We use a right-handed convention with U pointing toward the Galactic center and V toward $\ell = 90^\circ$.

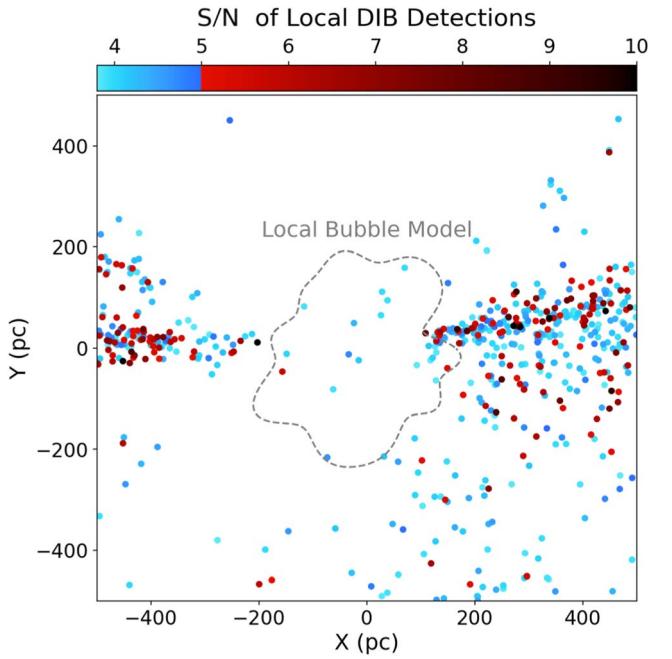


Figure 2. MADGICS DIB detections within the Local Bubble, where color indicates the S/N of the detections. A 2D slice at the midplane ($Z = 0 \text{ pc}$) of the boundary for a model of the Local Bubble (Pelgrims et al. 2020; Pelgrims 2022) is shown (dashed gray). In the MADGICS catalog, no confident detections (DIB S/N > 5) are present within the Local Bubble model. This is in contrast to previous work, though we find evidence that previously reported detections were associated with stellar features (Figure 1). The interactive version of the figure provides a manipulable 3D view of the DIB detections, Local Bubble model, location of the Sun, and estimates of the dust distribution from Leike et al. (2020), with the ability to turn each layer on/off individually. Hover tools over each DIB detection also provide the DIB S/N, EW_{DIB} , $v_{\text{DIB}}^{\text{LSR}}$, and σ_{DIB} . (An interactive version of this figure is available.)

The expected change in λ_{LSR} of the DIB as a result of Galactic rotation can be found using Equation (3) when given the distance of the Sun to the Galactic center R_0 , the circular (tangential) velocity of Galactic rotation V_c at the location of the DIB source \mathbf{R} , Galactic longitude of the DIB source referenced to the Galactic anticenter $\ell' = \ell - 180^\circ$, and distance from the Sun to the DIB source projected on to the Galactic plane d . This can be linearized in the small angle approximation to yield Equation (4).

$$\frac{\Delta \lambda_{\text{LSR}}}{\lambda_{\text{rest}}} = \frac{V_c(\mathbf{R}) \sin \ell'}{c} \times \left[1 - \frac{1}{\sqrt{\sin^2 \ell' + \frac{d^2}{R_0^2} + 2 \frac{d}{R_0} \cos \ell' + \cos^2 \ell'}} \right] \quad (3)$$

$$\frac{\Delta \lambda_{\text{LSR}}}{\lambda_{\text{rest}}} = \frac{V_c(\mathbf{R})}{c} \frac{d}{d + R_0} \ell' \quad \text{where } \ell' \ll 90^\circ \quad (4)$$

Because V_c is believed to be locally flat (Bovy et al. 2012) we will take $V_c(\mathbf{R})$ to be a constant. We fix $R_0 = 8.5 \text{ kpc}$ consistent with Reid et al. (2019). While we only know the distance to the background star, we will use this to approximate d , the distance to the DIB source. To help this approximation hold, we limit the fit to background stars within 3 kpc.

In Figure 6, we present a least-squares fit of λ_{LSR} as a function of ℓ' and d using Equation (4) and a constant term to measure λ_{rest} . We use inverse-variance weights that sum in

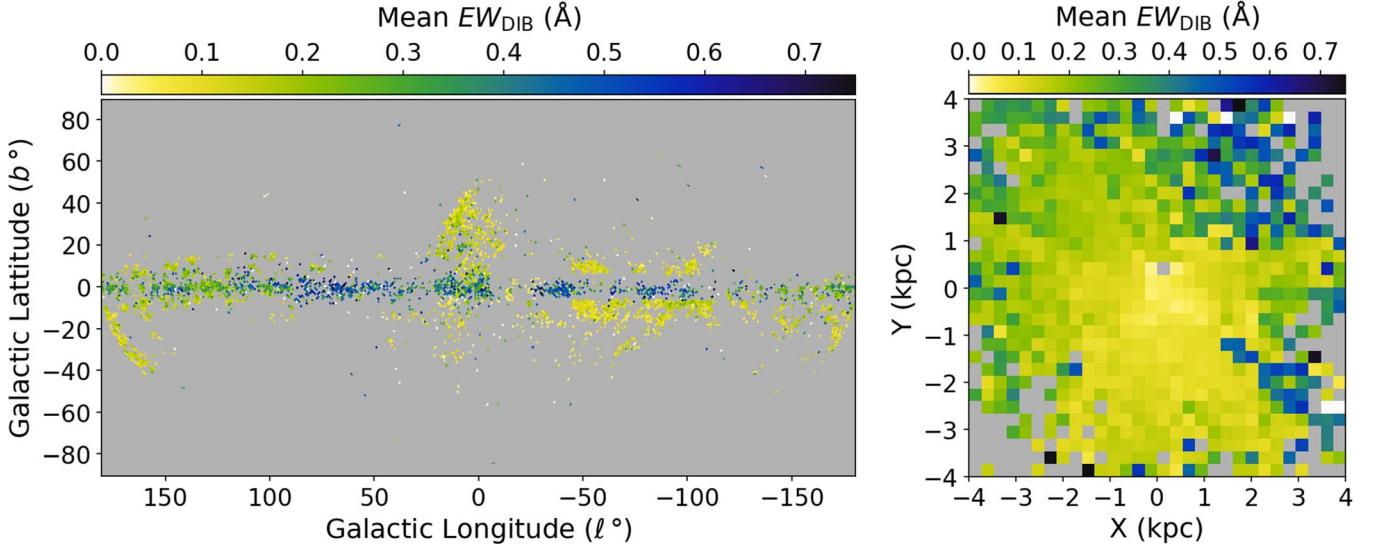


Figure 3. Projections of the 3D DIB EW map in (left) angular (ℓ, b) and (right) Cartesian (X, Y) Galactic coordinates. The inverse-variance weighted mean of EW_{DIB} is shown in color, averaging over dimensions not shown.

quadrature the MADGICS $\sigma(\lambda_{DIB})$ and the background star radial velocity uncertainty from Gaia. We restrict to $|\ell'| < 10^\circ$ so that we can safely apply the small angle approximation and $|b| < 10^\circ$ to help approximate the Galactic rotation as uniform in the z -direction. The intercept gives $\lambda_{\text{rest}} = 8623.14 \pm 0.087$ Å (vacuum).

To obtain the uncertainty estimate on the least-squares model parameters, we divide the data into $0 < d < 2$ kpc and $2 < d < 3$ kpc. We use the difference in the fit parameters relative to the formal uncertainties to determine a scalar by which to inflate the formal uncertainty so that the difference is only 1σ . We perform this procedure to help account for uncertainty in R_0 , resulting from the approximation of d , systematics associated with cuts in (ℓ', b) that are imposed, and the rather strong assumption of a purely tangential, azimuthally symmetric Galactic rotation.¹³ For example, measurements from Tchernyshyov & Peek (2017) show a net radial inflow past 1 kpc on the order of -3 km s $^{-1}$. We note that the 0.087 Å uncertainty reported here is comparable to that inflow velocity.

The value found in this work is bluer than that reported in Schultheis et al. (2023) by 4.7σ , using the uncertainties reported therein (8623.23 ± 0.019 Å vacuum), though only 1.03σ using the uncertainties reported in this work. We further validate the λ_{rest} and its uncertainty estimate reported here by comparison with ^{12}CO velocities in Section 3.4.2.

An alternative method for determining the rest-frame wavelength is referencing to known atomic ISM transitions, such as Na I (Krelowski 1988; Galazutdinov et al. 2000). While the wavelength of the Na I transition is known with high precision and its use alleviates assumptions about Galactic rotation, surveys measuring both DIBs and transitions for atomic ISM species have historically had far fewer samples. Further, referencing Na I requires assuming the Na I and DIB carriers are comoving, which can have large scatter (Farhang et al. 2015).

Other applications of the Galactic anticenter method to DIBs have conventionally performed a simple linear fit, holding d to be a fixed constant, and reported the intercept as λ_{rest} and the

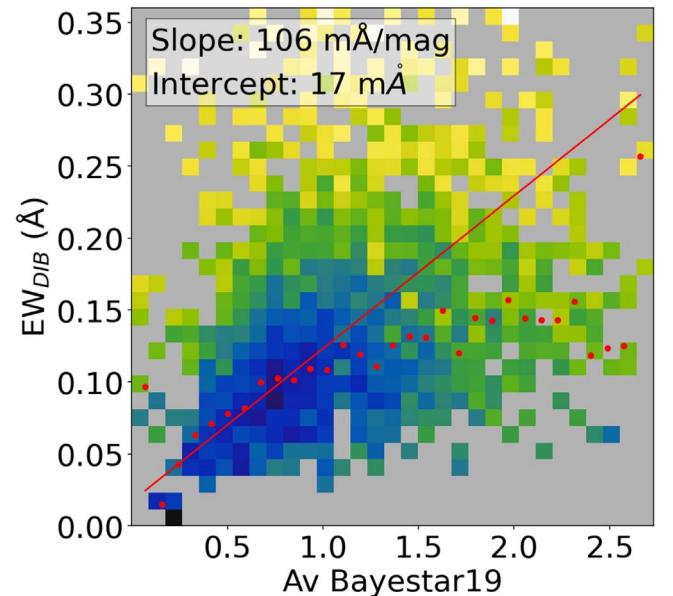


Figure 4. Weighted least-squares linear fit of EW_{DIB} vs. extinction to the background star as measured by Bayestar19 (Green et al. 2019) in the range $0 < Av < 1$. The color scale indicates EW_{DIB} inverse-variance weighted number density, where darker colors indicate more (confident) detections and a log normalization has been applied. Scatter points represent inverse-variance weighted mean EW_{DIB} for each Av bin.

slope in mÅ deg $^{-1}$. However, it is very difficult to compare these slopes between DIBs at different wavelengths or even the same DIB measured in samples with a different distribution of DIB carrier distances d (see Equation (4)). Instead, our fit parameter in Equation (4) is $V_c = 204 \pm 60$ km s $^{-1}$. This can be easily compared to Galactic rotation models. This value is $<1\sigma$ from that of Bovy et al. (2012), which is 216 ± 6 km s $^{-1}$. However, given the approximations made above, this should not be viewed as a measurement of V_c , but as a sanity check.

To compare to other work, we can compute the expected anticenter slope for a given d . Taking $d = 2$ kpc, we expect a slope of 19 mÅ deg $^{-1}$. Assuming a similar DIB carrier distance distribution, we can compare to Zasowski et al. (2015) by

¹³ The absolute uncertainty associated with V_{Sun} always remains.

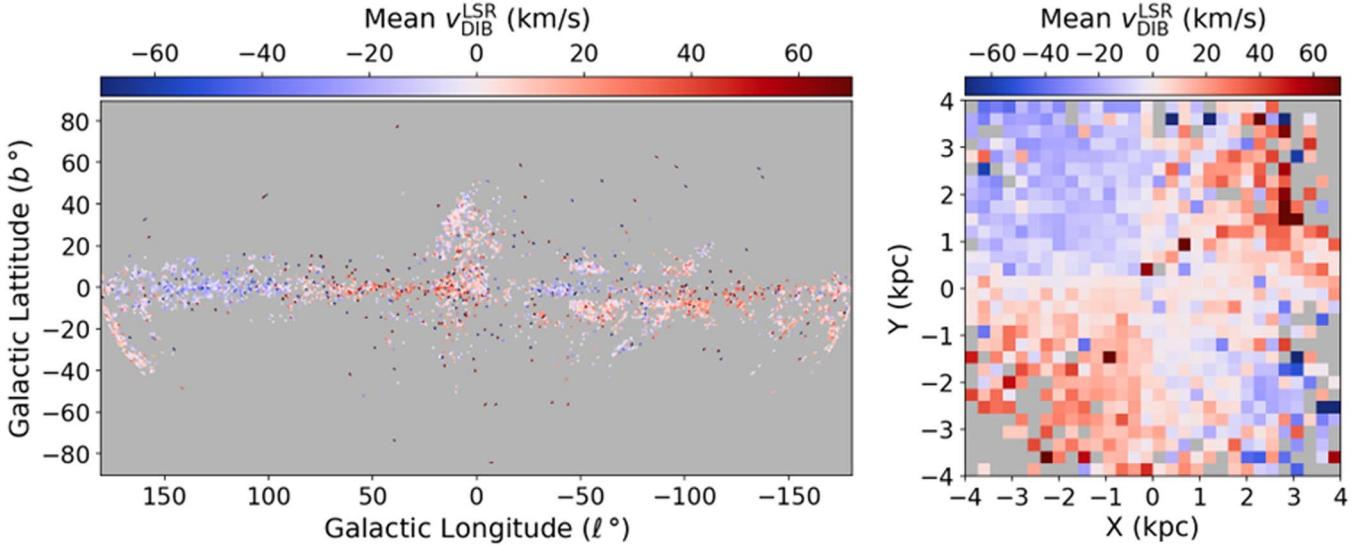


Figure 5. Projections of the 4D DIB radial velocity map in (left) angular (ℓ, b) and (right) Cartesian (X, Y) Galactic coordinates. The inverse-variance weighted mean of $v_{\text{DIB}}^{\text{LSR}}$ is shown in color, averaging over dimensions not shown.

multiplying by the ratio of the rest-frame wavelengths of the DIBs, $=0.56$.¹⁴ This would predict $32.1 \pm 4.5 \text{ m}\text{\AA} \text{ deg}^{-1}$. Schultheis et al. (2023) report $45 \pm 20 \text{ m}\text{\AA} \text{ deg}^{-1}$, which uses the same DIB, but a different sample of detections and a clearly different λ_{DIB} (see Figure 1). However, the comparison is made more difficult by the fact that other work is often only correct for the U_{Sun} velocity (Zasowski et al. 2015; Schultheis et al. 2023). The correction resulting from the V_{Sun} velocity is approximately linear with ℓ' near the Galactic anticenter and applying the correction reduces the resulting slope significantly.

With these validations (and those in Section 3.4.2), we can combine $v_{\text{DIB}}^{\text{LSR}}$ with the Gaia location and parallax measurements to present a 4D map showing the radial velocity of DIB carriers as a function of position. Figure 5 shows a projection of this map in angular (ℓ, b) and Cartesian (X, Y) Galactic coordinates, taking an inverse-variance weighted mean over the dimensions not shown. A lower noise map could be achieved with stricter DIB S/N cuts, though this comes at the expense of the number of detections and spatial coverage given the limited public Gaia RVS sample. The usual features (alternating sign quadrants) of a smooth Galactic rotation curve are observed. We defer searching for anomalies relative to a smooth Galactic rotation model until having a larger sample.

3.4.2. CO Comparison

^{12}CO is believed to be a good tracer of DIBs and dust because it has a well-known, easy-to-detect line and indicates the presence of carbonaceous molecules. We compare DIB kinematics to the Dame et al. (2001) moment-masked, composite ^{12}CO survey.¹⁵ We use the full galaxy cube ($-180^\circ < \ell < 180^\circ$, $|b| < 30^\circ$) which is reported on a rectilinear Galactic Cartesian grid ($1/8^\circ$ spacing) and a velocity grid ($|v_{\text{CO}}^{\text{LSR}}| < 320 \text{ km s}^{-1}$) with 1.3 km s^{-1} spacing. The LSR frame assumes the “classic” solar motion of 20 km s^{-1} toward (R.A., decl.) = $(18^\text{h}, 30^\circ)$ (epoch 1900), which is approximately $V_{\text{Sun}} = (10.3, 15.3, 7.7) \text{ km s}^{-1}$

¹⁴ Zasowski et al. (2015) restricted their fit to background stars within 2 kpc, which helps make this assumption more valid.

¹⁵ <https://lweb.cfa.harvard.edu/rtdc/CO/CompositeSurveys/>

private communication, Dame, 2022. Figure 7 shows the CO map, integrated along b .

In order to compare $v_{\text{DIB}}^{\text{LSR}}$ to the CO kinematics, we convert $v_{\text{DIB}}^{\text{LSR}}$ into the LSR frame convention used by Dame et al. (2001) and overplot the DIB detections with DIB S/N > 5.5 in Figure 7. The average $v_{\text{DIB}}^{\text{LSR}}$ as a function of ℓ follows the intensity-weighted average v_{CO} (Figure 8). Excitingly, we further observed a subpopulation of the DIBs that coincide with higher-amplitude oscillations in the CO map. Coloring the scatter plot by the distance to the background star, we find this subpopulation of DIB detections occurs along lines of sight to more distant stars (compared to the average within the sample). This is consistent with the interpretation of higher-amplitude oscillations in ^{12}CO as a kinematically coherent Galactic substructure at larger distances, sometimes attributed to spiral arms.

We quantitatively demonstrate the correlation between $v_{\text{DIB}}^{\text{LSR}}$ in Figure 8. Motivated by the multimodality shown in Figure 7, before taking the intensity-weighted average of v_{CO} , we find the peak in the CO spectrum nearest to $v_{\text{DIB}}^{\text{LSR}}$ for the (ℓ, b) -pixel in the CO map corresponding to the line of sight for a DIB detection. Peak finding has a tendency to slightly overstate the correlation, but the clumpiness of the CO necessitates such an approach. This peak-finding algorithm is discussed in detail in Appendix E, and we show there that a correlation similar to that shown in Figure 8 is obtained when restricting to low-velocity CO without any peak finding. A linear least-squares fit, restricted to $|v_{\text{CO}}| < 25 \text{ km s}^{-1}$, with inverse-variance weights using $\sigma(v_{\text{DIB}}^{\text{LSR}})$ is also shown in Figure 8. The slope of 0.95 ± 0.03 illustrates a very strong 1:1 relationship between the $v_{\text{DIB}}^{\text{LSR}}$ and v_{CO} , suggesting that where they coexist, DIB carriers and CO are predominately comoving.

The intercept of the linear correlation in Figure 8 provides a semi-independent check on the λ_{rest} and its associated uncertainty reported in Section 3.4.1 because the frequency of the CO transition is known to be high precision. An offset of $-0.52 \pm 0.25 \text{ km s}^{-1}$ corresponds to a $\Delta\lambda_{\text{rest}} = -0.015 \text{ \AA}$ or $< 0.2\sigma$ given the reported λ_{rest} uncertainty. Further, the scatter of residuals $/\sigma(v_{\text{DIB}}^{\text{LSR}})$, or Z-scores, was 0.6, indicating that the reported uncertainties slightly overestimated the scatter in the

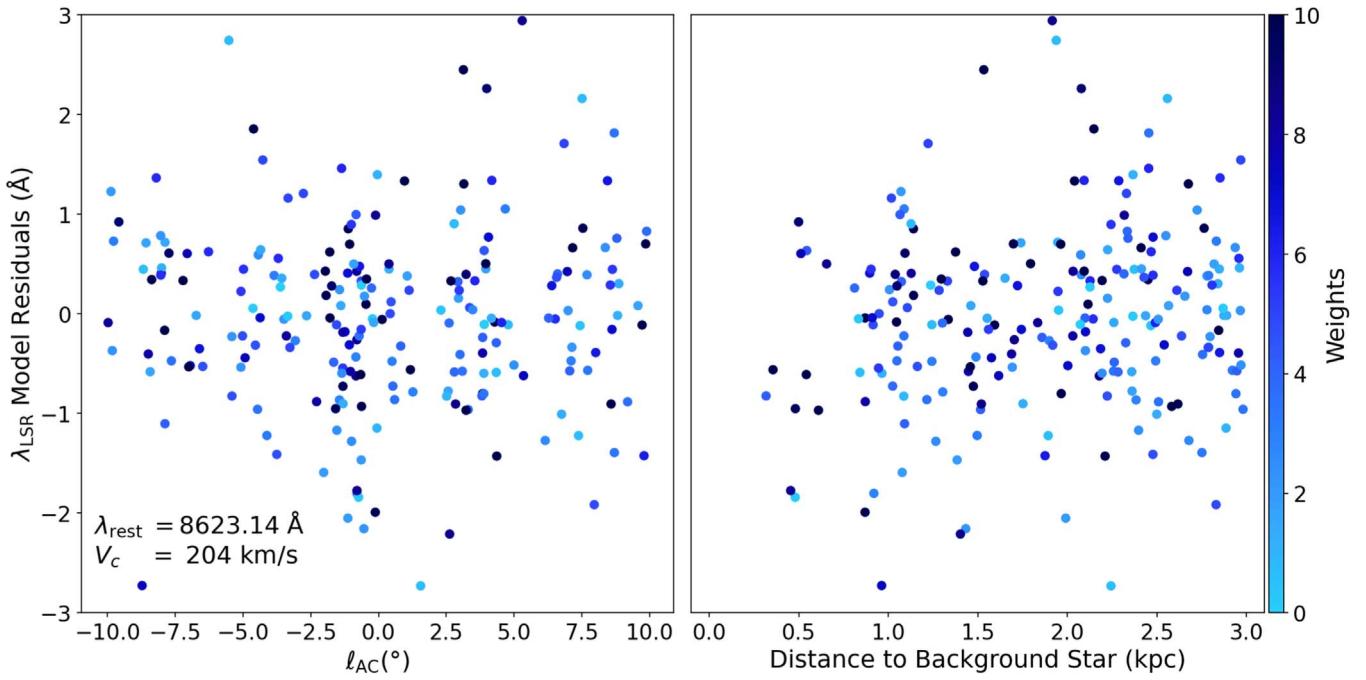


Figure 6. Scatter plot showing residuals after fitting λ_{LSR} vs. ℓ' and d where the color shows the weighting of each point in the fit. Only detections such that $|\ell'| < 10^\circ$, $|b| < 10^\circ$, and $d < 3$ kpc are shown. Our derived rest-frame wavelength of $\lambda_{\text{rest}} = 8623.14 \pm 0.087 \text{ \AA}$ is bluer than Schultheis et al. (2023), but our updated value is validated in Figure 8 by comparison with CO gas velocities.

residuals. However, given the selection bias toward correlation given the use of a peak-finding algorithm, a suppression of the scatter in the residuals is to be expected.

Thus, correlation with v_{CO} provides a secondary semi-independent validation of λ_{rest} , $v_{\text{DIB}}^{\text{LSR}}$, and their associated uncertainties, which lends confidence to the use of the MADGICS DIB catalog for kinematics studies.

3.5. Moment-2

The width of an absorption feature (moment-2) can often be used to infer properties about the physical environment or transition itself. In radio observations of the ISM, line widths are often attributed to a combination of Doppler, supersonic turbulence, and opacity broadening (Hacar et al. 2016). However, given the wavelength and observed line width for the DIBs, simple order of magnitude calculations show that these mechanisms are unlikely to contribute significantly to the line width (Edwards & Leach 1994). The dominant intrinsic source of broadening is lifetime broadening, which agrees with the very short excited state lifetimes predicted for most DIB carrier candidates (Snow et al. 2002; Campbell et al. 2015).

The observed line width can also appear broader because of heterogeneity in the central wavelength. For example, multiple dust clouds moving at different velocities along the line of sight can lead to multiple unresolved peaks (Tchernyshyov et al. 2018). As an order of magnitude, a difference in velocity of 20 km s^{-1} would lead to a wavelength shift that is only half of the typical σ_{DIB} , making resolving two components in the spectra extremely difficult. We expect this to be one of the largest sources of variability in observed σ_{DIB} . Heterogeneity in the chemical environment, which can slightly shift the transition wavelength for some carriers in the column density probed by the line of sight, is another possible source of broadening. However, given that most DIB carriers remain

unidentified, it is difficult to make predictive theories about the role of chemical environment heterogeneity.

We observe (Figure 9) a weak positive correlation of σ_{DIB} with the extinction (Av) from Bayestar19 (Green et al. 2019). This could result from there being more components (clouds at different velocities) on average along the lines of sight with higher reddening. It could also result from different environments in higher extinction clouds. Because many of the highest reddening lines of sight are also low signal-to-noise and we show a positive detection bias in σ_{DIB} , we hold off on rigorous interpretation of σ_{DIB} until a larger sample of Gaia RVS spectra are available so as to have more statistical power.

4. DIB Model

We provide a brief introduction to MADGICS and the specific choices made in the application to Gaia RVS spectra, with a more detailed description of the method in future work.

4.1. Statistical Description

MADGICS is a component separation algorithm that decomposes an observation, a data vector x_{tot} , into a linear combination of components x_k . In order to speak about such a decomposition, one must have prior information about the different components, which we require to sum to the data, and about the properties of each component. In MADGICS, we express that prior as a pixel-pixel covariance matrix C_k in the data space. In this case, this is literally the covariance of wavelength bin i and wavelength bin j in the spectrum (Equation (5)). Prior information about the mean of a component μ_k can also be included, though in practice we often set $\mu_k = 0$. Here, we have implicitly assumed there is no cross-component covariance in the priors. For notational convenience, we define the sum of the covariance and mean

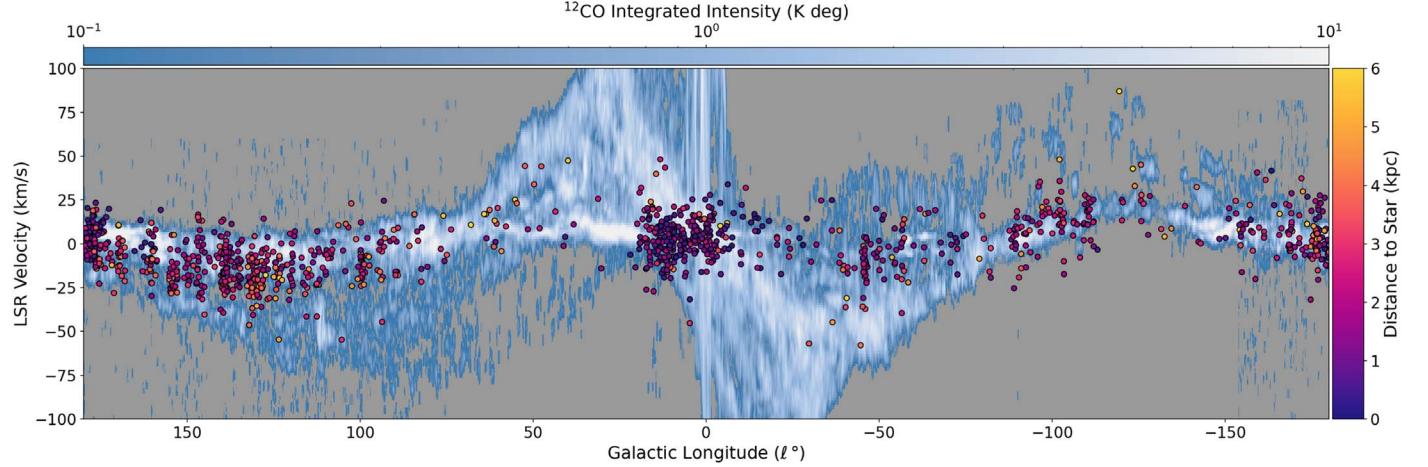


Figure 7. Scatter plot of $v_{\text{DIB}}^{\text{LSR}}$ for DIB detections with DIB S/N > 5.5 overlaid on longitude–velocity diagram of CO emission from Dame et al. (2001), integrated over b (color bar, top). Scatter points are colored by the distance to the background star (color bar, right).

priors over all components (Equation (6)).

$$x_{\text{tot}} = \sum_{k=1}^{N_{\text{comp}}} x_k \quad \text{where} \quad x_k \sim \mathcal{N}(\mu_k, C_k) \quad (5)$$

$$C_{\text{tot}} = \sum_{k=1}^{N_{\text{comp}}} C_k \quad \mu_{\text{tot}} = \sum_{k=1}^{N_{\text{comp}}} \mu_k \quad (6)$$

To obtain a posterior for the contribution of each component to the signal, we specify a likelihood that is the product of the Gaussian probability for \hat{x}_k given $\mathcal{N}(\mu_k, C_k)$ for each of the components and then impose the constraint that the components sum to the data to obtain a posterior. The posterior mean and covariance per component can be obtained analytically as

$$\hat{x}_k = C_k C_{\text{tot}}^{-1} (x_{\text{tot}} - \mu_{\text{tot}}) + \mu_k \quad (7)$$

$$\hat{C}_{kk} = (I - C_k C_{\text{tot}}^{-1}) C_k \quad (8)$$

$$\hat{C}_{km} = -C_k C_{\text{tot}}^{-1} C_m \quad (9)$$

where I is the identity matrix, \hat{x}_k is the predicted mean component, \hat{C}_{kk} is the predicted pixel–pixel covariance of pixels in component k , and \hat{C}_{km} is the predicted covariance of a pixel in component k with a pixel in component m . See Appendix H for an example of these predicted (posterior) covariances.

In terms of computational ease, Equations (7)–(9) show that the solution depends on only a single matrix inverse C_{tot}^{-1} , which can be computed quickly using low-rank approximations (see Section 4.2) and Woodbury updates (Woodbury 1950). In this work, processing 1 million Gaia RVS spectra took only 650 core hours.¹⁶

In MADGICS, all of the difficulty lies in specifying the models and priors on each component in the model. For Gaia RVS spectra, we perform two decompositions modeling the data as either star + residual or star + DIB + residual. The sum of \hat{x}_{DIB} and $\hat{C}_{\text{DIB}, \text{DIB}}$ give the estimates of EW_{DIB} and the uncertainty in EW_{DIB} , respectively.

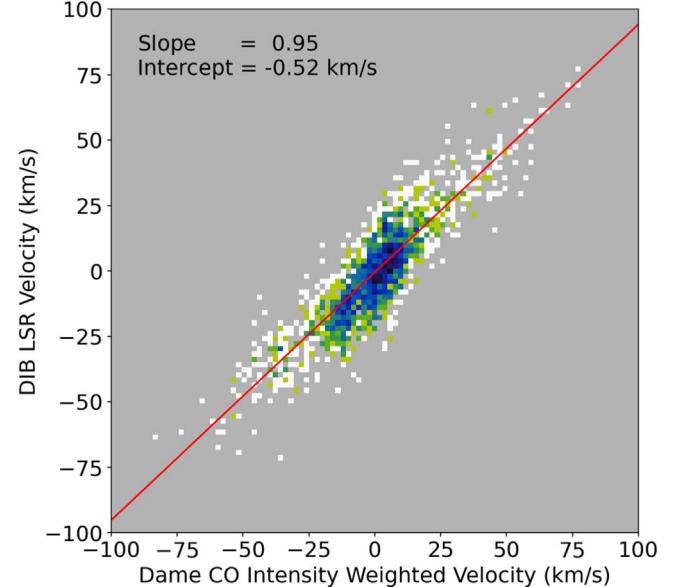


Figure 8. Histogram (2D) illustrating the correlation between the $v_{\text{DIB}}^{\text{LSR}}$ and intensity-weighted velocity of the nearest CO component in the Dame et al. (2001) CO map. The color scale indicates density, where darker colors indicate more detections and a logarithmic stretch has been applied.

4.2. Prior Construction

4.2.1. Stellar Prior

The stellar covariance is built from a high (> 70) stellar S/N and low reddening ($\text{SFD } E(B-V) < 0.05$ mag) training set, where we have used a cut on a 2D infrared emission-based dust map that acts as an upper bound on the dust along the line of sight (SFD; Schlegel et al. 1998). We also require GSP-Spec Flags 1–13 to be zero (the highest quality value) in order to be most exclusive against defects and outliers. After cuts, this is 39,657 spectra. For simplicity, we slightly restricted the spectral range of the stellar model to 8469.5–8688.3 Å (2189 versus 2401 pixels) where the vast majority of all Gaia RVS spectra had coverage so that we could work in a single common data space.

Since the Gaia RVS spectra are continuum normalized, we subtract an average constant continuum of 0.95. In the notation

¹⁶ Production runs were executed on AMD 64-core EPYC 7702P processors (AVX2) with 480 GB RAM running a Linux kernel v4.18.0 (Rocky Linux).

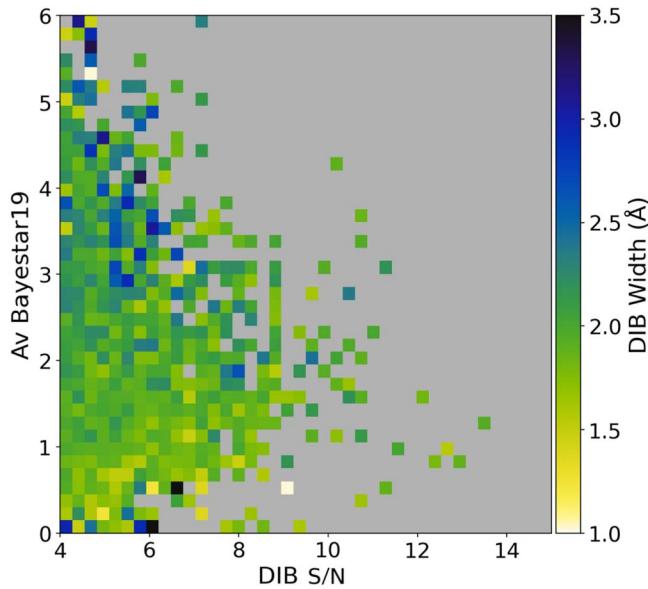


Figure 9. The median of the DIB width σ_{DIB} as a function of the DIB S/N and the extinction along the line of sight in the Bayestar19 3D dust map (Green et al. 2019). The weak correlation could be due to the presence of multiple cloud components at different velocities along the line of sight and/or the presence of velocity structure within individual clouds.

of Section 4.1, this is letting $\mu_{\text{star}} = 0.95$ times the ones vector. The covariance is then built by multiplying the matrix of these observations by its transpose, with the relative contributions weighted by stellar S/N.¹⁷ However, since the continuum normalization is not stable (see Appendix A), we add 0.1 times the one's matrix to the covariance so that the stellar component will absorb overall shifts in the continuum normalization.¹⁸

The resulting stellar covariance matrix (the MADGICS prior) is shown in Figure 10. Despite the complexity of stellar spectra, this covariance is in general well described by a low-rank approximation, which also significantly accelerates the computation. We keep the first 50 eigenvectors to describe the stellar covariance, one of which is shown in Figure 10.¹⁹ We explicitly check for and do not find significant peaks centered at the DIB wavelength.

In Figure 10, we overplot the same guidelines as in Figure 1. We show that there are significant peaks, comparable to those for lines carefully calibrated in the GSP-Spec line list, centered at the wavelengths of the pileups identified in Figure 1 as having the largest impact on biasing the Gaia DIB catalog. By plotting a higher-resolution solar spectrum from Kurucz et al. (1984) for reference, we see that the features in our stellar model that are coincident with those pileups also coincide with large amplitude features in the solar spectrum. Here we adopt the largest amplitude lines from the Kurucz archive (see

¹⁷ Weights were capped at 300 to prevent any one spectrum from dominating the covariance.

¹⁸ The value 0.1 was chosen to be comparable to the average pixel in the covariance matrix far from the diagonal and clear stellar features. However, the component separation is extremely stable to changes in the choice of this value. Similar results are also obtained by manually adding an eigenvector of ones to the low-rank approximation (described below) instead of adding a constant times the one's matrix to the full covariance matrix (as is described here).

¹⁹ The number of eigenvectors in the low-rank approximation was chosen in part by inspecting the steep falloff of the eigenspectrum. We also ran empirical tests that showed negligible changes in the stellar component obtained during the decomposition of test spectra as a function of the dimension of the low-rank approximation by 50.

footnote 7) contributing to those features in high-resolution spectra of both the Sun (Kurucz et al. 1984) and Arcturus (Hinkle et al. 2000) as labels for the respective pileup locations (see Appendix G). However, this comparison confirms that the contamination in the Gaia DIB catalog shown in Figure 1 is of stellar origin, regardless of the specific line assignments of those features.

4.2.2. Residual Prior

In all measurements, there is a contribution from noise that will follow its own covariance structure, and thus we also include a “residual” component to capture this noise. Ideally, these measurement uncertainties are uncorrelated and the residual prior is a simple diagonal. In that case, unmodeled features in the data which do not conform well to the prior of any other component in the model (such as a previously unidentified DIB) will simply appear in the residuals.

Gaia DR3 includes a per-pixel flux uncertainty vector, which could be used on the diagonal of the residual covariance prior. However, when doing so, one finds that the MADGICS residual component in fact has correlations, with a kernel that appears to be the result of upstream processing interpolating the spectra onto a common grid.²⁰ These correlations impact the uncertainty of EW_{DIB} , but we want to retain a diagonal residual prior, so we inflate the residual prior by multiplying by a constant (3.6) such that the sum of the correlated noise and uncorrelated noise covariances are the same.

4.2.3. DIB Prior

The contribution from DIBs is fundamentally multiplicative since it is an absorption of the background starlight. Since MADGICS is additive, it finds a linear decomposition of the data into components, we first model the spectra (even one containing a DIB detection) as being star + residual components (Figure 11). Then, we express a prior on the line shape of the DIB absorption. In this case, we use a Gaussian for simplicity, but since we need only specify the DIB prior as a pixel-pixel covariance matrix, we can accommodate arbitrary line shapes and/or degrees of asymmetry. Sampling over variations in the line shape, we construct a pixel-pixel covariance matrix.

Here we keep only a two-component, low-rank approximation for the DIB covariance, which we generate analytically using the functional form of a Gaussian and its derivative with respect to wavelength. We use a constant amplitude prior of 0.1 and reduce the amplitude of the derivative by a factor of 0.01. The amplitude value was chosen to be large relative to the expected DIB signal (a “loose” prior). MADGICS is able to use such “loose” priors and be insensitive to the specific value chosen because it leverages the correlation of many pixels. Stated another way, low-dimensional slices in high-dimensional spaces are fairly orthogonal. The derivative value was chosen to marginalize over shifts in the observed wavelength smaller than the step size of the radial velocity scan (see below).

However, this covariance is a prior on x'_{DIB} while the linear contribution of the DIB to the spectrum is $x_{\text{DIB}} = x'_{\text{DIB}} \times x_{\text{star}}$. Since we do not know x_{star} , we use \hat{x}_{star} from the star + residual decomposition and rescale the covariance for x'_{DIB} to obtain the

²⁰ The kernel has the form [0.1, 0.4, 0.8, 1.0, 0.8, 0.4, 0.1] centered on a given pixel.

covariance for x_{DIB} . Since the contribution from the DIB is predominantly in the residual component in the star + residual model, the \hat{x}_{star} changes only slightly in the star + DIB + residual model (see Figure 11). We can then update \hat{x}_{star} , using the component from the star + DIB + residual model, and iterate until convergence. In practice, we only iterate to refine the x_{DIB} once.

While MADGICS can easily marginalize over (co)variations in peak amplitudes, it struggles to marginalize over shifts in observed wavelength (radial velocity) or changes in peak width (σ_{DIB}) which are not well described in the data space representation. In order to optimize over radial velocity and σ_{DIB} , we build a grid of covariance priors over a range of $(\lambda_{\text{DIB}}, \sigma_{\text{DIB}})$ and evaluate the $\Delta\chi^2$ at each point, seeking the minimum (most negative) $\Delta\chi^2$.

In practice, we find the search of the 2D $(\lambda_{\text{DIB}}, \sigma_{\text{DIB}})$ -space well-approximated using line searches with a single iteration of refinement. The λ_{DIB} scan is ± 100 pixels, in 0.1 pixel steps (0.01 \AA) at the highest resolution. The σ_{DIB} scan is 0.4 to 4 \AA in 0.01 \AA steps at the highest resolution (see code in Section 6 for more). The resolution of the λ_{DIB} and σ_{DIB} scans was chosen such that the step size is $\lesssim 10\%$ of the typical uncertainty in each parameter estimate.

The λ_{DIB} and σ_{DIB} uncertainties are estimated from the shape of the $\Delta\chi^2$ surface around the minimum, using the Hessian estimated with finite differences. For a given point in the $(\lambda_{\text{DIB}}, \sigma_{\text{DIB}})$ -grid, we have estimates of EW_{DIB} and the uncertainty in EW_{DIB} given by the sum of \hat{x}_{DIB} and $\hat{C}_{\text{DIB}, \text{DIB}}$, respectively. However, in order to marginalize over σ_{DIB} , we combine these Gaussian estimates of the EW_{DIB} for several σ_{DIB} near the minimum. These estimates are combined with relative weights given by their likelihood ratios.

The final optimal decomposition for the example in Figure 11 is shown in black. The substructure in the DIB \times star component is the result of the stellar features, which is removed upon dividing by the stellar component. The DIB \times star and residual components are on the same scale and have comparable amplitudes. Note that the wavelength range in Figure 11 is intermediate to the full RVS spectral range and the window near the DIB shown in Figures 1 and 10. The DIB profile is fairly wide relative to the narrow window shown previously when focusing on stellar lines.

4.3. Stellar Validation

While we do not focus on modeling the stars, we briefly illustrate the quality of stellar modeling because it impacts which background stars can be used to detect DIBs and the magnitude of any possible biases resulting from stellar features. In Figure 12, we show both the distribution of the stellar training set and all public Gaia DR3 RVS spectra as a function of $\log g$, $[\text{M}/\text{H}]$, and $\log T_{\text{eff}}$ as estimated by GSP-Spec. The training set is an order of magnitude smaller, but has a similar shape compared to the full distribution. Only populations at the extremes in temperature and metallicity are notably missing, populations that form a small fraction of the overall sample.

Measuring how well spectra outside the training set are modeled (generalization) is one way to understand the flexibility and validity of the stellar prior. We do this by showing the median $\log(\chi_{\text{tot}}^2/\text{dof})$ for spectra as a function of the stellar parameters. The color scale is saturated at the value used as a “goodness of fit” cut in making the catalog (see Table 1 and Section 5.1). This highlights regions of stellar

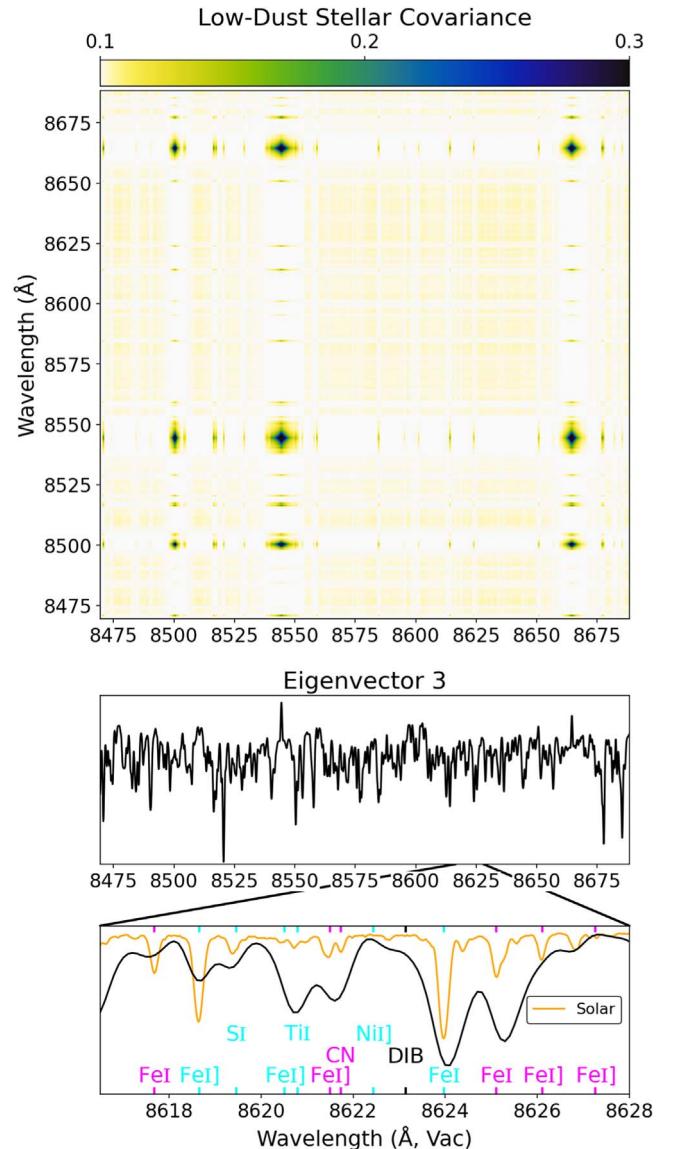


Figure 10. (Top) The covariance matrix of high stellar S/N and low-dust spectra used as MADGICS prior for the stellar component. (Middle) Third eigenvector of the stellar covariance matrix. (Bottom) Zoom-in of the third eigenvector near the DIB feature. Guidelines for the central wavelength of the DIB and stellar lines are labeled and color coded. A higher-resolution solar spectrum from Kurucz et al. (1984) is plotted in orange for reference.

parameter space that are so poorly modeled as to be excluded from consideration in making the DIB catalog. As shown in Figure 12, only a very small fraction of the sample at the most extreme stellar parameters are excluded.

Over the majority of the stellar parameter space, the $\chi_{\text{tot}}^2/\text{dof}$ is uniform and near one. This stable region is notably much larger than that spanned by the training set, indicating good model generalization. We also checked for any trends of the inferred DIB parameters as a function of stellar parameters and found no clear correlations. However, the occasional negative EW_{DIB} detections do preferentially occur in relatively less common stellar types (in lower density regions of the left two columns of Figure 12).

We also examine the residuals for evidence of systematic mismodeling of stellar features. We sort as a function of stellar S/N and average over blocks of 32 spectra to suppress random noise and emphasize features of stellar origin (Figure 13). Clear

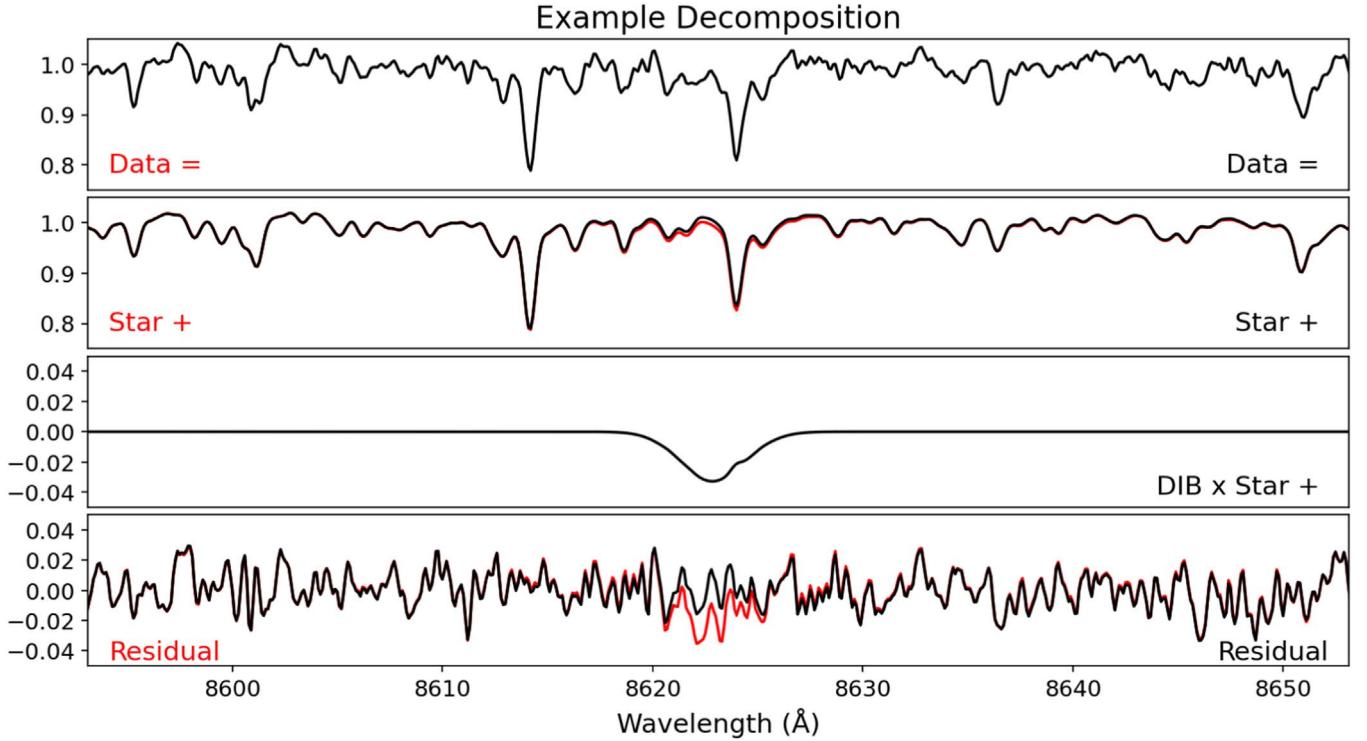


Figure 11. Example MADGICS decomposition of Gaia RVS spectrum under a star + residual (red) and star + DIB \times star + residual model (black).

features associated with the Ca triplet (8500.36, 8544.44, 8664.52 Å) appear and increase in magnitude with decreasing stellar S/N, reaching $\sim 2\%$ of the stellar component at the lowest stellar S/N.

The Ca triplet residuals resemble the derivative of the line shape and are indicative of a shift in the reported rest frame of the spectra (a shift in the apparent wavelength of the Ca triplet). As shown in Appendix I, this shift is $\sim 300 \text{ m s}^{-1}$ at low stellar S/N. The detection of this shift both demonstrates the extraordinary sensitivity of MADGICS for estimating radial velocity measures and is entirely consistent with the reported precision and magnitude dependence of Gaia RVS radial velocity measures (Katz et al. 2023). Other stellar features are visible to a lesser extent, $< 1\%$ in magnitude. Since synthetic models have residuals $\sim 5\%$ for the Ca triplet (Allende Prieto et al. 2013), the residuals shown in Figure 13 indicate high-quality modeling of the stellar component.

5. Catalog Building

5.1. Components to Catalog

Below we describe how to move from the MADGICS components to a DIB detection catalog (Figure 14). In obtaining the MADGICS components for all public spectra with the model described in Section 4, we implicitly drop spectra without full coverage in the slightly restricted spectral range of the stellar model (8469.5–8688.3 Å). These are predominantly stars with extraordinarily high radial velocities or defects in the mean RVS spectra; this cut removes only 754/999,645 spectra. We also cut DIB “detections” which occur too close to the boundary of the λ_{DIB} and σ_{DIB} parameter grids, such that the uncertainty of λ_{DIB} and σ_{DIB} cannot be estimated via a Hessian computed using finite differences. We also remove a small number of cases where the Hessian has inverted curvature along either the λ_{DIB} or σ_{DIB} direction for the step

size chosen. These are detections that do not have robust minima.²¹

Since the MADGICS DIB component need not be exactly Gaussian, we extract the Gaussian moments of the component through a simple least-squares fit of a Gaussian function. These λ_{DIB} and σ_{DIB} estimates agree quite well with the MADGICS estimates after the DIB S/N and χ_{tot}^2 cuts have been applied and there are no identifiable systematics. Before cuts, there are outliers associated with poorly modeled stellar spectra or cases where EW_{DIB} is in the noise. For that reason, we also cut “detections” where the least-squares fit λ_{DIB} and σ_{DIB} are at the edge the MADGICS search grid. After edge-of-grid cuts, 620,853 spectra remain.

We first apply the cut on $0.71 < \chi_{\text{tot}}^2/\text{dof} < 1.41$.²² The values for this cut were chosen to be a conservative definition of a “reasonably good” fit and were validated on injection tests (see Section 5.2). As shown in Figure 15, this cut can clearly separate errant detections that are locked to a specific wavelength in the stellar frame. This outlier population in $\chi_{\text{tot}}^2/\text{dof}$ also has unusually small or large σ_{DIB} . Clear outlier populations with respect to stellar S/N are also removed, the origin of which we have not specifically investigated, but appear related to the spectral normalization effects discussed in Appendix B. Note that all of these detections have a large corresponding $\Delta\chi^2$ and so would otherwise be easily mistaken as significant detections of DIBs. However, while the addition of the DIB component significantly improves the modeling of these spectra, these cases can be easily removed by recognizing that the overall goodness of fit remains poor. After the $\chi_{\text{tot}}^2/\text{dof}$ cut, 586,872 spectra remain.

²¹ Most of the above cases would be removed by the χ_{tot}^2 cut below, but we remove them first for clarity in Figure 15.

²² Throughout, we have rescaled the χ_{tot}^2 by the factor of 3.6 mentioned in Section 4.2.2 to obtain the usual interpretation of the $\chi_{\text{tot}}^2/\text{dof}$.

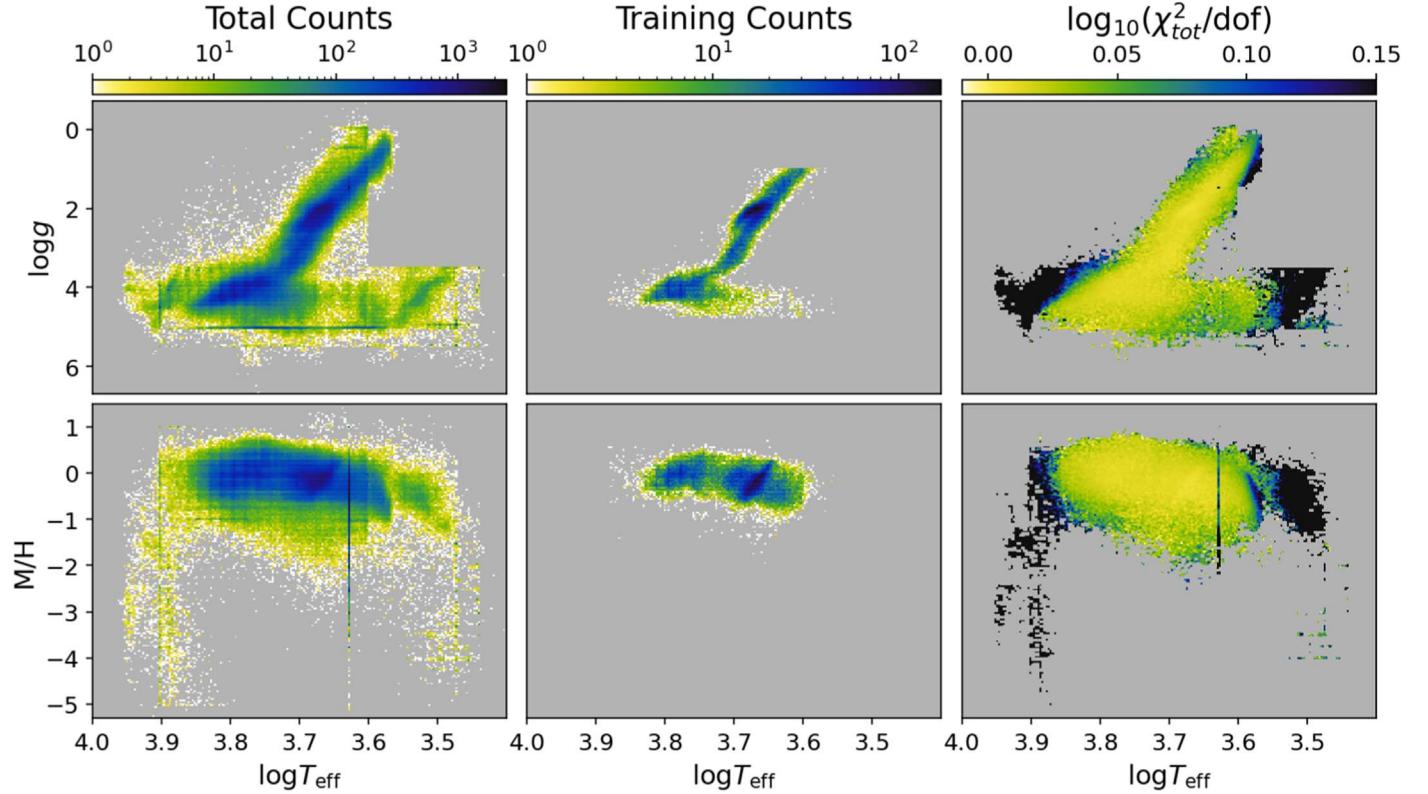


Figure 12. Histograms (2D) in the space of GSP-Spec stellar parameters for all (left) and training (middle) spectra. (Right) The overall “goodness of fit,” measured by how close χ^2/dof is to 1, is shown. Each pixel is the median of $\log_{10}(\chi^2/\text{dof})$, which is zero in the ideal case.

After these spuriously large $\Delta\chi^2$ cases have been removed, we make the final cut on $\Delta\chi^2$ such that the DIB S/N > 3.8 , where²³

$$\text{DIB SNR} = \sqrt{-\frac{\Delta\chi^2}{2}}. \quad (10)$$

The value of this cut was chosen based on injection tests and estimates of the expected false positive and false negative rates (see Appendix F). It is also clear from Figure 15 that this cut removes most of the spurious negative EW_{DIB} (emission-like rather than absorption-like) features.²⁴ After this choice of DIB S/N cut, 7789 spectra remain. For any choice DIB S/N cut, catalog impurity and noise are worse at low S/N. We view this DIB S/N cut as a lower bound that should be increased to obtain higher-quality samples as desired. The DIB S/N provides a well-motivated continuous scale on which to make such quality cuts.

The catalog created after these cuts is shown in Figure 14 as 2D histograms in the space of the moment-0, moment-1, moment-2, and DIB S/N of the detections. Clearly, the vast majority of the catalog has the expected (and likely true) positive absorption and moderate $\sigma_{\text{DIB}} \sim 1.9 \text{ \AA}$. Further, Figure 14 illustrates that spurious negative EW_{DIB} detections are only a small fraction of the catalog and are only present at low DIB S/N. This explicitly demonstrates how DIB S/N can be used as a well-motivated quality cut.

²³ This $\Delta\chi^2$ does not include the likelihood volume term associated with the DIB component.

²⁴ While DIB emission is not impossible, the only claimed example of emission related to DIB carriers is controversial and in the extreme conditions of the Red Rectangle nebula (Lai et al. 2020). Thus, DIB emission is not likely present in the sample considered here.

5.2. Injection Tests

To quantify biases and validate reported uncertainties in both the component separation and catalog creation steps, we created a set of injection tests. To generate these tests we start from a “parent” set of sources along lines of sight with low reddening, $\text{SFD} < 0.05 \text{ mag}$ (Schlegel et al. 1998), and with high-quality upstream processing (GSP-Spec flags 1–13 equal to zero). Then, we randomly draw a position, width, and amplitude (λ_{DIB} , σ_{DIB} , α_{DIB}) as well as an index of a “parent” spectrum for the synthetic DIB. Each of these parameters is uniformly sampled over a wide range: $\lambda_{\text{DIB}} \in [8618.47, 8628.47] \text{ \AA}$ $\sigma_{\text{DIB}} \in [0.6, 3.1] \text{ \AA}$ $\alpha_{\text{DIB}} \in [0.25, 0.0] \text{ \AA}$.

By uniformly sampling over “parent” spectra, the injection tests approximately follow the distribution of stellar S/N in the overall sample (though the SFD and GSP-Spec flag requirements modify the distribution). Since DIB absorption is multiplicative, absorbing the background starlight, to obtain the additive DIB component we must multiply the Gaussian profile with the drawn parameters by an estimate of the stellar component for the “parent” source. Since we have a stellar component estimate from a star+residuals model (see Section 4), we use that stellar component for the “parent” source. Then, the synthetic spectrum is the original parent spectrum plus the new DIB component, and we leave the “parent” per-pixel uncertainties unchanged.

These injection tests are then subjected to the same pipeline, component separation, and catalog creation cuts, as the real spectra. Given a choice of what a “good” detection is, we can compute a measure of the catalog completeness as a function of DIB parameters (see Appendix F). Since the injection tests have a ground truth, we can compute Z-scores, the predicted

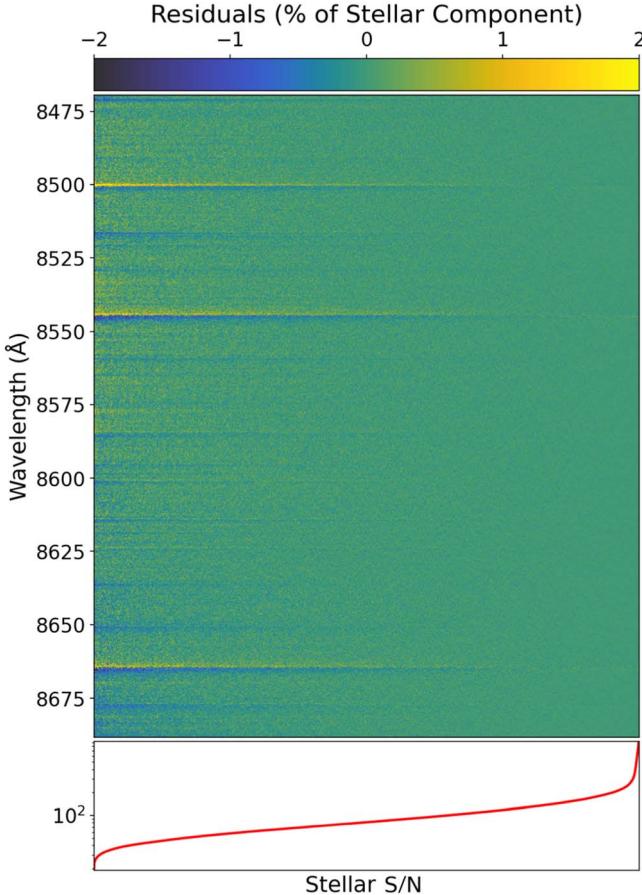


Figure 13. Residual components, sorted by stellar S/N, divided by the stellar component, and averaged over 32 adjacent spectra. The bottom panel shows the value of stellar S/N along the sorted axis.

minus true value, divided by the reported uncertainty. When the width of that distribution is one, the reported uncertainties are correctly estimated. Wider than one indicates underestimation and narrower than one indicates overestimation of the reported error bars.

Figure 16 shows the Z-score distributions for the three DIB moments (EW_{DIB} , λ_{DIB} , σ_{DIB}) as a function of the ground truth parameters, stellar S/N, and DIB S/N. As a function of injected EW_{DIB} , the estimated EW_{DIB} Z-score distribution is close to uniform, except at low injected EW_{DIB} where a positive bias is observed. The cause of this bias is the imposition of a detection threshold, as evidenced by the stronger bias at low DIB S/N (see Section 5.3 for more). The EW_{DIB} Z-score distribution is almost perfectly uniform with respect to the injected λ_{DIB} and σ_{DIB} . There is little bias in the EW_{DIB} estimate and error bars as a function of stellar S/N where there is a high density of samples and moderate stellar S/N. At low stellar S/N, there is a $\sim 0.5\sigma$ positive bias and underestimation of the error bars. This is likely because spectra with low stellar S/N tend to have lower DIB S/N.

The estimated λ_{DIB} distribution is almost perfectly uniform with respect to injected EW_{DIB} , λ_{DIB} , σ_{DIB} , and stellar S/N. As a function of DIB S/N, it is uniform except for at the lowest S/N, where there is a small ($\leq 50\%$) underestimation of the error bars.

The estimated σ_{DIB} distribution is almost perfectly uniform with respect to injected EW_{DIB} , λ_{DIB} , and σ_{DIB} , but shows a slight ($\leq 0.25\sigma$) positive bias at large injected EW_{DIB} and low

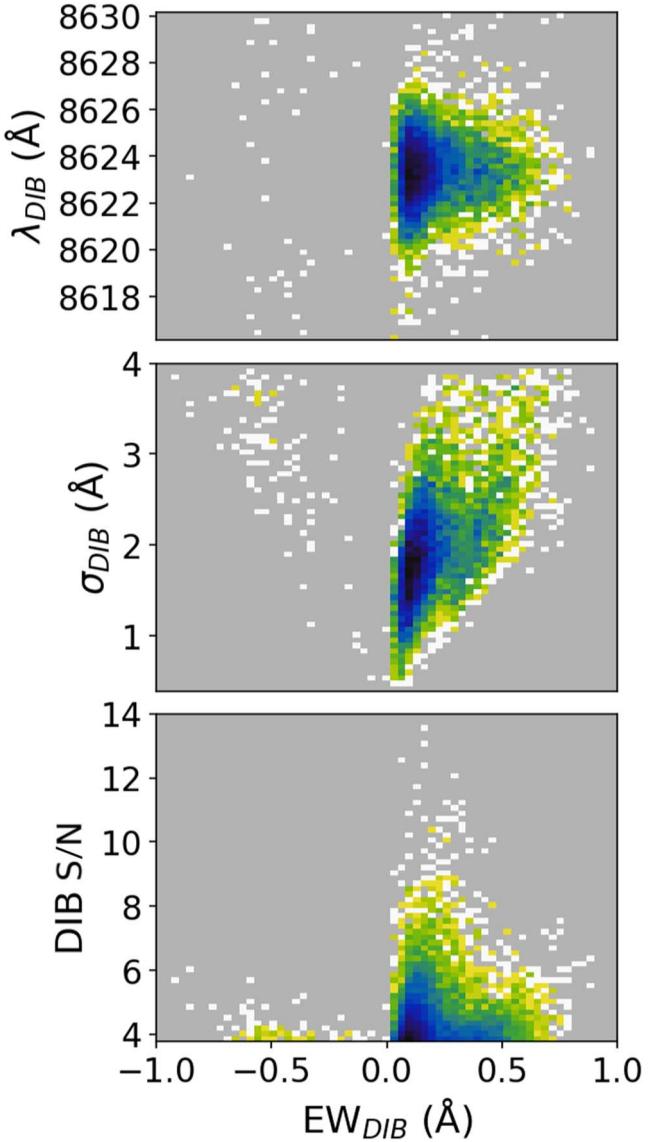


Figure 14. Histograms (2D) showing projections of the MADGICS catalog in the space of moment-0, moment-1, moment-2, and DIB S/N. The color scale indicates density, where darker colors indicate more detections, and each panel has its own logarithmic stretch.

injected σ_{DIB} . This is the result of a detection bias on σ_{DIB} , which is less commonly known compared to the EW_{DIB} (or flux) detection bias (see Section 5.3). At low DIB S/N, and low stellar S/N which tend to have lower DIB S/N, a larger positive bias on the order of 0.5σ is observed.

Figure 16 is intended primarily as a validation of the method and gives a slightly optimistic view of biases in the final catalog because it uses injections performed on high-quality spectra. An injection test releasing the GSP-Spec flag condition is explored in Appendix F. With this sample that is more representative of the spectral diversity in the catalog, errors can be $\sim 10\%$ underestimated and biases are slightly elevated (see Section 6).

Overall, the injection tests allow us to demonstrate impressive uniformity in the recovered Z-score distributions which lends confidence to the parameter estimates and uncertainties in the MADGICS DIB catalog.

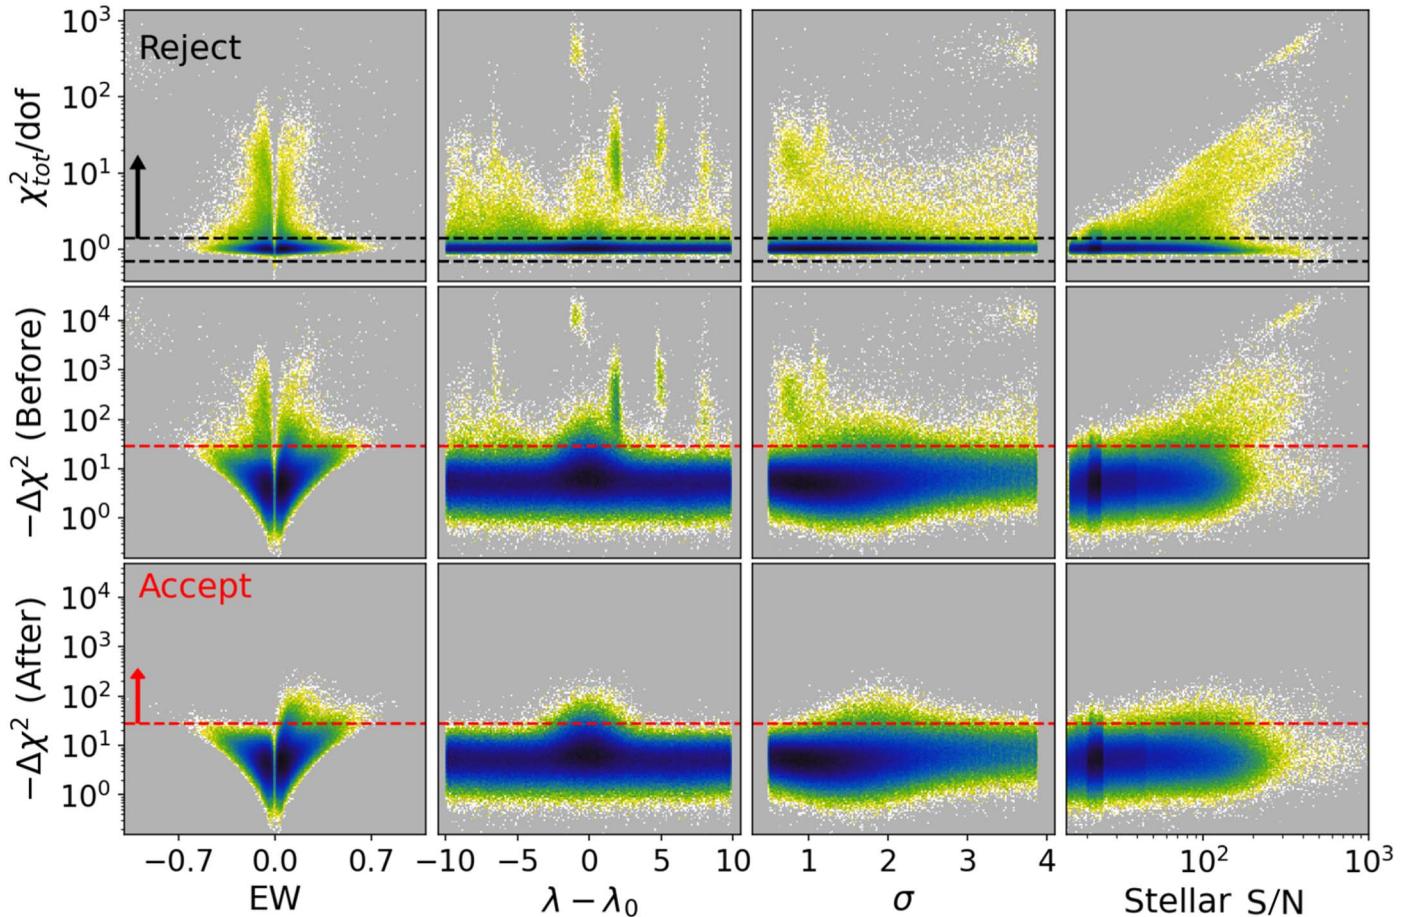


Figure 15. Histograms (2D) for DIB models as a function of EW_{DIB} , λ_{DIB} , σ_{DIB} , and stellar S/N in the columns, left to right respectively. The λ_{DIB} is reported relative to the reference center of the MADGICS parameter grid. Positive EW_{DIB} indicates absorption. The rows show this distribution for (top) the total χ^2 per degree of freedom of the component fit, the negative change in χ^2 as a result of adding the DIB component before (middle) and after (bottom) imposing cuts on χ^2_{tot}/dof . Guidelines indicate cuts imposed in each space. The color scale indicates density, where darker colors indicate more detections, and each panel has its own logarithmic stretch.

5.3. Low S/N Behavior

Injection tests subjected to the same detection cuts as the catalog provide an empirical measure of the biases imposed by the creation of a catalog (Figure 17). Understanding this detection bias function aids in interpreting population-level results from the catalog and can be incorporated into downstream inference. Inference wishing to avoid biases imposed by the detection cuts should work further upstream of the detection cuts (components of all spectra are available in Section 6).

Toward low true DIB S/N, we see the expected positively diverging detection bias causing an overestimation of EW_{DIB} . The σ_{DIB} has an overall envelope of increasing magnitude toward low true DIB S/N. However, the σ_{DIB} goes through a sign change for sufficiently large true σ_{DIB} . This sign change is driven primarily by the implicit prior we have imposed by the MADGICS parameter search over a fixed grid of σ_{DIB} . This grid truncates the larger sigma wing of the true posterior for σ_{DIB} , causing a negative σ_{DIB} bias. From high S/N measurements, we have a physical prior expectation that most true σ_{DIB} values lie within 1.5–2.3 Å for which the σ_{DIB} bias is predominantly positive.

The form of the low S/N bias can be reproduced theoretically only by combining the effects of both the fixed $\Delta\chi^2$ detection cut and bias correction for using parameter estimates in a Gaussian formalism for a likelihood that has a

nonlinear dependence on those parameters (Zanolin et al. 2010; Portillo et al. 2017). This is because we are approximating as Gaussian the parameter posteriors that are non-Gaussian due to that nonlinearity. The quantitative theoretical model will be described in a future work.

6. Data/Code Availability

The final catalog, code that implements MADGICS, and several small summary files are provided via doi: [10.5281/zenodo.7388333](https://doi.org/10.5281/zenodo.7388333) (2 GB). Jupyter notebooks that generate the priors, injection tests, and the catalogs as well as reproduce all figures are also provided. Individual files contained in the Zenodo, including the final catalog alone (3 MB), are available from the project website.²⁵ The quality of the final catalog can be further refined by imposing a cut on the DIB S/N field. The full set of intermediate data products, including the full star+residual components, star+DIB+residual components, and injection tests are also available from the project website (500 GB total).

7. Conclusion

In this work, we introduced a new technique based on Marginalized Analytic Data-space Gaussian Inference for

²⁵ <https://faun.rc.fas.harvard.edu/saydjari/GaiaDIB/>

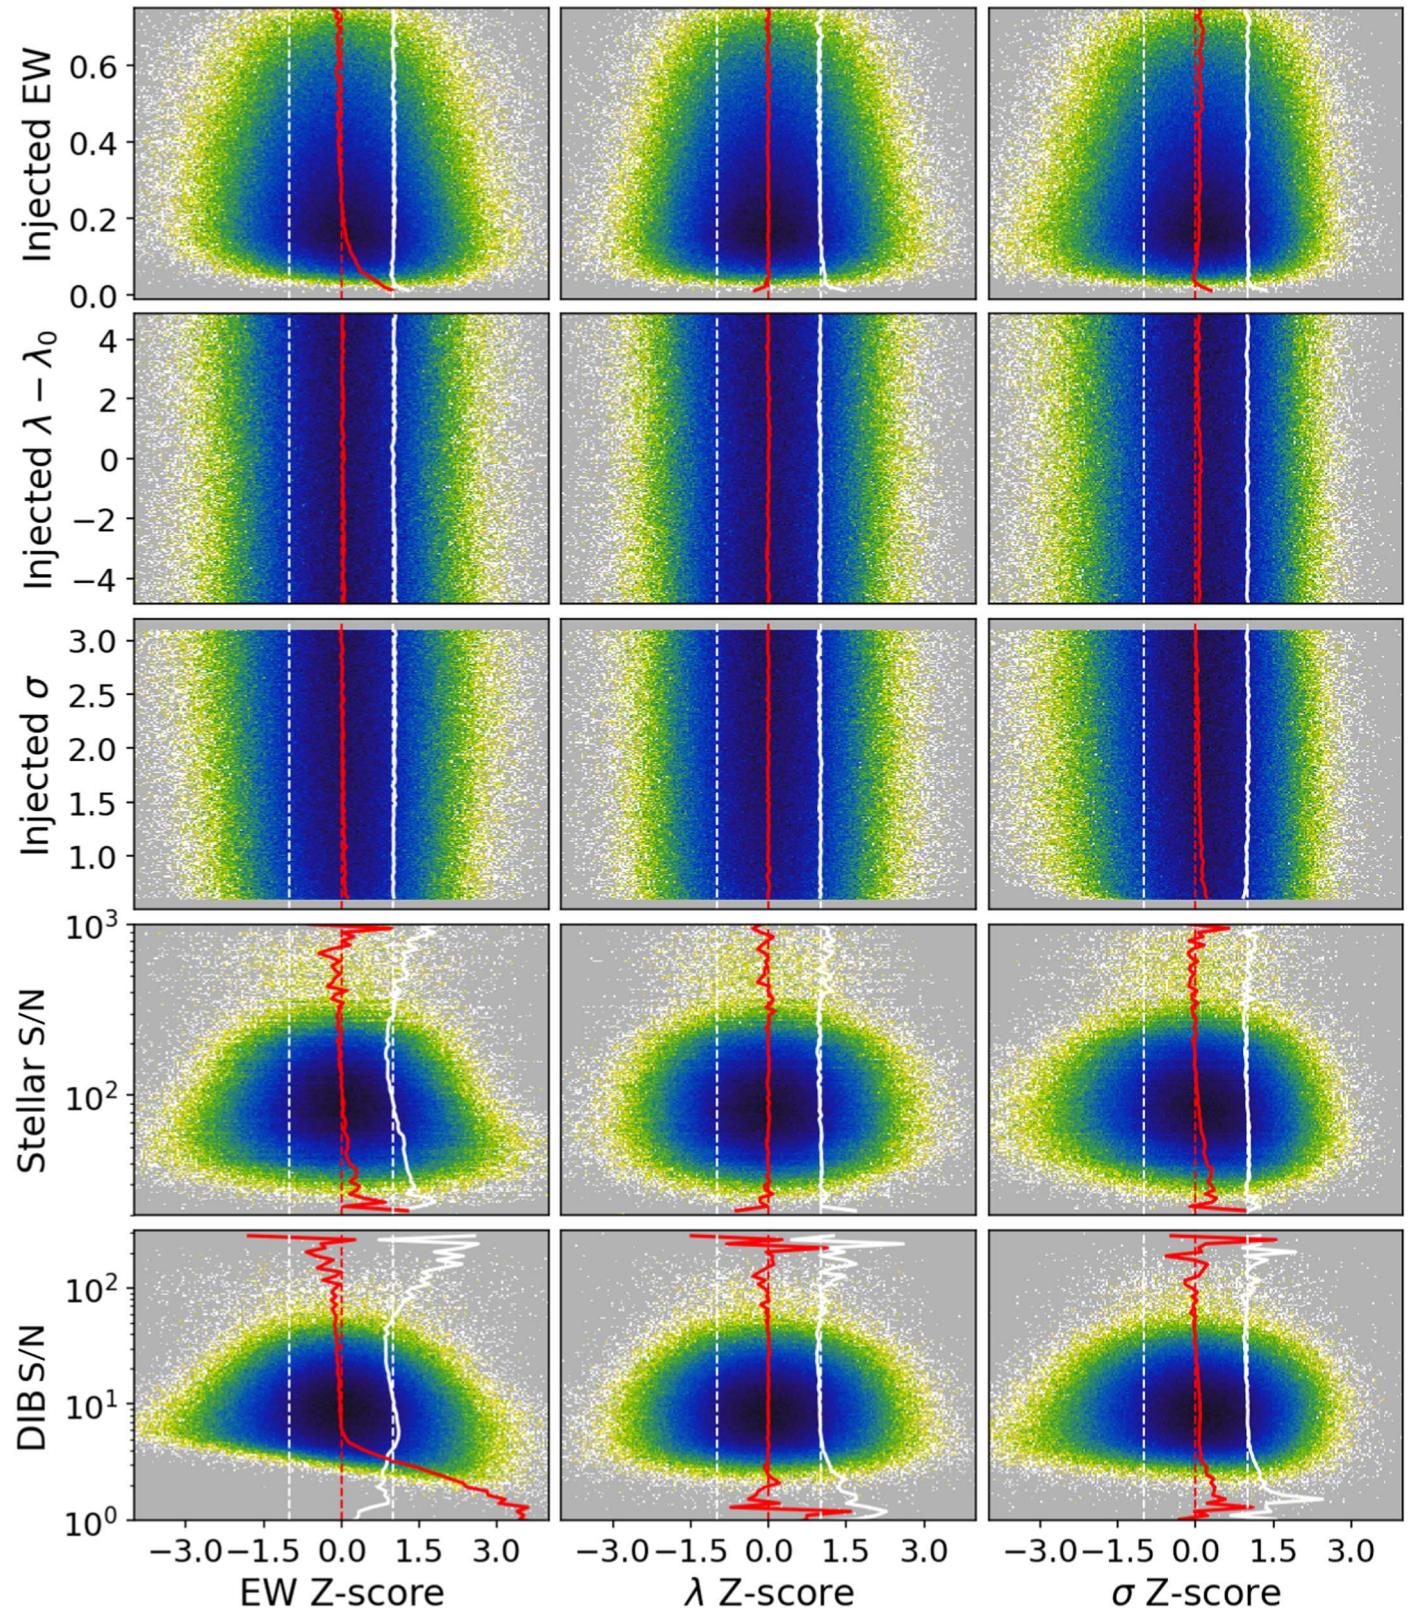


Figure 16. Histograms illustrating the bias of DIB detections relative to reported uncertainties (Z-scores) for injection tests where ground truth is known. Histograms are a function of injected EW_{DIB} , injected λ_{DIB} , injected σ_{DIB} , stellar S/N and DIB S/N in the rows, and show the bias in the recovered quantities EW_{DIB} , λ_{DIB} , σ_{DIB} in the columns (left to right). The λ_{DIB} is reported relative to the reference center of the MADGICS parameter grid. The color scale indicates density, where darker colors indicate more detections, and each panel has its own logarithmic stretch. Dashed guidelines provide a reference for 0 and $\pm 1\sigma$. Solid lines indicate the center (red) and spread (white) of the Z-score distribution.

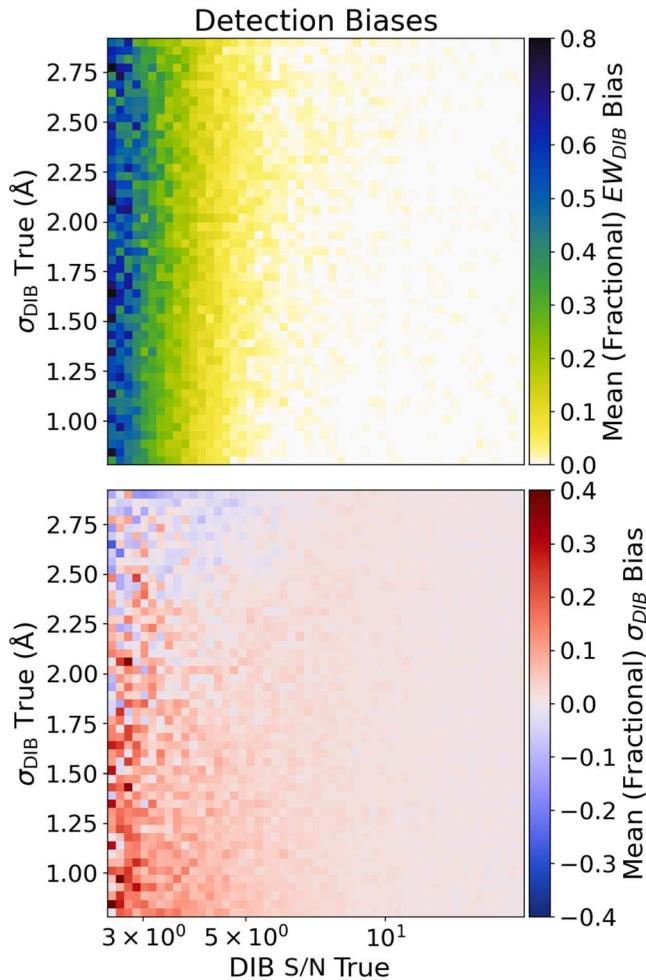


Figure 17. Fractional detection bias determined by injection tests for EW_{DIB} (top) and σ_{DIB} (bottom) as a function of the true DIB S/N and true σ_{DIB} .

Component Separation (MADGICS) to detect and measure the properties of DIBs in stellar spectra. Using MADGICS, we pushed detections to lower stellar S/N and marginalized over stellar features, resulting in a new Gaia RVS 8621 Å DIB catalog free from detectable stellar contamination. In this catalog, we find no evidence of DIB detections within the Local Bubble.

We present and validate a 4D map of ISM kinematics based on the first moment of the DIBs. As part of this work, we adjusted the measured rest-frame wavelength to 8623.14 ± 0.087 Å with a precision accounting for modeling uncertainties. Further, we show an unprecedented correlation of the DIBs with kinematic substructure in Galactic CO maps. An interpretable and physically realistic trend in the second moment of the DIBs was also presented. Rigorous validation of the catalog, its reported uncertainties, and the choices made in constructing it are provided via a series of synthetic injection tests. This allows modeling of the low DIB S/N detection biases in both EW_{DIB} and σ_{DIB} .

In future work, MADGICS could be extended to more complicated spectral component separation tasks, including identifying multiple stellar components (and the associated radial velocities), removing sky emission lines (from ground-based observations), and modeling DIB detections with

multiple cloud components (velocities) along the line of sight. Looking forward to increased source densities in future data releases, larger high-quality 4D kinematic maps of the ISM will enable detailed studies of ISM substructure and dynamics (Tchernyshyov et al. 2018).

Acknowledgments

A.K.S. gratefully acknowledges support from a National Science Foundation Graduate Research Fellowship (DGE-1745303). D.P.F. acknowledges support by NSF grant AST-1614941, “Exploring the Galaxy: 3-Dimensional Structure and Stellar Streams.” D.P.F. acknowledges support by NASA ADAP grant 80NSSC21K0634 “Knitting Together the Milky Way: An Integrated Model of the galaxy’s Stars, Gas, and Dust.” C.Z. acknowledges that support for this work was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51498.001 awarded by the Space Telescope Science Institute (STScI), which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. The support and resources from the Center for High Performance Computing at the University of Utah are gratefully acknowledged. A portion of this work was also enabled by the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

This work was supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions). The authors acknowledge Interstellar Institute’s program “With Two Eyes” and the Paris-Saclay University’s Institut Pascal for hosting discussions that nourished the development of the ideas behind this work.

We acknowledge helpful discussions with Mathias Schultheis, He Zhao, Morgan Fouesneau, Jos de Bruijne, Patrick de Laverny, Gordian Edenhofer, Torsten Enßlin, Gail Zasowski, Joel Brownstein, Tom Dame, Greg Green, Eddie Schlaflly, Jonathan Bird, Adam Wheeler, Andy Casey, and Phillip Cargile. A.K.S. acknowledges Sophia Sánchez-Maes for helpful discussions and much support.

This work has made use of data from the European Space Agency (ESA) mission Gaia (<https://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement. This research has made use of the VizieR catalog access tool, CDS, Strasbourg, France (Ochsenbein et al. 2000).

Facility: Gaia.

Software: Julia (Bezanson et al. 2017), FITSIO.jl (Pence et al. 2010), HDF5.jl (The HDF Group 1997–2022), Healpix.jl (Tomasi & Li 2021), BenchmarkTools.jl (Chen & Revels 2016), ASTROPY (Astropy Collaboration et al. 2013), IPYTHON (Perez & Granger 2007), MATPLOTLIB (Hunter 2007), SCIPY (Virtanen et al. 2020).

Appendix A Public RVS Selection Function

We illustrate the distribution of public RVS spectra in Figure 18. The highly heterogeneous distribution on the sky

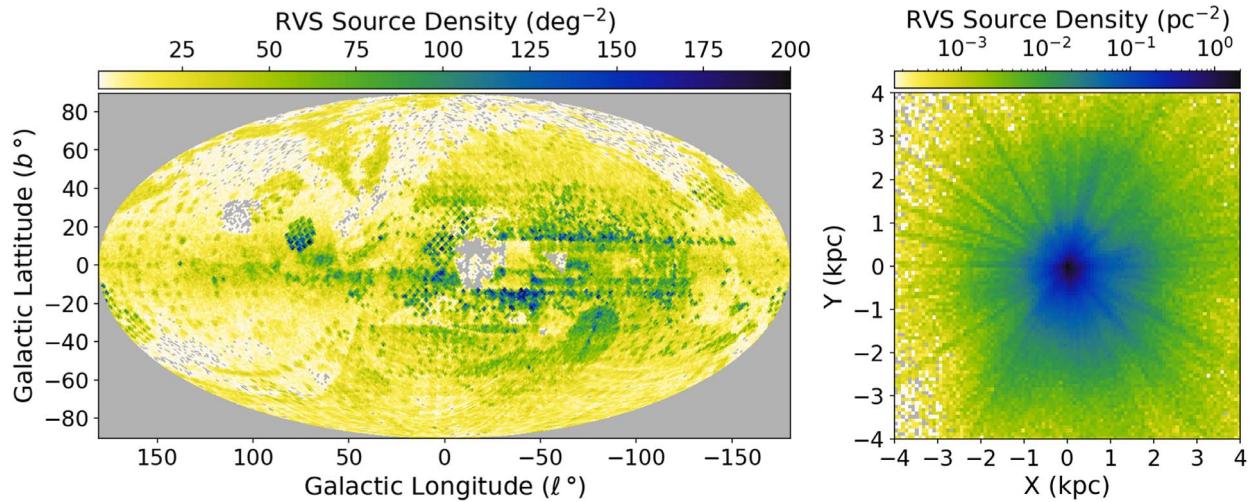


Figure 18. The density of public RVS spectra on the sky (left, HEALPix grid at NSide=64) and projected onto the Galactic plane in Cartesian coordinates (right).

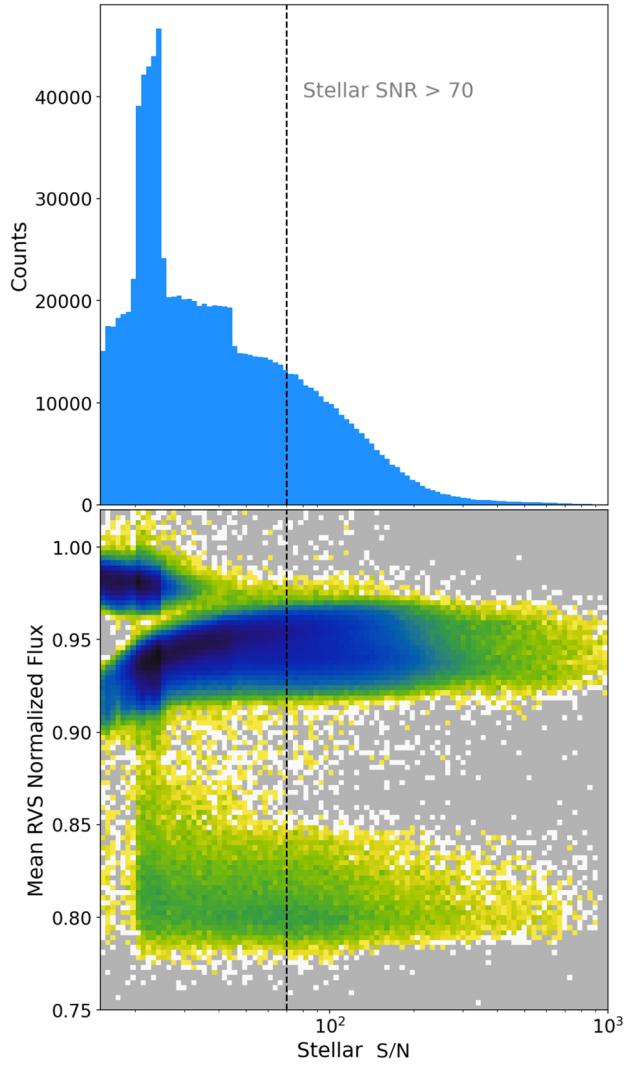


Figure 19. (Top) Histogram of public Gaia RVS sources with respect to stellar S/N. (Bottom) 2D histogram of public Gaia RVS sources with respect to stellar S/N and the mean normalized flux in the RVS spectrum. The color scale indicates density, where darker colors indicate more sources, and is a logarithmic stretch. A guideline is shown at stellar S/N 70, which is the cutoff imposed by the Gaia DIB catalog.

(shown in Galactic coordinates) shows clear patterns associated with the footprints of other surveys, such as Kepler (Borucki et al. 2010), Galactic Archaeology with Hermes (GALAH, Buder et al. 2021), the Large sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST, Luo et al. 2015), and the Apache Point Observatory Galactic Evolution Experiment (APOGEE, Majewski et al. 2017). We understand this to be because the majority of the public RVS spectra were those used for validation by comparing to other spectroscopic measurements of the same source (G. Seabroke et al. 2022, in preparation; Vallenari et al. 2023).

Heterogeneity is also seen with respect to the stellar S/N, as shown in Figure 19. Plotting as a function of stellar S/N also highlights the significant heterogeneity in the normalization of the RVS spectra as a result of a discrete choice in the Gaia pipeline to either pseudocontinuum normalize or rescale the spectra by a constant (in the case of cool stars or low stellar S/N; Vallenari et al. 2023).

Appendix B DIB S/N Distribution

It is useful to describe the MADGICS DIB catalog detections as a function of S/N. At the top of Figure 20, we show the cumulative detections as a function of the reported DIB S/N. This provides an idea of the number of detections remaining in the catalog for some choice of quality cut using the DIB S/N. At the bottom of Figure 20, we show the 2D histogram of the detections as a function of both DIB S/N and stellar S/N. Clearly, many high-quality detections (large DIB S/N) were found below the stellar S/N cutoff imposed by the Gaia DIB catalog. Because a large fraction of the spectra have low S/N, it is important to leverage their statistical power. The peak at stellar S/N ~ 20 is simply a reflection of the unusually large number of spectra at that stellar S/N in the public Gaia DR3 RVS spectra. We have confirmed that the fraction of DIB detections passing our DIB S/N cut actually decreases as a function of stellar S/N. This lends confidence that MADGICS is correctly finding the large DIB features which have high DIB S/N even on low stellar S/N spectra, without overfitting noise as DIB detections.

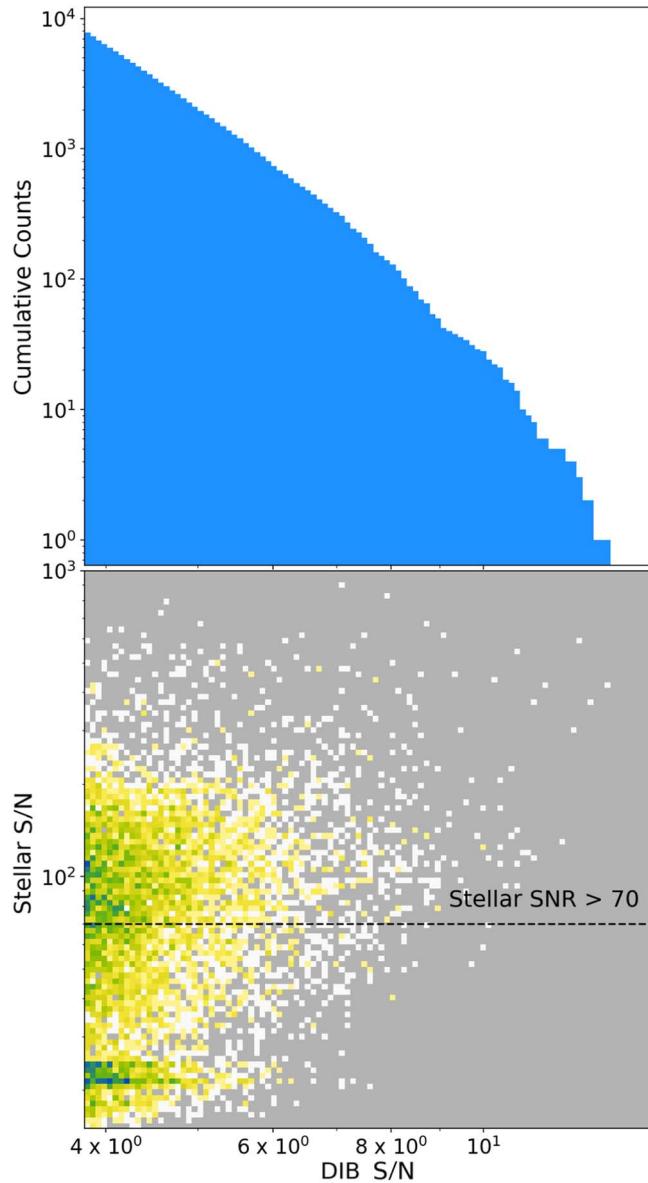


Figure 20. (Top) Cumulative counts of MADGICS DIB detections as a function of DIB S/N. (Bottom) 2D histogram of MADGICS DIB detections as a function of DIB and stellar S/N. The color scale indicates density, where darker colors indicate more sources, and is a logarithmic stretch. A guideline is shown at stellar S/N 70, which is the cutoff imposed by the Gaia DIB catalog.

Appendix C Removing Goodness of Fit Cuts

As stated in the main text, we aim to (1) impose goodness of fit cuts that eliminate any influence of stellar lines and (2) use methods that are maximally robust to stellar mismodeling so that real DIBs can be detected and well modeled in the largest number of spectra possible. To help compare the catalogs with respect to (2), we replicate Figure 1 without any goodness of fit cuts imposed. Without goodness of fit cuts, MADGICS would return detections in 19,071 spectra compared to the 50,787 detections in the Gaia Full catalog on public spectra.

Clearly, there are detections associated with several of the stellar lines identified in Figure 1 within the MADGICS detections when the goodness of fit cuts are not imposed. This is to be expected. When the star is severely mismodeled, removing the stellar component and finding the broad multiplicative absorption of the DIB is nearly impossible. The question is then “How robust is the DIB identification to the stellar mismodeling?” The lower density of detections associated with stellar features in the MADGICS catalog panel of Figure 21 suggests that the MADGICS approach is also more robust (2), in addition to having the high-quality goodness of fit cut (1) demonstrated in Section 3.1.

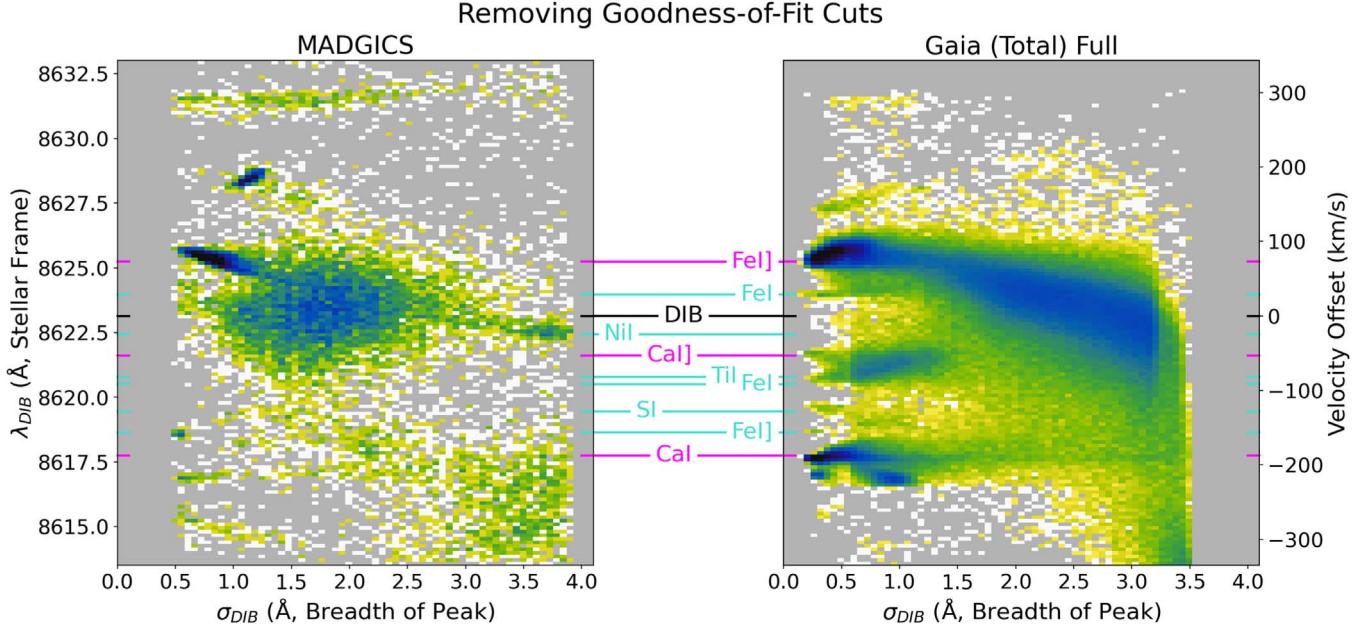


Figure 21. The density of detections in the MADGICS (left) and Gaia (right) 8623 \AA DIB catalog as a function of λ_{DIB} and σ_{DIB} , removing cuts imposed on the goodness of fit of the model to the overall spectrum. This removes the second block of cuts in Table 1 for both catalogs. The Gaia catalog shown is the full DIB sample, including DIB detections in the public Gaia DR3 DIB catalog that were identified on spectra that are not public. The color scale indicates density, where darker colors indicate more detections, and is a relative logarithmic stretch for both catalogs, with the maximum rescaled by relative catalog sizes. Guidelines for the central wavelength of the DIB and stellar lines are labeled and color coded between the plots.

Appendix D Dusty Stellar Prior

It is difficult to prove that a data-driven stellar model (the stellar covariance herein) does not have any small contributions from low EW_{DIB} lines of sight that are present in the training data despite the cut on $\text{SFD} < 0.05$. However, we can easily show the impact of the inverse, where we loosen the cut to $\text{SFD} < 0.1$ during the construction of the stellar prior, thereby including much dustier lines of sight that are likely to have more contributions from low EW_{DIB} DIBs. We then use this “dusty” stellar prior to detect DIBs on the synthetic injection tests described in the main text, and show the resulting change in the DIB catalog in Figure 22.

Figure 22 shows the median change in EW_{DIB} as a function of the true EW_{DIB} and true DIB S/N of the injected DIBs. We take this difference between DIB detections appearing in both the original and “dusty stellar prior” DIB catalogs built from the injection tests. Regions of parameter space where $>25\%$ of detections in the original catalog were not refound are colored green. Only a very small number (557/1,000,000) of injections were detected in the “dusty stellar prior” catalog, but not the original and are excluded from this plot. Regions of parameter space where no DIBs were detected in either catalog are shown in light gray, and regions not probed by the injection tests are shown in dark gray.

Over the majority of the detections, there is a <0.005 \AA reduction in the reported EW_{DIB} . Just near the detection threshold, a <0.005 \AA increase is observed. This can be interpreted as a shift in the detection threshold carrying with it the usual positive detection bias (see Section 5.3). For moderate EW_{DIB} (>0.15 \AA), there is very little change in the detectability of the synthetic injections, even at low DIB S/N. At low EW_{DIB} (<0.15 \AA), there is a clear increase in the true DIB S/N required in order for an injected DIB to be detected. This makes sense as some small contribution of the DIB peak is present in

the stellar model, though averaged over the relative velocity distribution between the star and dust. However, almost all of the DIBs that are lost have a true DIB S/N below the DIB S/N cut imposed in making the catalog, and thus were only detected previously because they fluctuated high. Further, it is important to note the persistence of low EW_{DIB} detections at higher DIB S/N.

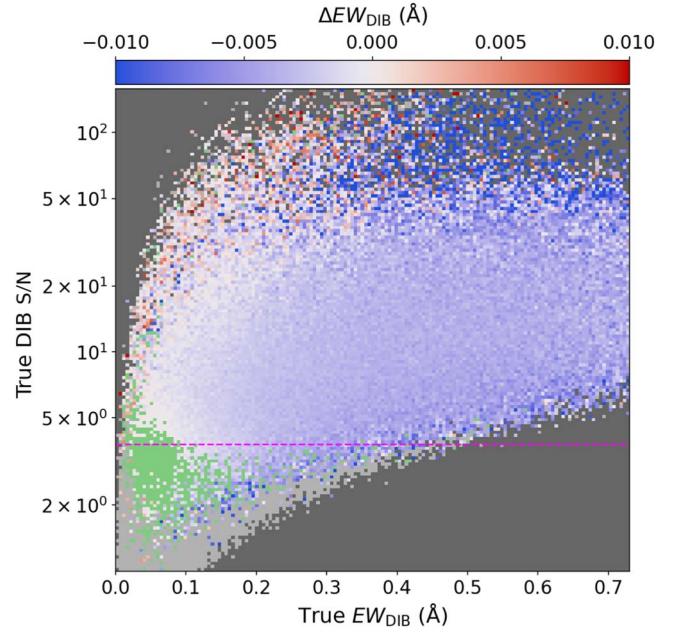


Figure 22. Median change in EW_{DIB} as a function of EW_{DIB} and DIB S/N after including dustier lines of sight in the training data for the stellar covariance. Green indicates $>25\%$ of detections not found, gray indicates detections not found with the clean stellar covariance, and dark gray indicates no injection tests present. The guideline (magenta, dashed) indicates the detection cut on the estimated DIB S/N imposed in making the MADGICS catalog.

Resolving the question of the presence of DIBs in the Local Bubble is important. In this work, we show that approaches with synthetic models can incorrectly identify stellar features as DIBs. However, it is difficult to prove that a data-driven stellar model is completely free of contributions from low EW_{DIB} DIBs. Yet, Figure 22 makes it clear that data-driven stellar models can solve this question if we push to higher stellar S/N spectra, since we can then push confident DIB detections to lower EW_{DIB} limits.

Appendix E Dust–CO Correlation

To account for the clumpiness of CO, we find the nearest (in velocity) peak in CO to the DIB peak before calculating the intensity-weighted mean CO velocity v_{CO} that corresponds to that DIB. First we restrict to only DIBs within the footprint of the Dame et al. (2001) CO survey (3602). Further, we restrict to DIBs such that there is at least some detection of CO within 1σ of the DIB velocity (2272). Then, we find the nearest connected set of velocity channels with CO detections above the threshold, where the threshold starts at zero, and the threshold increases in steps of 0.05 until the connected region (island) is less than 9 velocity pixels wide. After this stopping condition is reached, we take the intensity-weighted average over the island to obtain v_{CO} .

While such a peak-finding algorithm will necessarily overestimate the correlation between the CO and DIBs, we find a similar correlation when restricting to low-velocity CO without any peak finding, only making the two broad cuts outlined above (Figure 23). Computing v_{CO} as the intensity-weighted velocity over all velocity channels, we restrict to $|v_{CO}| < 15 \text{ km s}^{-1}$ to reduce the multimodality of the CO and DIB distributions. Then, we find a slope of 0.83 ± 0.04 and an intercept of $-0.51 \pm 0.3 \text{ km s}^{-1}$. Both of these values are consistent with those discussed in the text, though the slope is slightly flatter and there is larger scatter due to the additional averaging in v_{CO} . However, the scatter of residuals/ $\sigma(v_{DIB}^{\text{LSR}})$, or

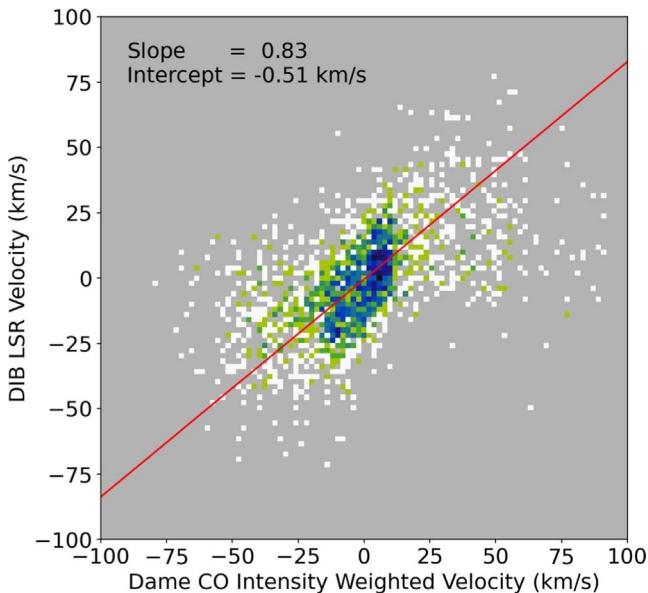


Figure 23. Histogram (2D) illustrating the correlation between the v_{DIB}^{LSR} and intensity-weighted velocity of CO over all velocity channels in the Dame et al. (2001) CO map. The color scale indicates density, where darker colors indicate more detections and a logarithmic stretch has been applied.

Z-scores, was 0.96, indicating that the reported uncertainties well describe the scatter in the residuals. Thus, we conclude the peak finding to determine v_{CO} does not significantly change the conclusions and it helps illustrate that strong correlations are present for high-velocity DIBs and CO after accounting for the multimodality.

Appendix F Purity and Completeness

Combining injection tests with a choice for what a “good” detection is allows an estimate of the completeness of the catalog (Figure 24). Here, we take an arbitrary choice for good detections to be $|\Delta\alpha_{DIB}|/\alpha_{DIB} < 0.2$, $|\Delta\lambda_{DIB}| < 2.5 \text{ \AA}$, $|\Delta\sigma_{DIB}|/\sigma_{DIB} < 0.2$, where the Δ indicates differences with respect to the ground truth values for the injected DIB profile. We show this completeness as a function of both EW_{DIB} and σ_{DIB} to illustrate that it is more difficult to detect wider DIBs given the same EW_{DIB} . For $EW_{DIB} \geq 0.2 \text{ \AA}$ the catalog is $\geq 80\%$ complete over the range of relevant σ_{DIB} here. Even though completeness is reduced at lower EW_{DIB} , we still achieve $\sim 20\%$ completeness at $0.05\text{--}0.1 \text{ \AA}$.

In creating Figure 24, we are treating the detection of synthetic injections as “true positives” (TP). We can leverage this to help motivate a choice for the $\Delta\chi^2$ cut used to make the catalog. To do this, we compute the TP rate as a function of the $\Delta\chi^2$ cut chosen, restricting to $\alpha_{DIB} > 0.07$ as motivated by Figure 24. Similarly, we can use the training set used for generating the stellar prior (low reddening, high stellar S/N, no GSP-Spec flags) as a set of ground truth negatives. Then, “detections” of a DIB on those spectra give an idea of the “false positive” (FP) rate. We can then produce the FP rate as a function of the $\Delta\chi^2$ cut chosen, again restricting to $\alpha_{DIB} > 0.07$.

Both of these reference data sets are somewhat idealized. To better capture how the heterogeneity in the full data set impacts the TP rate, we created a synthetic injection test using a

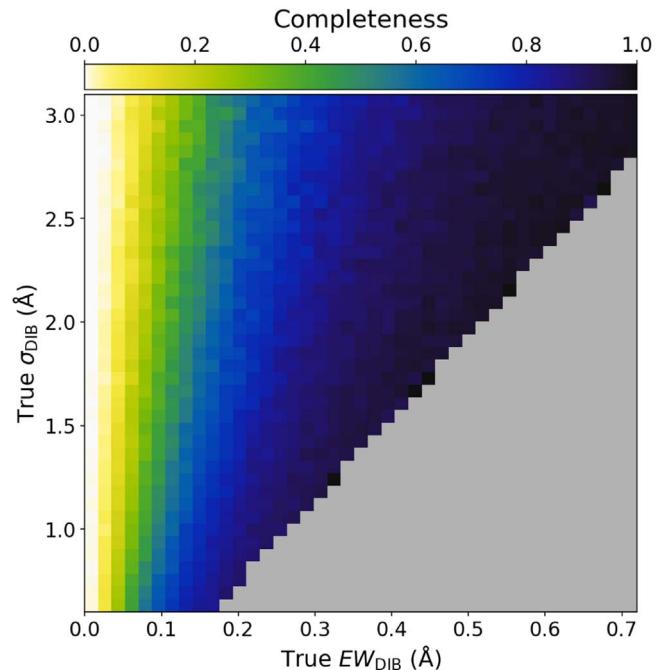


Figure 24. Estimated completeness of DIB catalog as a function of EW_{DIB} and σ_{DIB} based on quality cuts on injection tests.

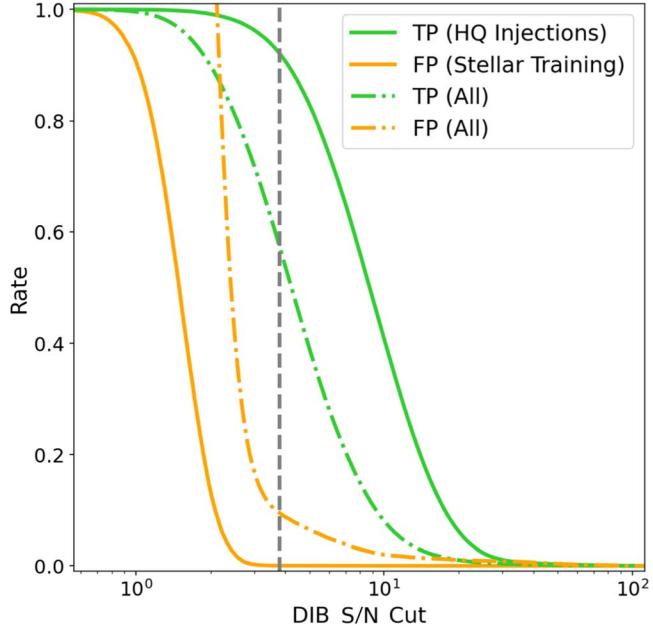


Figure 25. Estimated true positive (TP) and false positive (FP) rates as a function of DIB S/N cut using the injection tests and stellar training set, respectively (solid lines). Estimated TP and FP rates using a broader set of low reddening injection tests/spectra are shown in dotted-dashed lines. Vertical guideline illustrates the value chosen in this work (gray, dashed).

“parent” set that only requires SFD <0.5 mag (low dust) and that the spectra not have been used in the stellar prior training set. Similarly, we can estimate a more realistic FP rate using all low reddening spectra (without an injection step).

In Figure 25 we show these TP and FP rates as a function of the $\Delta\chi^2$ cut chosen, which we have converted into a cut on the DIB S/N. The value used in this work is shown by a dashed vertical line. This corresponds to (FP, TP) = (0.04%, 92%) for the high-quality spectra tests and (FP, TP) = (9%, 57%) for tests more representative of the average spectra in the Gaia DR3 RVS release. The exact choice of the test sets and definitions of TP and FP will modify the exact rates, but Figure 25 demonstrates the qualitative choice we have made to optimize for low false positive rates.

Appendix G Line Labels

As discussed in Section 4.2.1, we labeled the pileups in the stellar rest-frame wavelength found in the Gaia DIB catalog by the transition at those locations with the largest amplitude in models of the solar spectrum using the Kurucz archive (see footnote 7). In Table 2 we list these lines along with the set of stellar lines in the GSP-Spec model carefully calibrated by Contursi et al. (2021a) in the wavelength range near the DIB feature. Of the lines that were not calibrated as part of the GSP-Spec modeling, we were able to confirm that two of the Fe I] lines were part of the (private) GSP-Spec atomic line list (private communication, P. de Laverny, 2023). However, the molecular line list is proprietary, and thus the inclusion of the ^{13}CN lines in the GSP-Spec line list remains unknown to us (private communication, P. de Laverny, 2023). However, we reiterate that regardless of the inclusion of any given line in the line list, the Gaia DIB catalog clearly has contamination from poorly modeled stellar features centered on these wavelengths.

Table 2
Stellar Line Labels

Species	Wavelength (Å)	$\log(gf)$	Included?
			Calibrated GSP-Spec Lines
Fe I	8614.169	-1.850	Y
Fe I	8616.306	-1.246	Y
Fe I]	8618.647	-0.707	Y
S I	8619.457	-0.763	Y
Fe I]	8620.508	-1.868	Y
Ti I	8620.792	-1.980	Y
Ni I]	8622.439	-1.376	Y
Fe I	8623.969	-2.320	Y
S I	8628.911	-0.894	Y
Si I]	8631.971	-3.168	Y
Lines Not Calibrated or Included in GSP-Spec			
^{13}CN	8615.040	-1.345	?
Fe I	8617.678	-0.950	N
Fe I]	8621.491	-1.330	N
^{13}CN	8621.732	-1.269	?
Fe I	8625.112	-0.770	N
Fe I]	8626.109	-2.089	Y
Fe I]	8627.260	-3.177	Y
^{13}CN	8628.620	-1.208	?
^{13}CN	8631.382	-1.104	?

Appendix H Example Full Posteriors

As discussed in Section 4.1, we obtain \hat{C}_{kk} , the predicted pixel-pixel covariance of pixels in component k , and \hat{C}_{km} , the predicted covariance of a pixel in component k with a pixel in component m . In Figure 26, we show a subset of these covariance matrices that illustrate the covariance for pixels

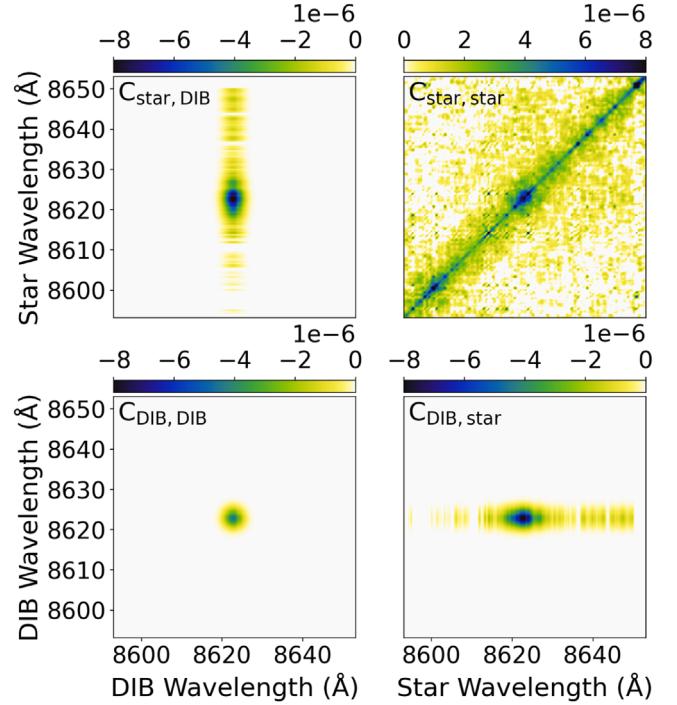


Figure 26. Posteriors for pixels within either the DIB component C_{DIB} or stellar component C_{stellar} . Covariance between pixels spanning the two components is shown as the cross terms $C_{\text{DIB},\text{stellar}}$ and $C_{\text{stellar},\text{DIB}}$.

within and between the stellar and dust components for the decomposition of the spectrum shown in Figure 11.

Because MADGICS imposes that the sum of the components be exactly equal to the data, the covariance on the sum of the components must be identically zero. This can be confirmed directly from Equations (8)–(9).

Appendix I Velocity Shift

To measure the magnitude of the rest-frame shift indicated by the residuals in Figure 13, we performed a least-squares fit of the residuals with the derivative of a Lorentzian line shape and compared that with the amplitude of a Lorentzian fit to the stellar components for the same stars. Each Ca triplet line was fit independently in a 100 pixel wide window (see Section 6 for more details). The result is plotted as a function of stellar S/N in Figure 27 (top) and suggests a velocity shift of $\sim 50 \text{ m s}^{-1}$ at S/N ~ 100 and $\sim 500 \text{ m s}^{-1}$ at S/N 20.

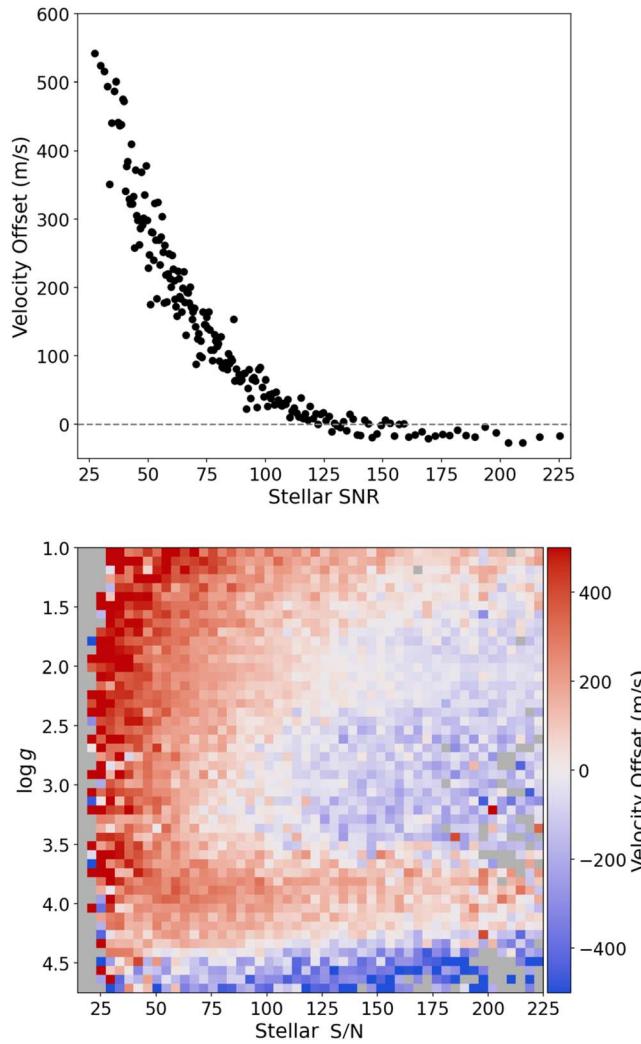


Figure 27. (Top) Velocity shift in the stellar rest frame as measured from the MADGICS residuals by comparing the magnitude a fit of a Lorentzian derivative to a Lorentzian fit of the stellar component of the spectra. (Bottom) Measurements of the average shift in the core of the Ca triplet lines using Lorentzian fits to Gaia data alone. Zero velocity reference arbitrarily chosen to be the $(\text{stellar S/N}, \log g) \approx (150, 2)$ bin. A comparison shows MADGICS residuals are sensitive enough to show small velocity shifts below the reported precision in Gaia radial velocities that are present in the data themselves.

In order to confirm this apparent velocity shift is a result of the data, and not a bias introduced by the MADGICS methodology introduced here, we examined the velocity shift more closely in the data from Gaia without using MADGICS at all. We separated the Gaia data (with the same cuts as described in Section 2) into 50 bins each in $\log(g)$ and S/N. The spectra in each bin were averaged and Lorentzian curves were fit independently to each of the three Ca lines. We then calculated the average velocity shift across all three Ca lines for each bin relative to high S/N giants, $(\text{stellar S/N}, \log g) \approx (150, 2)$, which is shown in Figure 27 (bottom). This analysis shows the same range of velocity offsets (around $\sim 600 \text{ m s}^{-1}$) as fitting the MADGICS residuals, suggesting that this is an effect of the Gaia data propagating through MADGICS.

The asymmetry as a function of S/N could arise from the asymmetric line shape of the Ca triplet, which Allende Prieto et al. (2013) showed leads to apparent shifts of a similar order of magnitude in the radial velocity measure as a function of spectral resolution. Given the Gaia RVS resolution and the type-dependent convective shifts $\sim 300 \text{ m s}^{-1}$ predicted by Allende Prieto et al. (2013), these residuals are impressive and entirely consistent with the reported precision of Gaia radial velocities. A similar bias in the radial velocity measure as a function of stellar magnitude was also shown by Katz et al. (2023), which proposed charge traps in the CCDs were the cause. So, the clarity of these small velocity offsets in the MADGICS residuals is a testament to the sensitivity of the method for determining radial velocity measures.

ORCID iDs

Andrew K. Saydjari <https://orcid.org/0000-0002-6561-9002>
 Ana Sofia M. Uzsoy <https://orcid.org/0000-0001-9308-0449>
 Catherine Zucker <https://orcid.org/0000-0002-2250-730X>
 J. E. G. Peek <https://orcid.org/0000-0003-4797-7030>
 Douglas P. Finkbeiner <https://orcid.org/0000-0003-2808-275X>

References

Allende Prieto, C., Koesterke, L., Ludwig, H. G., Freytag, B., & Caffau, E. 2013, *A&A*, **550**, A103
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Demleitner, M., & Andrae, R. 2021, *AJ*, **161**, 147
 Bailey, M., van Loon, J. T., Farhang, A., et al. 2016, *A&A*, **585**, A12
 Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, *SIAMR*, **59**, 65
 Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
 Bovy, J., Allende Prieto, C., Beers, T. C., et al. 2012, *ApJ*, **759**, 131
 Buder, S., Sharma, S., Kos, J., et al. 2021, *MNRAS*, **506**, 150
 Campbell, E. K., Holz, M., Gerlich, D., & Maier, J. P. 2015, *Natur*, **523**, 322
 Chen, J., & Revels, J. 2016, arXiv:1608.04295
 Contursi, G., de Laverny, P., Recio-Blanco, A., & Palicio, P. A. 2021a, *A&A*, **654**, A130
 Contursi, G., de Laverny, P., Recio-Blanco, A., & Palicio, P. A. 2021b, GSP-spec Line List, Version 21-Sep-2022 (Last Modified) (Strasbourg: CDS)
 Dame, T. M., Hartmann, D., & Thaddeus, P. 2001, *ApJ*, **547**, 792
 Edwards, S. A., & Leach, S. 1993, *A&A*, **272**, 533
 Edwards, S. A., & Leach, S. 1994, in AIP Conf. Ser. 312, Molecules and Grains in Space, ed. I. Nenner (Melville, NY: AIP), 589
 Farhang, A., Khosroshahi, H. G., Javadi, A., et al. 2015, *ApJ*, **800**, 64
 Farhang, A., van Loon, J. T., Khosroshahi, H. G., Javadi, A., & Bailey, M. 2019, *NatAs*, **3**, 922
 Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, **595**, A1
 Galazutdinov, G. A., Musaev, F. A., Krełowski, J., & Walker, G. A. H. 2000, *PASP*, **112**, 648

Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *ApJ*, **887**, 93

Hacar, A., Alves, J., Burkert, A., & Goldsmith, P. 2016, *A&A*, **591**, A104

Hinkle, K., Wallace, L., Valenti, J., & Harmer, D. 2000, Visible and Near Infrared Atlas of the Arcturus Spectrum 3727-9300 Å (San Francisco, CA: ASP)

Hunter, J. D. 2007, *CSE*, **9**, 90

Katz, D., Sartoretti, P., Guerrier, A., et al. 2023, *Gaia* Data Release 3, *A&A*, **674**, A5

Kollmeier, J. A., Zasowski, G., Rix, H.-W., et al. 2017, arXiv:1711.03234

Kos, J., & Zwitter, T. 2013, *ApJ*, **774**, 72

Krelowski, J. 1988, *PASP*, **100**, 896

Kurucz, R. L., Furenlid, I., Brault, J., & Testerman, L. 1984, Solar Flux Atlas from 296 to 1300 nm, National Solar Observatory Atlas, No 1 (Washington, DC: US. National Solar Observatory)

Lai, T. S. Y., Witt, A. N., Alvarez, C., & Cami, J. 2020, *MNRAS*, **492**, 5853

Lan, T.-W., Ménard, B., & Zhu, G. 2015, *MNRAS*, **452**, 3629

Leike, R. H., Glatzle, M., & Ensslin, T. A. 2020, VizieR Online Data Catalog, *J/A+A/639/A138*

Luo, A. L., Zhao, Y.-H., Zhao, G., et al. 2015, *RAA*, **15**, 1095

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, **154**, 94

Munari, U., Tomasella, L., Fiorucci, M., et al. 2008, *A&A*, **488**, 969

Ochsenbein, F., Bauer, P., & Marcout, J. 2000, *A&AS*, **143**, 23

Pelgrims, V. 2022, The shape of the shell of the Local Bubble, V1, Harvard Dataverse, doi:10.7910/DVN/RHPVNC

Pelgrims, V., Ferrière, K., Boulanger, F., Lallement, R., & Montier, L. 2020, *A&A*, **636**, A17

Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A., & Stobie, E. 2010, *A&A*, **524**, A42

Perez, F., & Granger, B. E. 2007, *CSE*, **9**, 21

Portillo, S. K. N., Lee, B. C. G., Daylan, T., & Finkbeiner, D. P. 2017, *AJ*, **154**, 132

Recio-Blanco, A., de Laverny, P., Palicio, P. A., et al. 2023, *A&A*, **674**, A29

Reid, M. J., Menten, K. M., Brunthaler, A., et al. 2019, *ApJ*, **885**, 131

Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525

Schlutheis, M., Zhao, H., Zwitter, T., et al. 2023, *A&A*, **674**, A40

Snow, T. P., Zukowski, D., & Massey, P. 2002, *ApJ*, **578**, 877

Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *AJ*, **132**, 1645

Tchernyshyov, K., & Peek, J. E. G. 2017, *AJ*, **153**, 8

Tchernyshyov, K., Peek, J. E. G., & Zasowski, G. 2018, *AJ*, **156**, 248

The HDF Group 1997–2022, Hierarchical Data Format, version 5, <https://www.hdfgroup.org/solutions/hdf5/>

Tomasi, M., & Li, Z. 2021, Healpix.jl: Julia-only port of the HEALPix library, 3.0, Astrophysics Source Code Library, ascl:2109.028

Vallenari, A., Brown, A. G. A., Prusti, T., et al. 2023, *A&A*, **674**, A1

van Loon, J. T., Bailey, M., Tatton, B. L., et al. 2013, *A&A*, **550**, A108

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Walker, G. A. H., Bohlender, D. A., Maier, J. P., & Campbell, E. K. 2015, *ApJL*, **812**, L8

Welty, D. E., & Hobbs, L. M. 2001, *ApJS*, **133**, 345

Woodbury, M. A. 1950, Memorandum Rep.42, Statistical Research Group, Princeton Univ.

York, D. G., Adelman, J., Anderson, J. E. J., et al. 2000, *AJ*, **120**, 1579

Zanolin, M., Vitale, S., & Makris, N. 2010, *PhRvD*, **81**, 124048

Zasowski, G., Ménard, B., Bizyaev, D., et al. 2015, *ApJ*, **798**, 35

Zhao, H., Schlutheis, M., Rojas-Arriagada, A., et al. 2021, *A&A*, **654**, A116

Zucker, C., Goodman, A. A., Alves, J., et al. 2022, *Natur*, **601**, 334