



A Deep, High-angular-resolution 3D Dust Map of the Southern Galactic Plane

Catherine Zucker^{1,2}, Andrew K. Saydjari^{1,3,4,14}, Joshua S. Speagle (沈佳士)^{5,6,7,8,15}, Edward F. Schlafly², Gregory M. Green⁹, Robert Benjamin¹⁰, Joshua Peek², Gordian Edenhofer^{11,12}, Alyssa Goodman¹, Michael A. Kuhn¹³, and Douglas P. Finkbeiner^{1,3}

¹ Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA; catherine.zucker@cfa.harvard.edu

² Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

³ Department of Physics, Harvard University, 17 Oxford St., Cambridge, MA 02138, USA

⁴ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁵ Department of Statistical Sciences, University of Toronto, 9th Floor, Ontario Power Building, 700 University Ave, Toronto, ON M5G 1Z5, Canada

⁶ David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, 50 St George Street, Toronto, ON M5S 3H4, Canada

⁷ Dunlap Institute for Astronomy & Astrophysics, University of Toronto, 50 St George Street, Toronto, ON M5S 3H4, Canada

⁸ Data Sciences Institute, University of Toronto, 17th Floor, Ontario Power Building, 700 University Ave, Toronto, ON M5G 1Z5, Canada

⁹ Max Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany

¹⁰ Department of Physics, University of Wisconsin-Whitewater, 800 W Main St, Whitewater, WI 53190, USA

¹¹ Max Planck Institute for Astrophysics, Karl-Schwarzschildstraße 1, 85748 Garching, Germany

¹² Ludwig-Maximilians-Universität, Geschwister-Scholl Platz 1, 80539 Munich, Germany

¹³ Centre for Astrophysics Research, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK

Received 2025 March 3; revised 2025 August 12; accepted 2025 August 13; published 2025 October 3

Abstract

We present a deep, high-angular-resolution 3D dust map of the southern Galactic plane over $239^\circ < l < 6^\circ$ and $|b| < 10^\circ$ built on photometry from the DECaPS2 survey, in combination with photometry from VISTA Variables in the Via Lactea, the Two Micron All Sky Survey, and “Unofficial” Wide-field Infrared Survey Explorer and parallaxes from Gaia Data Release 3 where available. To construct the map, we first infer the distance, extinction, and stellar types of over 700 million stars using the BRUTUS stellar inference framework with a set of theoretical MESA Isochrone and Stellar Tracks (MIST) stellar models. Our resultant 3D dust map has an angular resolution of $1'$, roughly an order of magnitude finer than existing 3D dust maps and comparable to the angular resolution of the Herschel 2D dust emission maps. We detect complexes at the range of distances associated with the Sagittarius-Carina and Scutum-Centaurus arms in the fourth quadrant, as well as more distant structures out to a maximum reliable distance of $d \approx 10$ kpc from the Sun. The map is sensitive up to a maximum extinction of roughly $A_V \approx 12$ mag. We publicly release both the stellar catalog and the 3D dust map, the latter of which can easily be queried via the Python package `dustmaps`. When combined with the existing Bayestar19 3D dust map of the northern sky, the DECaPS 3D dust map fills in the missing piece of the Galactic plane, enabling extinction corrections over the entire disk $|b| < 10^\circ$. Our map serves as a pathfinder for the future of 3D dust mapping in the era of LSST and Roman, targeting regimes accessible with deep optical and near-infrared photometry but often inaccessible with Gaia.

Unified Astronomy Thesaurus concepts: [Interstellar dust \(836\)](#); [Milky Way Galaxy \(1054\)](#); [Surveys \(1671\)](#)

Materials only available in the [online version of record](#): interactive figures

1. Introduction

The distribution of interstellar dust in the Milky Way has profound implications not only as a foreground contaminant for a broad range of astronomical observations, but also as a tracer of Galactic spiral structure and the sites of star formation within dense molecular clouds.

Interstellar dust scatters and absorbs starlight at near-infrared, optical, and ultraviolet wavelengths, causing a reddening effect due to the preferential attenuation of higher-energy photons. Near bright stars, interstellar dust clouds also reflect starlight, generating so-called “reflection nebulae” (K. Sellgren 1984). Due to the nonspherical nature of grains

whose major axes align perpendicular to the magnetic field, interstellar dust polarizes starlight (J. L. Han 2017). Interstellar dust also reradiates at mid- and far-infrared wavelengths, acting as a critical foreground for the cosmic microwave background (D. P. Finkbeiner et al. 1999; Planck Collaboration et al. 2016).

Through its absorption, scattering, and reprocessing of starlight, interstellar dust also plays a pivotal role in the physics and chemistry of the interstellar medium (ISM). For example, dust helps regulate the temperature of the ISM, heating the gas via photoelectric heating (J. C. Weingartner & B. T. Draine 2001). Interstellar dust also catalyzes the formation of molecular hydrogen—the key ingredient in molecular clouds—and allows star formation to occur by removing the gravitational energy of collapsing clouds via far-infrared radiation (B. T. Draine 2003).

Historically, most of our understanding of the distribution of Galactic dust has come from emission, either from associated H I or from the dust itself. These maps are inherently 2D with varying angular resolution, where information on the structure

¹⁴ Equal contribution Hubble Fellow.

¹⁵ Equal contribution.



of the dust as a function of distance has been projected onto the plane of the sky. A number of 2D dust maps have been built over the past half century. For example, D. Burstein & C. Heiles (1978) combined galaxy counts with H I column density measurements derived from 21 cm spectral-line maps to probe integrated dust column density under the assumption that gas and dust are well mixed (see also D. Burstein & C. Heiles 1982; D. Lenz et al. 2017). D. J. Schlegel et al. (1998) derived the dust column density from far-infrared emission at 100 and 240 μm using IRAS (G. Neugebauer et al. 1984) and DIRBE data (N. W. Boggess et al. 1992), calibrating the far-infrared flux to dust reddening using a sample of elliptical galaxies, and later Sloan Digital Sky Survey (SDSS) photometry and spectra of 250,000 stars (E. F. Schlafly & D. P. Finkbeiner 2011), producing a map with an angular resolution of $6'$.

More recently, Planck Collaboration et al. (2014) used a similar technique to D. J. Schlegel et al. (1998) for modeling far-infrared emission, combining the 100 μm IRAS data with Planck data between 353 and 857 GHz to derive an all-sky 2D map of dust reddening at an angular resolution of $5'$. Concurrently, A. M. Meisner & D. P. Finkbeiner (2015) produced all-sky maps of diffuse Galactic dust emission based on Wide-field Infrared Survey Explorer (WISE) 12 μm images at an angular resolution of $15''$. The WISE-based maps are roughly a factor of $4\times$ higher than the next highest angular resolution 2D dust emission maps from Herschel. Compared to WISE, the Herschel maps achieve a lower $1'$ angular resolution ($37''$ at 500 μm ; M. J. Griffin et al. 2010), but suffer from less contamination from stars and galaxies.

Alongside 2D emission-based maps, the recent rise in wide-field surveys has also enabled the construction of 2D dust maps based on photometry. These 2D dust maps primarily rely on measuring the infrared color excess toward red clump (RC) and red giant branch (RGB) stars, because these stellar populations act as so-called “standard crayons,” with very consistent intrinsic colors. By estimating the color excess to many standard crayons, these 2D star-based approaches can reconstruct the integrated extinction out to large distances, commensurate with the completeness limit of the underlying stellar populations. A number of such pioneering 2D star-based maps have been produced, particularly toward the inner Galaxy. For instance, F. Surot et al. (2020) use RC and RGB stars detected in the VISTA Variables in the Via Lactea (VVV) survey to map the color excess toward a 300 sq. deg² region of the bulge with very fine angular resolution ranging between $10''$ and $2'$ (see also O. A. Gonzalez et al. 2012; D. M. Nataf et al. 2013; J. Alonso-García et al. 2017). More recently, using intrinsic colors drawn from the StarHorse stellar parameter catalog (F. Anders et al. 2019) rather than purely RC and RGB stars, M. Zhang & J. Kainulainen (2022) built a 2D dust map over the full VVV survey area, achieving an angular resolution of $30''$ and sensitivity to extinctions up to $A_V \approx 10\text{--}20$ mag.

While these 2D emission and 2D star-based maps have been critical for correcting for the effects of dust obscuration primarily on extragalactic observations—where the entire Milky Way is a foreground—they are largely insufficient for corrections *within* the Milky Way, or to probe the internal structure of the Milky Way’s ISM. For these applications, it is necessary to map the distribution of dust in three spatial dimensions (3D).

3D dust mapping has flourished over roughly the past decade, adding the critical third dimension (distance) to our understanding of the interstellar dust distribution. Similar to the 2D star-based dust maps, 3D dust mapping relies on the principle that dust reddens stellar photometric colors, so these 3D maps are based on dust extinction rather than dust emission. By modeling this cumulative reddening effect for stars at different distances along an individual line of sight in the Milky Way, it is possible to infer the differential reddening along that line of sight. By grouping hundreds of millions of stars sightline by sightline across the Galaxy, one can reconstruct the 3D distribution of dust.

One of the primary challenges of early 3D dust mapping efforts is simultaneously inferring the type of the star (and thus its intrinsic colors), its distance, and reddening from photometry alone. Despite the challenges, several such 3D dust maps have been built solely on the optical and/or infrared photometry readily available in large-scale surveys. D. Marshall et al. (2025) combined the near-infrared photometric colors of stars detected by the Two Micron All Sky Survey (2MASS; M. F. Skrutskie et al. 2006) with models for stars’ intrinsic colors and distances obtained from the Besançon Stellar Population Synthesis Model of the Galaxy (A. C. Robin et al. 2003) to infer the 3D distribution of reddening toward the inner Galactic plane. S. E. Sale et al. (2014) used a hierarchical Bayesian model applied to photometry from the IPHAS survey (J. E. Drew et al. 2005) to construct a 3D dust map toward the northern Galactic plane. G. M. Green et al. (2015) probabilistically inferred the distances, reddenings, and stellar types of ≈ 800 million stars using a combination of Pan-STARRS1 (PS1; K. C. Chambers et al. 2016) and 2MASS photometry to produce a 3D dust map over three-quarters of the northern sky (the “Bayestar” map; see also G. M. Green et al. 2014, 2018).

By providing constraints on the distances of a billion stars— independent of their colors—the second data release (DR2) of the Gaia mission ushered in a new era of 3D dust maps with unprecedented distance resolution. A number of efforts have combined the broadband photometric colors of stars with Gaia astrometry in pursuit of even more highly resolved maps with improved distance resolution. For example, R. Lallement et al. (2019) applied a hierarchical inversion algorithm to a combination of Gaia and 2MASS photometry and Gaia parallax measurements to produce a 3D dust map out to 3 kpc from the Sun with a distance resolution (in Cartesian space) of ≈ 25 pc (see also R. Lallement et al. 2018, 2022; J. L. Vergely et al. 2022). Applying a random forest regression to a combination of Gaia, 2MASS, and WISE photometry alongside Gaia parallax measurements, B. Q. Chen et al. (2019) inferred the properties of ≈ 50 million stars and constructed a 3D dust map of the full Galactic plane out to ≈ 5 kpc with an angular resolution of $6'$. G. M. Green et al. (2019) incorporated Gaia parallaxes into their “Bayestar” dust mapping pipeline and implemented a new spatial prior (compared to G. M. Green et al. 2018) for correlating neighboring sightlines, resulting in a 3D dust map with 4 times finer distance resolution than their previous maps and a typical angular resolution of $\approx 7'$. Incorporating Gaia-informed estimates of stellar distances and reddenings (obtained from the StarHorse catalog; F. Anders et al. 2019), R. H. Leike et al. (2020) combined metric Gaussian variational inference (J. Knollmüller & T. A. Enßlin 2019), Gaussian processes, and information field theory to map the 3D

distribution of dust at ≈ 1 pc Cartesian resolution out to ≈ 400 pc from the Sun—resolving the detailed internal structure of local molecular clouds for the first time (see also R. H. Leike & T. A. Enßlin 2019). Using a new catalog of distance and extinction measurements from X. Zhang et al. (2023) based on Gaia BP/RP spectra, G. Edenhofer et al. (2024) built upon the methodology of R. H. Leike et al. (2020) to produce a parsec-scale 3D dust map out to a distance of 1.25 kpc from the Sun at $14'$ angular resolution.

While Gaia has been transformational to the field of 3D dust mapping, the gains have largely been limited to the solar neighborhood, within a few kiloparsecs of the Sun. This limitation stems from the underlying requirement that stars must be detected *behind* dense dust clouds in order for the cloud to be detected in 3D maps. If the dust column is so high—as in the inner Galactic plane—that it extinguishes the light from background stars entirely, the cloud goes undetected in 3D maps (G. M. Green et al. 2019; C. Zucker et al. 2019). Because most existing maps rely on either easily extinguished optical photometry and astrometry from Gaia and/or shallow near-infrared photometry (from, e.g., 2MASS), most structure in the inner Galactic plane at distances $\gtrsim 2$ kpc has largely gone unresolved in 3D. This structure includes some of the richest regions in the inner Galaxy, including the Scutum-Centaurus arm, which hosts much of the Milky Way’s most active star formation.

In this work, we leverage two of the deepest infrared surveys toward the inner Galaxy, the Dark Energy Camera Plane Survey 2 (DECaPS2; A. K. Saydjari et al. 2023b) and the VISTA Variables in the Via Lactea (VVV; D. Minniti et al. 2010), to infer the distances, reddenings, and stellar types of hundreds of millions of stars. Incorporating Gaia parallax distances when available nearby, but relying on the deep infrared photometry in heavily dust-enshrouded regions at greater distances, we construct a 3D dust map of the southern Galactic plane. Our 3D dust map achieves three main goals:

1. By incorporating deep infrared photometry into an improved stellar modeling pipeline, we resolve hitherto-undetected structure toward the inner Galactic plane at distances $\gtrsim 2$ kpc.
2. By leveraging the sheer stellar density of DECaPS2 (≈ 800 million stars over only 6% of the sky), we produce the highest-angular-resolution 3D dust map ($\approx 1'$) to date, constituting almost an order of magnitude improvement in angular resolution over current 3D dust maps (B. Q. Chen et al. 2019; G. M. Green et al. 2019).
3. By targeting the portion the Galaxy unreachable by PS1 ($\delta < -30^\circ$), we can combine our map with the “Bayestar19” 3D dust map (G. M. Green et al. 2019) to provide full coverage of the Galactic plane at $|b| < 10^\circ$ and enable extinction corrections over the entire disk.

In Figure 1, we provide a schematic overview of our pipeline. In Section 2, we describe the various data sets and the survey-level cuts applied to ensure data quality. In Section 3, we describe how the photometric and astrometric data from these surveys are combined into a single catalog. Leveraging the assembled catalog, we describe the methodology underpinning our stellar modeling and 3D dust mapping pipelines in

Section 4. In Section 5, we present the 3D distribution of stars and dust in the southern Galactic plane. We discuss the implications of our map in Section 6 and compare with existing efforts. In Section 7, we discuss the availability of the data generated and software produced under this work. Finally, we conclude in Section 8. For a summary of known artifacts in our 3D dust map, we refer readers to Appendix A. In this work, our goal is to present the method of construction of the map, while future works will use this map to further investigate the spatial and dynamical organization of the ISM in the inner Galaxy.

2. Data

Our analysis is based on a combination of photometric and astrometric data from five surveys:

1. The Dark Energy Camera Plane Survey 2 (DECaPS2; A. K. Saydjari et al. 2023b).
2. The Vista Variables in the Via Lactea Survey (VVV; D. Minniti et al. 2010).
3. The Two Micron All Sky Survey (2MASS; M. F. Skrutskie et al. 2006).
4. The “Unofficial” Wide-field Infrared Survey Catalog (unWISE; E. F. Schlafly et al. 2019).
5. The third data release of the Gaia Mission (Gaia DR3; Gaia Collaboration et al. 2023).

2.1. DECaPS2

The second release of the Dark Energy Camera Plane Survey (DECaPS2; A. K. Saydjari et al. 2023b) is a deep, five-band optical and near-infrared survey of the southern Galactic plane ($239^\circ < l < 6^\circ$, $|b| < 10^\circ$).¹⁶ It builds upon the first data release (DECaPS1; E. F. Schlafly et al. 2018) by improving the photometric reduction and extending the latitude range from $|b| < 4^\circ$ to $|b| < 10^\circ$, for a total of 2700 deg^2 of sky coverage (6.5%) with an average of three epochs per band. The DECaPS2 survey totals 3.32 billion objects and achieves average single-exposure AB magnitude depths of 23.5, 22.6, 22.1, 21.6, and 20.8 mag in the g , r , i , z , and Y bands, respectively, though with considerable variation due to crowding. The filters span the wavelength range 400–1065 nm. We use the CloudCoverr-“corrected” crowdsource photometry where the full background covariance is modeled, improving the uncertainty estimates in regions with structured backgrounds (see E. F. Schlafly 2021; A. K. Saydjari & D. P. Finkbeiner 2022, for details on the methods). On a band-by-band basis, we require an object to have an “OK” detection in at least one epoch ($\text{nmag_cflux_ok} > 0$), errors < 0.1 mag, and $\text{fracflux} > 0.75$, indicating that at least 75% of the point-spread function (PSF)-weighted fraction of flux at the star’s location is derived from itself (as opposed to neighboring sources). The DECaPS2 survey complements the PS1 survey, which covers three-quarters of the northern sky ($\delta > -30^\circ$) at wavelengths similar to DECaPS2, which was previously used in the construction of the Bayestar19 3D dust map (G. M. Green et al. 2019).

¹⁶ All the DECaPS2 photometry utilized here is available for download at http://decaps.skymaps.info/release/data/files/DR2_REDUX/DATABASE/dbfits/ and accessible via NOIRLab’s AstroDataLab.




Figure 1. Schematic overview of the process of turning images of the sky into a 3D map of dust. The DECaPS2 survey forms the foundation of our 3D dust map, whose photometric colors are combined with complementary photometric (VVV, 2MASS, unWISE) and astrometric surveys (Gaia) where available and fed into our stellar inference framework. We use the brutus stellar inference framework to infer the distance, extinction, and stellar type of hundreds of millions of stars. We then group stars into pixels and fit the set of distance–reddening measurements along the line of sight in each pixel to generate a 3D map of dust.

2.2. VVV

The Vista Variables in the Via Lactea Survey (VVV) survey (D. Minniti et al. 2010) is a near-infrared survey targeting 562 deg^2 of the bulge ($-10^\circ < l < 10^\circ$, $-10^\circ < b < 5^\circ$) and southern Galactic plane ($-65^\circ < l < -10^\circ$, $-2^\circ < b < 2^\circ$). The original PSF reduction of the VVV survey by J. Alonso-García et al. (2018) used DoPhot (P. L. Schechter et al. 1993) to derive a catalog of 846 million sources with coverage in the Z , Y , J , H , and K_s filters, spanning 0.84–2.5 μm . We utilize a more recent PSF reduction by M. Zhang & J. Kainulainen (2019), which applies Daophot (P. B. Stetson 1987) to the J , H , and K_s bands only to obtain a 926 million source catalog that goes roughly 1 magnitude deeper than J. Alonso-García et al. (2018). The M. Zhang & J. Kainulainen (2019) reduction achieves 5σ limiting magnitudes of 20.8, 19.5, and 18.7 mag in the J , H , and K_s bands, respectively, in the Vega system. M. Zhang & J. Kainulainen (2019) flag sources as spurious detections by sigma-clipping outlying sources using the magnitude–error relations in different band combinations.

The DECaPS2 error modeling implemented by A. K. Saydjari et al. (2023b) debiased and improved uncertainty estimates for photometry on structured backgrounds using the CloudCovErr algorithm. Injection tests performed on every DECam exposure in the survey indicate that the photometric error estimates are correct except in regions of the most extreme crowding. Without marginalizing over structured backgrounds, uncertainty can often be underestimated by a factor of 2 or more. Therefore, we adopt a more stringent error

cut for VVV, 2MASS, and unWISE since their uncertainties may be underestimated. On a band-by-band basis, we require that no source is flagged as spurious in any magnitude–error relation using that band, and that the error is <0.05 mag.

2.3. 2MASS

The Two Micron All Sky Survey (2MASS; M. F. Skrutskie et al. 2006) is an all-sky infrared survey in the J , H , and K bands, achieving a 10σ point-source depth of 15.8, 15.1, and 14.3 mag, respectively, in the Vega system. The point-source catalog contains a total of 470 million objects, of which roughly 340 million sources are considered good quality. On a band-by-band basis, we only utilize detection that meet the 2MASS “high-reliability” criteria¹⁷ with errors <0.05 mag. We also exclude sources that are flagged as having possible contamination from nearby bright point sources (requiring `cc_flg==0`) and galaxies (requiring `gal_contam==0`).

2.4. unWISE

The “Unofficial” Wide-field Infrared Survey Explorer (unWISE; E. F. Schlafly et al. 2019) catalog analyzes the unblurred “unWISE” coadds (D. Lang 2014) derived from the WISE images to detect 2 billion sources over the full sky in the W1 and W2 bands at 3.4 μm and 4.5 μm , respectively.

¹⁷ For a description of the high-reliability criteria, see the 2MASS All Sky Data Release Explanatory Supplement (https://irsa.ipac.caltech.edu/data/2MASS/docs/releases/allsky/doc/sec1_6b.html).

Compared to the official ALLWISE catalog (R. M. Cutri et al. 2013), the unWISE catalog is based on deeper imaging and uses the `crowdsources` algorithm (E. F. Schlafly 2021) to improve modeling of crowded regions. The unWISE catalog extends 0.7 mag deeper than ALLWISE, achieving a 5σ point-source depth of ≈ 20.7 and 20.0 mag for W1 and W2 in the AB system, respectively. On a band-by-band basis, we require that no flags are set (`flags_unwise==0`), errors < 0.05 mag, and `fracflux` > 0.85 , indicating that at least 85% of the PSF-weighted fraction of flux at the star’s location is derived from itself rather than contamination from adjacent sources.

2.5. Gaia DR3

Gaia (Gaia Collaboration et al. 2016) is an all-sky optical survey providing astrometry (parallaxes and proper motions) and photometry (in the *G*, BP, and RP bands) for over 1 billion stars. Given the breadth of the Gaia passbands (330–1050 nm in *G*, 330–680 nm in BP, and 630–1050 nm in RP) in comparison to the DECaPS2 filters covering a similar wavelength range, we only utilize the Gaia astrometry (specifically the Gaia parallax measurements) in this work. We leverage the parallax measurements from the third data release (Gaia DR3; Gaia Collaboration et al. 2023), which provides a median parallax uncertainty of 0.02–0.03 mas for $G = 9$ –14 mag and 0.5 mas uncertainty for $G = 20$ mag (L. Lindegren et al. 2021). Following C. Fabricius et al. (2021), we exclude stars with renormalized unit weight error `ruwe` > 1.4 .

3. Assembling a Final Catalog

To assemble a final catalog, we cross-match the DECaPS2, VVV, 2MASS, unWISE, and Gaia DR3 surveys using the Large Survey Database (LSD) architecture (M. Juric 2012). DECaPS2 serves as the primary catalog for the astrometry, meaning we find photometry and astrometry for sources in the other catalogs that cross-match to detections in DECaPS2, adopting a cross-match radius of $0''.5$. After applying the band-by-band survey-level photometric quality cuts described in Sections 2.1–2.4, we require that the star be detected in at least four bands (*g*, *r*, *i*, *z*, *Y*, *J*, *H*, *K*, W1, W2), at least one of which must be a DECaPS2 band. To ensure continuity across the VVV survey boundary, we do consider sources that have detections in both VVV and 2MASS, but do not double count the *J*, *H*, and *K* detections for the purposes of imposing the four-band minimum. For example, if a star is detected in 2MASS *J*, *H*, *K* and VVV *J*, *H*, *K* alongside a DECaPS *Y*-band detection, this star would be labeled as a four-band detection when filtering, but both the 2MASS and VVV data would be incorporated in the fit. We do not require a Gaia parallax measurement but include these distance constraints in the catalog when available. In Figure 2, we break down the number of photometric bands per star incorporated into the stellar modeling, tabulated across the full DECaPS2 footprint. The number of bands ranges from four (the minimum) to 13 (the full filter coverage, consisting of *g*, *r*, *i*, *z*, *Y*, J_{VVV} , H_{VVV} , K_{VVV} , J_{2MASS} , H_{2MASS} , K_{2MASS} , W1, and W2), with an average of five bands per star.

Our resulting catalog contains 793 million sources. Despite targeting only one-third of the Galactic plane, our source count is just shy of the 799 million stars used to reconstruct the Bayestar19 3D dust map (covering the remaining two-




Figure 2. Number of photometric bands per star incorporated into the stellar modeling, tabulated across all stars in the full DECaPS2 footprint. The number of photometric detections ranges from four (the minimum) to 13 (the maximum), with an average of five bands incorporated per star.

thirds of the plane and the rest of the northern sky), underlining the increased source density in the DECaPS2 footprint (Figure 3). In Figure 4, we break down the number of stars detected per band. There are between 605 and 772 million detections in any individual DECaPS2 *grizY* band. Roughly 20% of the sample have detections in a VVV band (138–186 million detections per band), compared to roughly 3% for 2MASS (25–35 million detections per band). Less than 2% of stars are detected in unWISE (15 million detections per band). In contrast, roughly half of the sample (427 million stars) have a Gaia parallax measurement, though only 5% of stars (40 million) have a parallax that alone will strongly constrain the star’s distance (signal-to-noise ratio $\text{SNR} > 5$). In many cases the low-SNR parallaxes are still useful for breaking the dwarf–giant degeneracy. Considering only the region of the sky where there is overlap between the VVV and DECaPS2 footprints (purple polygon in Figure 3, encompassing part of the bulge and the plane and the southern plane $|b| < 2^\circ$), roughly 75% of stars have both DECaPS2 and VVV detections, underlining the importance of the deeper near-infrared photometry of VVV compared to 2MASS in dust-enshrouded regions.

4. Methods

Here we outline our methodology for mapping the 3D distribution of dust in the southern Galactic plane. As described in Section 4.1, we start by inferring the stellar parameters, dust extinction, and distance on a star-by-star basis using our assembled photometric and astrometric catalog from Section 3. We then group the stars into discrete pixels on the celestial sphere and use the per-star distance and extinction estimates within each pixel to model the distribution of dust as a function of distance along the line of sight, as described in Section 4.2. The per-star and line-of-sight inference required a significant number of CPU hours, and the computational resources required to generate these data products are described in Appendix B.

4.1. Stellar Modeling

To perform the stellar modeling, we leverage the open-source Python package BRUTUS (J. S. Speagle et al. 2025).




Figure 3. Distribution of source density on the plane of the sky for stars we model in Section 4.1.1, the majority of which will be used in the line-of-sight dust reconstruction. The VVV footprint (purple polygon) shows a modest increase in source density in comparison to the rest of the DECaPS2 footprint. The median source density is 70 stars per arcmin², with an interquartile range of 37–125 stars per arcmin².

BRUTUS derives estimates of extrinsic stellar parameters ϕ (distance, extinction, total-to-selective extinction ratio) over a grid of intrinsic stellar parameters θ (e.g., initial mass). The BRUTUS pipeline has already successfully been applied to estimate distances, extinctions, and stellar properties of 170 million stars at high galactic latitude (the *Augustus* catalog) as presented in J. S. Speagle et al. (2024). In Section 4.1.1, we briefly summarize the core inference framework from J. S. Speagle et al. (2025) and highlight specific adaptations implemented in this work with respect to J. S. Speagle et al. (2024). In Section 4.1.2, we describe the BRUTUS setup (version v0.8.3) used to perform the inference on the assembled catalog from Section 3.

4.1.1. Statistical Framework

We assume that the observed magnitudes of a star, denoted by $\mathbf{m} \equiv \{m_i\}_{i=1\dots b}$, measured across a set of b photometric passbands, can be modeled as

$$\mathbf{m}_{\theta,\phi} \equiv \mathbf{M}_{\theta} + \mu + A_V \times (\mathbf{R} + R_V \times \mathbf{R}'), \quad (1)$$

where $\mathbf{M}_{\theta}$ is the set of absolute magnitudes for a star as a function of its intrinsic stellar parameters θ ; μ is the distance modulus; A_V is the dust extinction, measured as the total attenuation in magnitudes in the V band; $\mathbf{R}$ is the reddening vector; R_V is the total-to-selective extinction ratio, measured as the attenuation in the V band A_V relative to the color excess in the B and V bands, E_{B-V} ($R_V = \frac{A_V}{E_{B-V}}$); and $\mathbf{R}'$ is the differential reddening vector that sets the shape of the reddening vector $\mathbf{R}'$.

Following J. S. Speagle et al. (2024), to parameterize θ we utilize the MESA Isochrone and Stellar Tracks (MIST) models (J. Choi et al. 2016), which are a set of theoretical isochrones that relate intrinsic stellar evolutionary parameters θ to surface-level parameters θ_* via stellar atmospheric models (see P. A. Cargile et al. 2020). The intrinsic stellar parameters θ include the initial mass M_{init} , the initial metallicity $[\text{Fe}/\text{H}]_{\text{init}}$, and the equivalent evolutionary phase (EEP), which breaks down the various stages of stellar evolution into an equal number of steps to ensure efficient sampling of phases even when stars are rapidly evolving on the post main sequence. The surface-level parameters θ_* include the effective temperature T_{eff} , the surface gravity $\log g$, the surface metallicity $[\text{Fe}/\text{H}]_{\text{surf}}$, and the surface α -abundance

enhancement $[\alpha/\text{Fe}]_{\text{surf}}$. As in J. S. Speagle et al. (2024), $[\alpha/\text{Fe}]_{\text{surf}} = 0$ by default since the current models assume solar-scaled abundance patterns. However, we hope to explore the effect of α -abundance enhancements—and particularly any potential degeneracies between $[\alpha/\text{Fe}]_{\text{surf}}$ and extinction—in future work.

We utilize version 1.2 of the nonrotating MIST models. We do not model the secondary mass fraction and assume all stars are single stars. The exact grid of models for the stellar evolutionary parameters θ utilized in this work (`grid_mist_v10`) is summarized in Table 1 and available for download via the BRUTUS GitHub page. Notably, we extend the mass grid down to $0.2 M_{\odot}$ from the minimum mass of $0.5 M_{\odot}$ considered in J. S. Speagle et al. (2024) to account for the abundance of M dwarfs now being detected at greater distances in the DECaPS2 survey. Note that we do not include any models on the pre main sequence ($\text{EEP} < 202$) or beyond the start of the thermally pulsing asymptotic giant branch (AGB; $\text{EEP} > 808$).

The quantities $\mathbf{R}$ and $\mathbf{R}'$ are fixed to the extinction curve of E. F. Schlafly et al. (2016) and published filter curves for the different bandpasses.¹⁸ Unlike J. S. Speagle et al. (2024), we assume that $\mathbf{R}$ and $\mathbf{R}'$ are the same for all models, such that $\mathbf{R}_{\theta} = \mathbf{R}$ and $\mathbf{R}'_{\theta} = \mathbf{R}'$. Therefore, the extinction curve does not change as a function of spectral type. The gray component of the extinction curve is not measured in E. F. Schlafly et al. (2016), and instead relies on the study of R. Indebetouw et al. (2005), which finds $A_H/A_K = 1.55$. The components of $\mathbf{R}$ are computed as derivatives of A_{542} at $A_{542} = 2.3$ mag for a 4500 K solar-metallicity giant, where A_{542} is the extinction at 542 nm. The values of $\mathbf{R}'$ give the derivative of the extinction at $R_V = 3.32$ for each filter for the same dust column and solar spectrum. We renormalize the extinction curve to scale with A_V , and the extinction coefficients normalized to $A_V = 1$ mag for the mean $R_V = 3.32$ are summarized in Table 2. Note that the extinction coefficient for the WISE W2 band, A_{W2} , is larger than typically reported in the literature, on par with A_J (see, e.g., S. Wang & X. Chen 2019). The W2 band is the only infrared band (and the only band other than the g band) to have a negative derivative R' , causing this value to be slightly larger than literature values when we renormalized the extinction curve to scale with A_V (see Equation (1)). Less than 2% of stars

¹⁸ We document the extinction curve script originally published in E. F. Schlafly et al. (2016) on Zenodo at doi:10.5281/zenodo.16813633




Figure 4. Breakdown of the number of stars detected in each photometric band (shown in blue for DECaPS2, purple for VVV, pink for 2MASS, orange for unWISE) and with an available Gaia parallax measurement (shown in yellow). For Gaia, we further subdivide the stars into all those with a Gaia parallax detection and only those with a signal-to-noise ratio on the parallax detection >5 . The top panel shows the breakdown of band coverage for the entire DECaPS2 footprint, while the bottom panel shows the breakdown for the subset of the DECaPS2 footprint that overlaps with the VVV footprint (see purple polygon in Figure 3).

Table 1

Parameter Grid for the MIST Models (grid_mist_v10)

Minimum	Maximum	Spacing
Initial Mass (M_{init})		
$0.2 M_{\odot}$	$2.0 M_{\odot}$	$0.025 M_{\odot}$
Initial Metallicity ($[\text{Fe}/\text{H}]_{\text{init}}$)		
-3.0	-2.0	0.10
-2.0	0.45	0.05
Equivalent Evolutionary Point (EEP)		
202	454	6
454	808	2

have detections in W2 (and the extinction behaves as expected when scaling with A_V in Equation (1)) so this effect should be negligible on the resulting stellar inference. This model for the extinction curve is an approximation, assuming that a particular parcel of dust will change the magnitude of any observed star in exactly the same way. This model would be correct if the bandpasses were infinitely narrow or the extinction curve constant in wavelength over a bandpass. In fact, variation of the extinction curve over the bandpass introduces small dependences of $\mathbf{R}$ and $\mathbf{R}'$ in Table 2 on extinction and stellar type. Here we neglect that variation, which is quite small for the bandpasses considered in this work but is much larger for the very broad Gaia G , BP, and RP bands, for example. This model has the advantage that it is simpler and can be implemented more efficiently. While we technically allow the extinction curve to vary from star to star (by sampling in R_V), any inference surrounding the dust

Table 2

Adopted Reddening Vector

Filter	$\mathbf{R}$	$\mathbf{R}'$
DECam g	1.3285	-0.0169
DECam r	0.7000	0.0631
DECam i	0.3529	0.0982
DECam z	0.1814	0.0989
DECam Y	0.1297	0.0933
VISTA J	0.0823	0.0609
VISTA H	0.0492	0.0371
VISTA K	0.0318	0.0241
2MASS J	0.0835	0.0624
2MASS H	0.0489	0.0370
2MASS K	0.0315	0.0239
WISE W1	0.0484	0.0052
WISE W2	0.0746	-0.0080

Note. Parameterization of the mean reddening vector $\mathbf{R}$ (corresponding to $R_V = 3.32$) and the differential reddening vector $\mathbf{R}'$ in this work (see Equation (1)). The extinction coefficients have been normalized to $A_V = 1$ mag with the gray component fixed assuming $\frac{A_H}{A_K} = 1.55$ following R. Indebetouw et al. (2005).

extinction curve is highly prior driven, as we discuss further in Appendix C.

Based on our model, the posterior probability $P(\theta, \phi | \hat{m}, \hat{\omega})$ that the set of intrinsic stellar parameters θ and extrinsic parameters ϕ is consistent with the observed magnitudes $\hat{m}$, measured Gaia parallax $\hat{\omega}$, and prior information on θ and ϕ is

given by Bayes' theorem:

$$P(\theta, \phi | \hat{m}, \hat{\omega}) \propto P(\hat{m}, \hat{\omega} | \theta, \phi) P(\theta, \phi) \\ \equiv \mathcal{L}_{\text{phot}}(\theta, \phi) \mathcal{L}_{\text{astr}}(\phi) \pi(\theta, \phi), \quad (2)$$

where $P(\theta, \phi | \hat{m}, \hat{\omega})$ is the posterior probability for θ and ϕ , $P(\hat{m}, \hat{\omega} | \theta, \phi) \equiv \mathcal{L}_{\text{phot}}(\theta, \phi) \mathcal{L}_{\text{astr}}(\phi)$ is the likelihood—broken into the photometric $\mathcal{L}_{\text{phot}}(\theta, \phi)$ and astrometric likelihood $\mathcal{L}_{\text{astr}}(\phi)$ —and $P(\theta, \phi) \equiv \pi(\theta, \phi)$ is the Galactic prior. Our photometric likelihood $\mathcal{L}_{\text{phot}}(\phi)$ compares the set of observed magnitudes $\mathbf{m}$ with the predicted magnitudes $\mathbf{m}_{\theta, \phi}$ given the associated observational errors $\hat{\sigma}$. We assume $\mathcal{L}_{\text{phot}}(\theta, \phi)$ to be independent and roughly Gaussian in the measured magnitude in each band b such that

$$\mathcal{L}_{\text{phot}}(\theta, \phi) \equiv \prod_{i=1}^b \frac{1}{\sqrt{2\pi\hat{\sigma}_i^2}} \exp \left[-\frac{1}{2} \frac{(m_i(\theta, \phi) - \hat{m}_i)^2}{\hat{\sigma}_i^2} \right]. \quad (3)$$

Likewise, the astrometric likelihood $\mathcal{L}_{\text{astr}}(\phi)$ compares the predicted parallax $\varpi(\phi)$ with the observed parallax $\hat{\omega}$ given the associated parallax error $\hat{\sigma}_{\omega}$. We also assume $\mathcal{L}_{\text{astr}}(\phi)$ to be Gaussian such that

$$\mathcal{L}_{\text{astr}}(\phi) \equiv \frac{1}{\sqrt{2\pi\hat{\sigma}_{\omega}^2}} \exp \left[-\frac{1}{2} \frac{(\varpi(\phi) - \hat{\omega})^2}{\hat{\sigma}_{\omega}^2} \right]. \quad (4)$$

The Galactic prior $\pi(\theta, \phi)$ encompasses prior beliefs on the 3D distribution and properties of stars and dust in the Milky Way. The prior is broken up into several components that are assumed to be independent:

$$\pi(\theta, \phi) \propto \underbrace{\pi(M_{\text{init}})}_{\text{IMF}} \\ \times \underbrace{\pi(d|\ell, b)}_{\text{3D number}} \\ \times \underbrace{\pi([\text{Fe}/\text{H}]_{\text{init}}|d, \ell, b)}_{\text{3D metallicity}} \\ \times \underbrace{\pi(t_{\text{age}}|d, \ell, b)}_{\text{3D age}} \\ \times \underbrace{\pi(A_V|d, \ell, b)}_{\text{3D extinction}} \\ \times \underbrace{\pi(R_V)}_{\text{Dust extinction curve}}. \quad (5)$$

Here d is the distance to the star and (ℓ, b) are the star's Galactic coordinates. We adopt the same default priors for the initial mass function (IMF), metallicity, age, and dust extinction curve as described in Appendix A of J. S. Speagle et al. (2025). Briefly, we assume that the initial masses of stars follow a broken power law following P. Kroupa (2001). For the 3D number density prior, we adopt a three-component thin disk, thick disk, and halo model on stellar distance, metallicity, and age informed by previous studies (X.-X. Xue et al. 2015; J. Bland-Hawthorn & O. Gerhard 2016; F. Anders et al. 2019). Our 3D number density prior is conceptually similar to C. A. L. Bailer-Jones et al. (2018), except they adopt a Galactic model based on a Gaia mock catalog from J. Rybizki et al. (2018). Following E. F. Schlafly et al. (2016), for the prior on the dust extinction curve we assume a mean R_V of $\mu_{R_V} = 3.32$ and standard deviation of $\sigma_{R_V} = 0.18$. This is a very tight prior on R_V , so any significant variations we may see

in R_V are likely due to our fit compensating for limitations in other areas (e.g., the stellar models) rather than capturing meaningful variations in the dust extinction curve. Unlike J. S. Speagle et al. (2024), who adopt a 3D extinction prior based on the 3D dust map of G. M. Green et al. (2019), we instead place a flat prior on extinction over the range $A_V = 0\text{--}24$ mag. The range for the flat prior was chosen to match the reddening range of G. M. Green et al. (2019), who sampled for the stellar reddening between 0 and 7 mag.

4.1.2. Application to Data

We apply the BRUTUS pipeline to all 793 million stars in the assembled catalog described in Section 3. As discussed in J. S. Speagle et al. (2025, 2024), there are systematic offsets between the MIST models and the underlying photometric data that are not fully captured by the photometric errors. We thus apply both a zero-point correction and an error floor to the photometry.

Full details on both the zero-point correction and the error floor are summarized in Table 1 of J. S. Speagle et al. (2024). Briefly, the photometric zero-point corrections range from 1% to 4% on a band-by-band basis, and are multiplied to the observed flux densities after transforming from magnitude space to flux density space. The error floors are added in quadrature to the reported observational uncertainties on the magnitudes and range from 0.02 mag for DECaPS2 up to 0.04 mag for unWISE.

In addition to the photometric corrections, we also apply a correction to the Gaia astrometry. Specifically, we implement a parallax zero-point correction (as described in L. Lindegren et al. 2021) dependent on the star's magnitude, color, and ecliptic latitude. To apply the correction, we utilize the `gaiadr3-zero-point` package.¹⁹

The `brutus` configuration used to generate our stellar modeling results, including all photometric and astrometric corrections, is available online via Zenodo (J. Speagle et al. 2025).

We generate 1000 random samples from the posterior to construct a 2D binned posterior on distance and reddening for each star, which will be used to fit the line-of-sight 3D dust model as described in Section 4.2. To generate the 2D posteriors, we use the built-in `bin_pdfs_distred` function in BRUTUS, evaluating the posterior density of each star on a grid with distance modulus spanning $\mu = 4\text{--}19$ mag in steps of 0.125 mag and reddening spanning $E = 0\text{--}7$ mag in steps of 0.01 mag. We convert from extinction to reddening using the R_V samples. Following G. M. Green et al. (2019), we smooth the posteriors, adopting a Gaussian kernel with a standard deviation equal to 1% of the total range.

In addition to the 2D distance–reddening posteriors, we use the 1000 random samples to compute the 2.5th, 16th, 50th, 84th, and 97.5th percentiles (median, 1σ , and 2σ errors) of the intrinsic and extrinsic stellar parameters. After saving the marginalized distance–extinction 2D posteriors and the percentiles, we thin the samples and save a subset of random samples to enable other science cases beyond those described in this work.

¹⁹ https://gitlab.com/icc-ub/public/gaiadr3_zeropoint

4.2. 3D Dust Modeling

In this section, we describe our approach to utilizing the posterior density estimates of distance and reddening for individual stars generated as part of the modeling in Section 4.1.2 to constrain the line-of-sight distribution of dust in the Milky Way. In Section 4.2.1, we explain the quality cuts imposed on individual stellar posteriors, and in Section 4.2.2 how we group stars into individual pixels on the celestial sphere to fit for the extinction as a function of distance along each line of sight. Finally, in Section 4.2.3, we discuss the inference framework for the line-of-sight fits, and in Section 4.2.4, we describe how we use data-driven Gaussian process priors to infill missing pixels.

4.2.1. Filtering per-Star Posteriors

To fit the line-of-sight dust distribution, we partition the sky using a HEALPix pixelization (K. M. Górski et al. 2005) with $N_{\text{side}} = 8192$ and a nested ordering scheme. Before grouping stars by HEALPix pixel, we filter out stars with unreliable modeling. First, we remove stars with poor best-fit chi-square (χ^2) values. Because χ^2 depends on the number of bands detected per source, which varies from four to 13 bands, we remove all stars for which

$$P(\chi_{n_{\text{bands}}}^2 > \chi_{\text{best}}^2) < 0.01, \quad (6)$$

with n_{bands} being the number of detected bands. Equation (6) is designed to filter stars that fail to achieve even a single reasonable fit. The term $P(\chi_{n_{\text{bands}}}^2 > \chi_{\text{best}}^2)$ denotes the χ^2 survival function, or the probability of observing a χ^2 value larger than χ_{best}^2 assuming the number of degrees of freedom is equal to the number of detected photometric bands. By filtering at the $P(\chi_{n_{\text{bands}}}^2 > \chi_{\text{best}}^2) < 0.01$ threshold, we remove stars whose best-fit model is statistically inconsistent with the data at the 99% confidence level, roughly equivalent to removing outliers at the 2σ – 3σ level.

Next, we impose a cut to remove low-mass M dwarfs from the sample. The theoretical MIST isochrones are known to exhibit large systematic biases in predicted colors at masses below $M_{\text{init}} < 0.5 M_{\odot}$ (see, e.g., J. Choi et al. 2016; J. S. Spaggle et al. 2025, 2024). Thus, we remove all stars where the 50th percentile of the M_{init} samples is less than $0.5 M_{\odot}$. As noted in Section 4.1.1, we extend the initial mass grid down to $M_{\text{init}} = 0.2 M_{\odot}$. Although stars with masses between $M_{\text{init}} = 0.2$ – $0.5 M_{\odot}$ are ultimately filtered out, extending the mass grid down to lower masses improves the line-of-sight fits, since these stars are more effectively excluded rather than potentially being mis modeled as higher-mass solutions due to known degeneracies between, for instance, dwarf and giant (see, e.g., G. M. Green et al. 2014).

Next, we impose a cut on the maximum distance to a star, removing all stars where the 2.5th percentile of the distance samples (2σ below the median distance) is greater than 30 kpc. This cut roughly corresponds to the edge of the stellar disk on the far side of the Milky Way and removes a very small fraction ($<1\%$) of spurious fits, which largely manifest in nebulous and crowded regions near $b = 0^\circ$, likely due to source blending.

Finally, we apply a cut to remove a small fraction of stars displaying a prominent, very nearby ($\mu \lesssim 8$ mag), high-reddening ($E_{B-V} \gtrsim 3$ mag) mode. This mode appears

predominantly within a few degrees of the Galactic center and does not correspond with any nearby known dense clouds. We explored several pathways to explain the existence of this mode, ruling out any obvious issues in our pipeline such as poor stellar model coverage, the quality of the stellar photometry in the most crowded area of the sky, or the more extreme variation in the extinction law toward the Galactic center. While we could not definitively identify the root cause of this issue, we know the modeling of this mode is nonphysical based on the absence of such extreme, nearby dust in high-Cartesian-resolution 3D dust maps of the solar neighborhood (see, e.g., R. H. Leike et al. 2020; G. Edenhofer et al. 2024). To address this issue, we remove stars where a majority of the probability in their 2D binned posterior on distance and reddening is inconsistent with the G. Edenhofer et al. (2024) 3D dust map. G. Edenhofer et al. (2024) provide 12 samples of the extinction density at parsec-scale resolution out to 1.25 kpc from the Sun. For each μ bin in our binned posteriors, we determine the maximum extinction density of the G. Edenhofer et al. (2024) map (over the entire sky and all 12 samples) at that distance, considering only distance bins out to the limit of the G. Edenhofer et al. (2024) map ($\mu = 10.4375$ mag).²⁰ We then sum over all probability within $\mu = 10.4325$ mag and above the maximum reddening from G. Edenhofer et al. (2024) at each distance, removing stars where $>50\%$ of their probability lies in this unphysical regime. For the remaining stars, we set this unphysical portion of the binned posterior equal to zero, and renormalize each posterior before running the line-of-sight fits.

Of the 793 million stars in the catalog, 84 million are removed via this filtering scheme, leaving 709 million stars to constrain the line-of-sight dust distribution. We hereafter refer to these 709 million stars as the “high-quality” stellar sample. Over the DECaPS2 survey footprint, we sample the 3D dust distribution in roughly 51 million $N_{\text{side}} = 8192$ pixels and 120 distance bins (over 6.1 billion total voxels) using this high-quality stellar catalog.

4.2.2. Pixelization

The angular resolution of extinction-based 3D dust maps depends on both the stellar density and chosen pixelization. If there are insufficient reliable posteriors for stars in a given resolution element, then the line-of-sight reddening for that pixel will be poorly constrained. Previous works have assigned stars to pixels with a binary weight function, including stars that fall within a HEALPix pixel and excluding them from all other pixels (e.g., G. M. Green et al. 2019).

One of the primary reasons for the high angular resolution of our map is our use of the DECaPS2 catalog, which goes deeper than PS1 and carefully deblends crowded fields to obtain an extremely high-stellar-density catalog. The other key to our high angular resolution is methodological, choosing to use a Gaussian weighting function when assigning stars to pixels, rather than a binary one. Specifically, before a star’s posterior contributes to the likelihood for a given HEALPix pixel (see Section 4.2, Equation (7)), its posterior is multiplied by a Gaussian weighting function on the great-circle distance between the star and the HEALPix pixel center, with FWHM

²⁰ We convert from the unitless extinction of the G. Edenhofer et al. (2024) map to E_{B-V} using the published extinction coefficients from X. Zhang et al. (2023).

= 2.5 pixels ($1'.07$ at $N_{\text{side}} = 8192$, which has $26''$ pixels) so it is well sampled. We truncate this weighting function at $2 \times \text{FWHM}$ (5 pixels, 4.7σ , $2'.15$ at $N_{\text{side}} = 8192$) for computational efficiency, because stars far in the tails of the weighting function do not change the line-of-sight inference significantly. Note that we implement this Gaussian weighting scheme *not* due to uncertainty on the star’s sky coordinates, but rather because it acts as a form of regularization that correlates the inference across neighboring pixels.

While this weighting scheme can increase the accessible angular resolution with a given stellar density, the maximum angular resolution is still limited by the stellar density of the photometric catalogs used. We find that $\sim 1\%$ of pixels have too few stars (<10) for a reliable line-of-sight inference when using a HEALPix pixelization with $N_{\text{side}} = 8192$. For comparison, this angular resolution is $5 \times$ higher than the integrated Planck reddening maps (Planck Collaboration et al. 2014).

One added benefit of our Gaussian weighting scheme is that, in the limit of uniform stellar density, the dust map presented here has a well-defined PSF²¹ with $\text{FWHM} = 1'.07$. This facilitates comparisons with (near-)infrared emission-based dust maps and is in contrast to many other 3D dust maps (where angular resolution varies across the sky and as a function of distance), despite most astronomers expecting image data to have a well-defined PSF. Previous 3D dust mapping works have imposed correlations between independent pixels after the fact (G. M. Green et al. 2019) or imposed a spatial correlation kernel prior at inference time (R. H. Leike & T. A. Enßlin 2019; R. H. Leike et al. 2020; G. Edenhofer et al. 2024), but neither approach yields a well-defined PSF on the sky.

4.2.3. Line-of-sight Inference

For each $N_{\text{side}} = 8192$ pixel, we split the line of sight up into discrete distance bins, evenly spaced in distance modulus from $\mu = 4$ to 19 mag, with a bin spacing of 0.125 mag. We model the dust distribution as a step function, where the parameter α denotes the increases in reddening in the set of distance bins, discretized as integer multiples of 0.01 mag. Recall that we precomputed the joint posterior density on distance modulus μ and reddening E for individual stars in Section 4.1.2, which we now denote by $\tilde{p}(\mu, E)$.

Following G. M. Green et al. (2019), the posterior probability of α is given by

$$p(\alpha | \{\mu, E\}) \propto p(\alpha) \prod_{i \in \text{stars}} \int \tilde{p}_i(\mu_i, E(\alpha, \mu_i)) d\mu_i. \quad (7)$$

$p(\alpha)$ is a prior that encompasses our expectations on the differential reddening. Following G. M. Green et al. (2019), we place a log-normal prior on the increase in reddening in each distance bin, adopting a smooth model of the distribution of dust in the Galaxy based on R. Drimmel & D. N. Spergel (2001). The term $\int \tilde{p}_i(\mu_i, E(\alpha, \mu_i)) d\mu_i$ is our likelihood function for the i th star in the pixel, and is equivalent to taking the line integral over μ through the star’s precomputed distance–extinction posterior following the distance–reddening

curve defined by α . The total likelihood is the product of the individual stellar likelihoods over all stars.

We sample for α using a custom C++ implementation²² of parallel tempered Markov Chain Monte Carlo (MCMC: J. S. Speagle 2019) that fits a model in which line-of-sight extinction is discretized at the 0.01 mag level (G. M. Green et al. 2019). We begin with a randomized initial extinction profile for each line of sight. In each distance bin, the initial guess in the jump in extinction is drawn from a chi-square distribution (with 1 degree of freedom) and the total extinction (in the large-distance limit) is normalized to equal ($\pm 20\%$) of the 90th percentile of the extinctions of the individual stars along the line of sight.

We use chains at four different temperatures (with inverse temperatures of 0.9^0 , 0.9^1 , 0.9^2 , and 0.9^3). Each chain is generated using Metropolis–Hastings sampling, with a mix of four proposal types: 33.3% “shift” steps (moving differential extinction up or down by 0.01 mag in a given distance bin), 16.7% “absolute” steps (setting the total extinction at a given distance to a random value between the extinctions of the neighboring distance bins), 33.3% “swap” steps (swapping the differential extinction at two randomly chosen distances), and 16.7% “swap neighbor” steps (swapping the differential extinction at two neighboring distances). We sample the ensemble of temperature chains in rounds of updates. During each update round, we first take 2400 (2400, 4800, 4800) steps in each temperature chain. Then, we propose a “temperature-swap” step, which swaps the state of two randomly selected temperature chains. We perform a total of 2500 (5000, 5000, 10,000) rounds of updates. The first 20% of update rounds are discarded as burn-in.

Convergence was assessed through the autocorrelation times of three variables: the prior, the likelihood, and the strength of the first principal component of the parameter vector in the MCMC chain. We require that none of the autocorrelation times exceed 5% of the total number of MCMC steps. Pixels that failed to converge were rerun a maximum of three times per setting, under each of four increasingly expensive MCMC settings (settings listed parenthetically). Empirically, we found most pixels that could “converge” under a given configuration would do so within three attempts, so we repeat computationally cheaper sampling configurations before changing configurations. Even if pixels did not “converge” under this autocorrelation metric, the result of the final (lowest-temperature) line-of-sight sampling is reported in our map, though it is flagged by an accompanying bitmask.

4.2.4. Infilling

A small fraction ($\sim 1\%$) of our pixels have too few stars for a reliable line-of-sight fit (<10 stars). These pixels occur most frequently at the edges of the DECaPS2 imaging footprint or for lines of sight with sufficient extinction to significantly suppress detection of stars behind the dust. One could simply excise such pixels from the map or fill them with an average of neighboring pixels, but both of these options are unsatisfying. We do have *some* information about the dust extinction in these pixels, based on their neighbors and the spatial correlation of dust density. We can use the extinction in neighboring pixels to make an informed prediction that gives a mean and variance for the extinction in the missing pixels,

²¹ We caution that a pixel on the edge of a dense cloud or filament may have a very nonuniform angular distribution of stars. In this case, the PSF is still somewhat ambiguous.

²² <https://github.com/andrew-saydjari/bayestar.git> v0.2.0

which is often called “inpainting” or “infilling.” While principled infilling is better than returning no value for the map at those locations, the infilled values will be biased low because we include no prior to account for the fact that the stars/pixels we are missing are preferentially more extinguished. Such pixels are flagged in the bitmask, so the user can decide whether to ignore them, or use them as a “best guess.”

We perform this infilling with a method akin to conditional Gaussian process regression, using correlations only within a single distance slice (i.e., the cumulative reddening integrated to that slice). However, because these infills are conditioned on the pixels that are *not* missing, which *are* correlated along the distance axis, the infills will be reasonably correlated along the distance axis even though we have not imposed a prior on correlations in that dimension.

We adapt the A. K. Saydjari & D. P. Finkbeiner (2022) method, designed for Cartesian pixel images, to work on ring-ordered HEALPix images far from the poles. The restriction on ordering allows a simple translation operator to be defined, which is necessary to obtain local samples of images the same size as the region of interest. We use these samples to learn a completely data-driven model of the covariance of pixels in the region of interest. This translation operator is only valid at low latitude because HEALPix rings above $|b| \gtrsim 41.8^\circ$ have a decreasing number of pixels toward the poles. Fortunately, our map at $|b| < 10^\circ$ is well within the region of validity.

First, we do a preliminary moving median infill, replacing missing pixels with a rough guess so that every sample the size of the region of interest can be used to build a model of the pixel–pixel correlations. The rough first pass uses the median of missing pixels in an 11 pixel-wide region, as long as more than 5% of the pixels are not missing. This infilling is repeated seven times to completely fill larger holes in the map.

Then, we loop over pixels missing from the original map, infill all missing pixels in a region 21 pixels wide around a given missing pixel, and repeat until no missing pixels remain. We use samples of this 21 pixel-wide region that are translates of the original region within a larger 51 pixel-wide region to learn a local mean and pixel–pixel covariance matrix. This allows us to predict missing pixels in the region of interest conditioned on the nonmissing pixels. The prediction gives a posterior on the missing pixels, so we obtain not only a mean, but also realistic samples of the missing pixels with noise properties and fluctuations consistent with the local nonmissing pixels. We stably seed the random number generator used for the draws to ensure draw consistency between different distance bins.

To ensure the self-consistency of the infill, the infills for subsequent missing pixels are conditioned on previously infilled pixels. This is performed independently for each infill draw (and the mean). Thus, while each infill is guaranteed to make use of the same pixels in their conditional prediction, the values of those pixels can differ if they were infilled in a previous step in the loop. This self-consistency comes at the cost of introducing a (slight) dependence on the infill order. Given this dependence, we make the infills at least deterministic by infilling pixels in order of their (great-circle) distance from the approximate center of our survey footprint ($\ell = 301^\circ$, $b = 0^\circ$). This iterative infill, conditioned on past infills, could extend arbitrarily far from the survey edge, but we limit it by requiring all pixels in the training samples are




Figure 5. Spatial distribution of stars over the full catalog. Each panel shows a projection of the stellar density in heliocentric Galactic Cartesian coordinates. The Galactic center is marked with a “+” symbol and the Sun with a “ $\odot$ ” symbol in the XY and XZ projections. We detect stars out to distances of $d = 15$ kpc and beyond, but with an underdensity of stars near the midplane ($z = 0$ pc) due to high levels of dust extinction.

not missing after the preliminary (moving median) infill, which sets the boundary implicitly.

We provide the mean infill for the mean map and a random draw of the infill for the samples of the map, alongside the code used to generate the infill, in Section 7. Pixels that were infilled are identified in the associated bitmask.

5. Results

This work presents two main data products: a star catalog of distance, reddening, and other parameters of 709 million stars, and the 3D dust map derived from them.

5.1. Stellar Catalog

We release the stellar catalog of distance, extinction, total-to-selective extinction ratio (R_V), and stellar type, including the median, 1σ , and 2σ errors (expressed by the 2.5th, 16th, 50th, 84th, and 97.5th percentiles of the posterior samples). We also release five random samples from the posterior. Our catalog includes all 709 million high-quality sources used in the line-of-sight dust reconstruction, selected from the 793 million stars over which we performed the stellar inference. The table (`decaps_dr2.stellar_inference`) is also available to query via NOIRLab’s AstroDataLab²³ and accessible via TAP-accessible clients, including the `astroquery` Python package.

In Figure 5, we show XY, XZ, and YZ projections of the spatial distribution of stars in the catalog in a heliocentric Galactic Cartesian frame. The underdensity of stars along the $z = 0$ pc axis is due to the high levels of dust extinction in the

²³ https://datalab.noirlab.edu/data-explorer?showTable=decaps_dr2.stellar_inference




Figure 6. Comparison between the observed CMD (apparent magnitude vs. apparent color), the absolute CMD (absolute magnitude vs. absolute color) and the absolute, dereddened CMD (dereddened absolute magnitude vs. intrinsic color) for DECaPS i vs. $i - z$. We compute the CMDs over the full high-quality stellar sample that have detections in both i and z bands, totaling roughly 670 million stars. In the absolute, dereddened CMD, we see clear overdensities corresponding to the asymptotic giant branch (the “AGB bump”) and the red clump and red giant branches (the “RC/RGB bump”), highlighting the strength of our modeling in the post-main-sequence regime.

midplane, while the dearth of stars in the immediate solar vicinity ($d \lesssim 1$ kpc) is due to both the latitude limits of the survey and the M dwarf cut we impose, given the unreliability of the MIST models in the low-mass regime. We also detect a noticeable drop-off in stellar density around $l \approx 300^\circ$, which aligns with the edge of the VVV footprint.

In Figure 6, we show a comparison between an observed color–magnitude diagram (CMD), an absolute CMD (corrected for distance), and an absolute, dereddened CMD (corrected for distance and extinction) in the DECaPS i vs. $i - z$. We compute the CMD for high-quality stars that have detections in both i and z bands, totaling roughly 670 million stars. We derive the dereddened CMD using the median (50th percentile) of the distances and extinction samples. As expected, we see a pronounced tightening of the CMD around the main sequence and giant branches, with noticeable overdensities corresponding to a blended RC and RGB bump, as well as a bump in the AGB, as marked in Figure 6.

In Figure 7, we show a Kiel diagram, which displays the surface gravity, $\log(g)$, versus effective temperature, T_{eff} . Like the CMDs, we compute the Kiel diagram over the full catalog, and take the median of the surface gravity and effective temperature samples. As expected, we see a clear bifurcation in the Kiel diagram between the main sequence ($\log(g) \lesssim 4$) and the post main sequence ($\log(g) \gtrsim 4$). Roughly 12% of the sample (89 million stars) has evolved off the main sequence, corresponding to $\text{EEP} > 454$ in the context of the MIST models. This large fraction of bright, midplane giants is critical for constraining the 3D distribution of dust at large distances, as we will show in Section 5.2. Beyond the Kiel diagram shown in Figure 7, we are also able to broadly reproduce general trends in the age–metallicity relation. However, our inferred ages and metallicities are highly prior dominated and affected by our choice of binning when constructing the model grid described in Section 4.1.1. Therefore, unlike our estimates of distance and extinction, we caution future users not to overinterpret the stellar type parameters given that a majority of stars are fit with only five bands of photometry (see Figure 2).

Finally, in Appendix D, we show that we can recover the Gaia G -band photometry. We infer the intrinsic type, distances, and extinction of stars based on other broadband photometry from a combination of DECaPS2, VVV, 2MASS, and unWISE, and do not incorporate Gaia photometry (only the astrometry) into our




Figure 7. Kiel diagram, showing the distribution of two of the inferred stellar parameters, surface gravity, $\log(g)$, as a function of their effective temperature, T_{eff} . We compute the Kiel diagram across the full sample used in the line-of-sight dust reconstruction, totaling 709 million stars.

model. However, using these inferred stellar parameters, we can predict the G -band photometry and compare this prediction to the observed G -band photometry for stars detected in Gaia. While we find overall good agreement across all SNR regimes, we systematically underpredict the Gaia G -band magnitudes by ≈ 0.05 mag, which we discuss more in Appendix D.

5.2. 3D Dust Data Product

We produce a 3D dust map with an angular resolution of $\text{FWHM} = 1'$ (HEALPix $N_{\text{side}} = 8192$ pixels of 0.43) sampled in 120 logarithmic distance bins evenly spaced between 63 pc and 63 kpc. Our map is probabilistic, and we generate 100 samples of the line-of-sight reddening (see Section 4.2.3). Our primary data product is the mean line-of-sight reddening in each distance bin. Our map is released in units of $E(B - V)$ in magnitudes. $E(B - V)$ is derived from the underlying stellar inference on A_V and R_V , where $E(B - V) = \frac{A_V}{R_V}$, with a mean of $R_V = 3.32$ based on our prior on the variation in the extinction curve from E. F. Schlafly et al. (2016). For each pixel of the map, we provide an estimate of the minimum and maximum reliable distance, which we discuss more in Appendix E. Given that we sample over 51 million HEALPix pixels on the plane of the sky and 120 distance bins along the




Figure 8. Plane-of-sky reddening in $E_B - V$ of the DECaPS 3D dust map, integrated over discrete distance intervals.

line of sight, our primary data product constitutes the reddening inferred in over 6.1 billion voxels.

Given our high angular resolution, recall that some pixels lack the minimum number of stars (≥ 10 stars) requisite to perform the line-of-sight inference, and these pixels were infilled (at each distance slice) following the procedure described in Section 4.2.4. We flag these pixels and also release a corresponding HEALPix bitmask of the infilled pixels.

The map can be downloaded at the Harvard Dataverse (see doi:10.7910/DVN/J9JCKO). The map can also be easily accessed via the `dustmaps` Python package (G. M. Green 2018), which provides a standard interface to query this map—as well as a host of other 3D and 2D reddening maps—over a user-defined set of sky coordinates and distances.

In Figure 8, we show the plane-of-sky reddening integrated over discrete distance ranges ($d < 1$ kpc, $1 \text{ kpc} < d < 2$ kpc, $2 \text{ kpc} < d < 3$ kpc, $3 \text{ kpc} < d < 5$ kpc, $5 \text{ kpc} < d < 10$ kpc, and $d \rightarrow \infty$). We see a wealth of structure over the distance range probed by our map. Within 1 kpc, we detect nearby molecular clouds including Vela C, the Pipe Nebula, and Lupus. Beyond 1 kpc, we detect massive star-forming regions like NGC 6334 and RCW 120. We also detect a diversity of feedback-driven bubbles,

including several H I shells cataloged by S. Ehlerová & J. Palouš (2013).

Like every data product, our map contains artifacts that users should be aware of, many of which only occur in specific regions of the sky and/or manifest at specific distances. These artifacts may be less obvious to the reader, since our maps are not produced the same way typical astronomical images are produced. We show and discuss these artifacts in more detail in Appendix A, but note that the boundary of the VVV survey is visible near $b = \pm 2^\circ$ primarily between $d = 2\text{--}3$ kpc across part of the fourth quadrant in Figure 8. This artifact stems from the fact that VVV goes 3 magnitudes deeper in J , H , and K compared to 2MASS. Since highly reddened stars needed to be detected behind clouds for the cloud to be detected with 3D dust mapping, the vastly different depths translate to some dense clouds being detected in VVV within $|b| < 2^\circ$ that are largely undetected and/or placed at a different distance at higher latitudes where only 2MASS data existed at the time of the map’s generation.²⁴

²⁴ After generating our 3D dust map, the VVVx survey (R. K. Saito et al. 2024) has become publicly available, which extends the latitude coverage of the original VVV survey that we utilize here. Utilizing the VVVx survey in future 3D dust maps will likely shift this artifact to higher latitudes.




Figure 9. Top-down view of DECaPS 3D dust map, integrated over $z = \pm 300$ pc. The Sun is marked with a “ $\odot$ ” symbol at $(x, y) = (0, 0)$ and the Galactic center is to the right.

In Figure 9, we show a bird’s-eye view of the southern Galactic plane, integrated over $z = \pm 300$ pc. Despite evidence of “Fingers of God”—owing to our angular resolution being finer than our distance resolution—we detect a significant amount of dense structure in the fourth quadrant between $d = 2$ – 4 kpc previously unresolved in purely Gaia-based maps. Particularly at $l < 300^\circ$, we also detect discrete dust complexes between $d = 5$ – 10 kpc. We observe several kiloparsec-sized voids in this top-down view, akin to those seen in nearby face-on galaxies with JWST (J. C. Lee et al. 2022).

6. Discussion

Here we highlight the two greatest strengths of the map in the context of existing approaches: its high angular resolution (Section 6.1) and its depth (Section 6.2). In Section 6.2, we also combine the DECaPS 3D dust map with the complementary Bayestar19 3D dust map in the northern sky to provide complete coverage of the plane $|b| < 10^\circ$. The combination of the map’s angular resolution and its ability to resolve structure at large distances provides an important proof of concept for the future of 3D dust mapping in the era of LSST and Roman (see Section 6.3).

6.1. Herschel-resolution 3D Dust Mapping

As detailed in Section 4.2.2, the angular resolution of 3D dust maps is directly tied to the density of stars on the plane of the sky and the chosen pixelization. Thanks to the increase in plane-of-sky stellar density made possible by the DECaPS2 survey, the data support a pixelization at HEALPix $N_{\text{side}} = 8192$, with an angular resolution defined by a $1'$ Gaussian PSF ($2.5\times$ the 0.43 pixel scale at HEALPix $N_{\text{side}} = 8192$) across the full footprint.

In Figure 10, we show a cutout of our $1'$ resolution map alongside a cutout from the Bayestar19 3D dust map (G. M. Green et al. 2019) and the Herschel far-infrared emission map at $500\ \mu\text{m}$ (S. Molinari et al. 2016) toward a dense filament in the Galactic plane centered at $(l, b) = (242.5, -1^\circ)$. Both the Bayestar19 3D dust map and the DECaPS 3D dust map have been integrated out to a distance of 7 kpc, roughly the maximum reliable distance of the DECaPS map toward the filamentary dust structure in this region of the sky. While the angular resolution varies across the Bayestar19 3D dust map, G. M. Green et al. (2019) estimate a typical resolution of $6.8''$ (see their Section 2.3.1). For comparison, the Herschel emission map has an angular resolution of $37''$ at $500\ \mu\text{m}$.

Figure 10 clearly highlights the strength of our map’s angular resolution and sensitivity: We are able to capture a majority of the rich filamentary structure seen with Herschel in dense regions. Our map constitutes the highest-angular-resolution 3D dust map available today, achieving $7\times$ higher angular resolution than Bayestar19 and an angular resolution within a factor of 2 of the Herschel $500\ \mu\text{m}$ emission maps. However, unlike Herschel, we have distance information and can place the filament in 3D within its broader Galactic environment given the significant depth of our 3D dust map, as discussed in Section 6.2. While Figure 10 demonstrates the strength of our map, we also caution that compared to 2D approaches, we still underpredict the total amount of reddening in very dense regions, particularly toward the Galactic center. In Appendix A, we discuss the limitation of our map toward the Galactic center, by comparing our projected 3D dust map with the 2D reddening map of F. Surot et al. (2020), which relies solely on infrared VVV photometry.




Figure 10. Plane-of-sky comparison of a dense filament in the Galactic plane. The top two panels show the projected Bayestar19 3D dust map and the DECaPS 3D dust reddening maps, integrated out to a distance of 7 kpc. The bottom panel shows the 2D Herschel 500 μm emission map, which is by definition integrated to infinity.

6.2. Depth

Alongside possessing an angular resolution comparable to the Herschel 500 μm emission maps, our approach resolves

structure substantially deeper into the Galactic plane than most existing 3D maps. Unlike the angular resolution of 3D dust maps, which is tied to the density of stars on the plane of the sky (Section 6.1), the depth of 3D dust maps depends on the

ability to resolve a sufficient number of highly reddened stars at large distances. If reddened stars are not detected background to a dense cloud, that cloud will not appear in the resulting map, even if the cloud nominally lies within the distance range targeted in the reconstruction. In contrast to most existing approaches, 76% of our stellar sample (541 million stars) lies at $d > 3$ kpc, 43% (303 million stars) at $d > 5$ kpc, and 17% (119 million stars) at $d > 7$ kpc, underlining the ability of our map to resolve structure significantly beyond the solar neighborhood.

We can characterize the depth of our maps in the context of two regimes: those that require Gaia parallax measurements and those that do not. In contrast to Gaia-reliant approaches, our inference pipeline is flexible enough to infer stellar distances in the absence of parallaxes, allowing us to probe much larger volumes of the Galaxy. In addition, in comparison to maps that also do not depend on Gaia, we incorporate deeper photometry and more sophisticated modeling of highly reddened post-main-sequence stars, allowing us to penetrate further into the midplane. Our pipeline’s flexibility to incorporate and model deeper photometry in the absence of Gaia translates to a data product capable of resolving dense clouds up to a depth of roughly $d = 10$ kpc. However, the exact depth varies across the sky, particularly as a function of latitude, and is estimated in part by our maximum reliable distance calculation (see Appendix E).

6.2.1. Comparisons within the Solar Neighborhood

Most modern 3D dust maps (B. Q. Chen et al. 2019; R. Lallement et al. 2019; J. L. Vergely et al. 2022; T. E. Dharmawardena et al. 2024) rely on higher-quality Gaia parallax measurements to constrain stellar distances, a key ingredient for 3D dust mapping alongside stellar extinction estimates. As an optical instrument, Gaia has a limiting magnitude of $G \approx 21$ mag for stars with a full astrometric solution (Gaia Collaboration 2022). Maps that require Gaia detections typically lose a sufficient number of background sources at $d \gtrsim 2$ kpc in the inner Galaxy due to the high levels of dust extinction. Therefore, most Gaia-based maps focus on the solar neighborhood ($d < 2\text{--}3$ kpc; R. Lallement et al. 2019; R. H. Leike et al. 2020; G. Edenhofer et al. 2024; T. E. Dharmawardena et al. 2024) and caution that the detection of structure is inhomogeneous and sparsely sampled at the edges of their maps.

In Figure 11, we show a comparison between the Gaia-based 3D dust map from J. L. Vergely et al. (2022; extending out ± 3 kpc in X and Y with a Cartesian resolution of 10 pc) alongside the DECaPS 3D dust map. In comparison to the J. L. Vergely et al. (2022) map shown in Figure 11, our pipeline does *not* require a Gaia parallax, and 46% of stars in our sample (366 million stars) are not detected in Gaia and modeled using photometry alone. Thanks to the flexibility of our pipeline to model stars in the absence of parallaxes, Figure 11 illustrates the vastly different scale probed by the DECaPS map. Even focusing on the region of overlap, Figure 11 shows that we also detect more clouds between $d = 2\text{--}3$ kpc and recover more extinction per cloud, indicating that we are detecting more highly reddened stars background to each cloud complex. We further emphasize this point in Figure 12, where we compare the extinction profiles (the cumulative extinction as a function of distance) between J. L. Vergely et al. (2022) and our DECaPS map along three




Figure 11. Top-down view of DECaPS 3D dust map (top) alongside the J. L. Vergely et al. (2022) map (bottom, restricted to the DECaPS footprint), each integrated over $z = \pm 400$ pc. The J. L. Vergely et al. (2022) map has been restricted to $239^\circ < l < 6^\circ$ to facilitate comparison with DECaPS over the same longitude range. The Sun is marked with a “ $\odot$ ” symbol at $(x, y) = (0, 0)$. Both are displayed on the same relative extinction scale and saturate at the 99.95% percentile of the respective reconstruction. Spiral arm models from M. J. Reid et al. (2019) for the Sagittarius-Carina (Sgr-Car) and Scutum-Centaurus (Sct-Cen) arms are overlaid. Within $d \lesssim 5$ kpc, we find that a majority of the massive cloud complexes lie in the interarm region, deviating from log-spiral fits in this region of the Galaxy. An interactive version of this figure that allows you to flash back/forth between the panels is available. A version is also available at the author’s Harvard website (https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/DECaPS_Marshall_Comparison.html). An interactive version of this figure is available in the [online article](#).

lines of sight in the Galactic plane out to a distance of 3 kpc from the Sun. While the maps largely agree at low extinction within $d \lesssim 1$ kpc and $A_V \lesssim 2$ mag, we find significant differences at larger distances, with DECaPS detecting a factor of a few times more extinction per cloud compared to J. L. Vergely et al. (2022).²⁵

Therefore, while Gaia-based solar neighborhood 3D dust maps achieve a higher distance resolution in the nearest 2 kpc, the DECaPS map recovers more dense clouds in the southern Galactic plane at $d \gtrsim 2$ kpc, where most of the star formation is occurring along the Sagittarius-Carina and Scutum-Centaurus arms (see, e.g., the distribution of YSOs in the southern Galactic plane from M. A. Kuhn et al. 2021).

In Figure 11, we overlay the traces for the Sagittarius-Carina and Scutum-Centaurus arms, constrained by maser parallax measurements toward high-mass star-forming regions from M. J. Reid et al. (2019). We find that while we recover large cloud complexes at the correct *range* of distances, most of the dust lies in the *interarm* region (between the near Sagittarius-

²⁵ The J. L. Vergely et al. (2022) map is provided in units of $A_{550\text{nm}}$, equivalent to extinction in the V band. We query the J. L. Vergely et al. (2022) extinction profiles using the G-TOMO tool at <https://explore-platform.eu>. For this comparison, we convert our DECaPS 3D dust map from $E(B - V)$ to A_V assuming an $R_V = 3.32$.




Figure 12. Comparison of the cumulative extinction as a function of distance between the J. L. Vergely et al. (2022) map (orange profiles) and the DECaPS map (purple profiles) toward three lines of sight in the Galactic plane. The DECaPS map recovers significantly more extinction per cloud (particularly beyond 1 kpc) compared to J. L. Vergely et al. (2022).

Carina and near Scutum-Centaurus arms, at distances $d \lesssim 4$ kpc) as defined by M. J. Reid et al. (2019). Given the dearth of maser measurements in the fourth quadrant, we anticipate substantial revision may be needed to spiral arm models in the fourth quadrant, which can be informed by a combination of this 3D dust map, future maser parallax measurements, and complementary spatial-dynamical constraints on the Milky Way’s ISM (e.g., diffuse interstellar bands, or DIBS; K. Tchernyshyov et al. 2018; A. K. Saydjari et al. 2023a) in the southern Galactic plane.

6.2.2. Comparisons beyond the Solar Neighborhood

As established in Section 6.2.1, maps that do not depend on Gaia to infer stellar distances are able to probe much larger volumes of the Galaxy. Several such 3D dust maps currently exist in the literature, and we focus here on comparisons with G. M. Green et al. (2019) and D. Marshall et al. (2025), the latter of which is an update to the original 3D dust mapping work from D. J. Marshall et al. (2006). S. Rezaei Kh. et al. (2024) also reconstructs the distribution of dust out to $d = 10$ kpc. However, we forgo detailed comparisons with S. Rezaei Kh. et al. (2024) since they rely on APOGEE spectroscopy for their 3D dust reconstruction. APOGEE Data Release 16 is so sparsely sampled in the southern Galactic plane that no stars are incorporated in much of the fourth quadrant (see their Figure 2) and only 44,000 stars are included in the sample over the remainder of the plane.

The Bayestar19 3D dust map (G. M. Green et al. 2019) is a complementary map to the DECaPS map and reconstructs the 3D dust distribution over three-quarters of sky in the north, primarily based on PS1 photometry for 799 million stars. While G. M. Green et al. (2019) use an empirical set of stellar models rather than the theoretical MIST models we adopt here, the procedure for determining stellar distances is similar to the methodology outlined in Section 4.1.1: Measured parallaxes

and their corresponding uncertainties are incorporated into the per-star modeling as an additional likelihood where available; otherwise, stellar posteriors on distance and extinction are inferred using a minimum of four photometric passbands, as in this work. In Figure 13, we combine the Bayestar19 3D dust map with the DECaPS 3D dust map to provide complete 360° coverage of the Galactic plane. The combination of the two maps provides complete coverage of the plane $|b| < 10^\circ$. In Figure 13, we show a top-down view of the plane integrated over $|z| < 100$ pc to highlight differences in the maps near the midplane.²⁶ Close to $z = 0$ pc, we find that the DECaPS map detects more dense clouds at distances between $d = 2\text{--}6$ kpc compared to Bayestar19.

We attribute the greater depth of the DECaPS map to two causes. First, we incorporate significantly deeper infrared photometric data, allowing us to detect more stars background to dense clouds at larger distances. G. M. Green et al. (2019) combine PS1 photometry with 2MASS photometry, but due to its shallow depth, only 10%–20% of PS1 sources also have 2MASS detections. While we also utilize 2MASS data, we primarily rely on VVV photometry in the infrared near the midplane, which is 3 magnitudes deeper than 2MASS and provides complementary detections to DECaPS for a majority of sources in the VVV footprint. In addition to the incorporation of deeper infrared data, we also employ more sophisticated modeling of post-main-sequence stars, allowing us to detect highly reddened giants at larger distances. G. M. Green et al. (2019) derive empirical stellar templates based on observations of uniformly old stellar populations toward globular clusters with SDSS and only approximate the morphology of the giant branch (see discussion in G. M. Green et al. 2014). In contrast, the MIST models encompass a broader range of evolved evolutionary phases (up to the beginning of the thermally pulsing AGB phase) and extend to higher initial masses (for a comparison between the Bayestar19 and MIST post-main-sequence models, see Figure 8 in J. S. Speagle et al. 2024).

In Figure 14, we compare with the D. Marshall et al. (2025) 3D dust map, which is the most similar to the DECaPS 3D dust map in terms of depth. D. Marshall et al. (2025) adopts a 3D stellar distribution based on the Besançon Galaxy model (see A. C. Robin et al. 2003; M. A. Czekaj et al. 2014). The Besançon model provides the intrinsic (unreddened) colors of simulated stars in the 2MASS bands, which can be compared with the observed reddened 2MASS photometric colors to infer the near-infrared color excess along each line of sight. Our map is qualitatively similar to the D. Marshall et al. (2025) reconstruction, and we place several dense clouds in the third and fourth quadrants at similar distances. In some cases, we find the same features, but systematically offset in distance (see, e.g., dense clouds between $l = 240^\circ$ and 270°), which may be the result of systematic uncertainties in the Besançon model used in the reconstruction of D. Marshall et al. (2025). As raised in Section 6.2.1, we argue that there is scant delineation between the dust associated with the Scutum-Centaurus and Sagittarius-Carina arms within $d \lesssim 5$ kpc in the DECaPS dust map. However, D. Marshall et al. (2025) argues for more clearly defined arms and finds evidence of the Centaurus and Carina tangencies in the fourth quadrant,

²⁶ We mask out the Bayestar19 3D dust map at latitudes $|b| > 10^\circ$ before integrating over z to facilitate a one-to-one comparison to the DECaPS map, which only spans $|b| < 10^\circ$.




Figure 13. Top-down view of the Bayestar19 3D dust map (red; G. M. Green et al. 2019) combined with this DECaPS 3D dust map (blue), both integrated over $|z| < 100$ pc and limited to $|b| < 10^\circ$.

particularly on scales spanning 10 kpc in the bird’s eye view in Figure 14, where D. Marshall et al. (2025) is more sensitive to large-scale structure. While beyond the scope of this work, a more detailed comparison of the nature of spiral structure in the fourth quadrant is worthy of follow-up investigation. For now, see the interactive version of Figure 14 to flash back and forth between the D. Marshall et al. (2025) and DECaPS maps on the same grid.

6.3. The Future Era of LSST and Roman

The future of 3D dust mapping lies not in the era of Gaia but in the era of LSST and Roman. The fourth and fifth data releases (DR4 and DR5) of Gaia circa 2026 and 2030 will herald even more precise astrometry, with parallaxes expected to improve by factors of $1.4\times$ and $1.9\times$, respectively (Gaia Collaboration et al. 2023). However, future Gaia data releases will not substantially improve the limiting magnitude of the survey ($G \approx 21$ mag), so the number of DECaPS stars that have an accompanying parallax detection in Gaia DR4 and Gaia DR5 will not substantially increase. Therefore, studies

that rely on Gaia alone will never be able to push significantly beyond the solar neighborhood, either now or in the future.

Recall from Section 6.2 that 46% of stars in our sample (366 million stars) do not have a Gaia parallax measurement: Their stellar type, distance, and extinction are inferred using photometry alone. The future era of LSST and Roman will be an era of even deeper photometry than utilized here, and we need to be prepared to meet it. Expected to start science operations in 2025, LSST will target the southern Galactic plane between $\lambda = 320$ and 1050 nm in the u , g , r , i , z , and y bands. While the filter balance and footprint are still being finalized, the LSST strategy will consist of a high-visit region near the bulge and a thick strip of the Galactic plane in the fourth quadrant as part of its Galactic plane “Wide Fast Deep” (or WFD) survey, alongside a low-visit “Dusty Plane” survey that fills in the remainder of the third quadrant (LSST 2024). LSST is expected to produce a photometric catalog of roughly 20 billion stars over its 10 yr duration, with a typical 5σ single-visit point-source depth in r of 24.5 mag in the AB system (Ž. Ivezić et al. 2019).




Figure 14. Comparison of the DECaPS dust map (top panel) and the D. Marshall et al. (2025) 3D dust map (bottom panel, restricted to the DECaPS footprint). To better highlight structures over a narrow latitude range, we show the average differential extinction, rather than the integral (as in Figure 11). Both are displayed on the same relative extinction scale and saturate at the 99.5% percentile of the respective reconstruction. Arm models from M. J. Reid et al. (2019) for the Sagittarius-Carina (Sgr-Car) and Scutum-Centaurus (Sct-Cen) arms are overlaid. An interactive version of this figure that allows you to flash back/forth between the panels is available. A version is also available at [the author’s Harvard website](#).

An interactive version of this figure is available in the [online article](#).

Complementing LSST, Roman will target the Galactic plane in the first 2 yr of its mission as part of its General Astrophysics Survey program (R. E. Sanderson et al. 2024). Launching by mid-2027, Roman’s Wide-Field Instrument has eight science filters spanning $0.48\text{--}2.3\ \mu\text{m}$ (NASA 2024), where the longest-wavelength F213 filter will prove critical for peering through highly extinguished regions in the plane. Like LSST, the filter balance and footprint for the Roman Galactic Plane Survey are still being defined, but is expected to reach a 5σ point-source depth of 25.4 mag in F146 in the AB system, assuming a typical 57 s integration (NASA 2024). Roman has the potential to provide multiband photometric imaging for tens of billions of stars (R. Paladini et al. 2023). Therefore, either alone or in combination, LSST and Roman will produce photometry for at least an order of magnitude more stars in the Galactic plane than currently available.

To match the potential of these next-generation photometric surveys, we anticipate three required methodological improvements. First, we need to employ faster inference frameworks. Second, we need to develop more accurate templates for the intrinsic colors of stars that capture both low initial stellar masses and a broad range of post-main-sequence evolutionary phases. And third, we need to implement more nuanced modeling of the extinction curve.

The stellar inference pipeline we present here is effective but slow—on average, BRUTUS took 4 s/star^{-1} to generate our stellar catalog (see Appendix B). Machine learning approaches are more easily scalable and offer an alternative to our brute-force Bayesian inference pipeline, producing outputs in a

fraction of a second per star once the model is trained. For example, X. Zhang et al. (2023) developed a data-driven model to infer stellar atmospheric parameters, distances, and extinctions for stars in the Gaia XP catalog (see also data-driven models from G. M. Green et al. 2021). The model was trained on the 1% of Gaia XP stars with stellar atmospheric labels (T_{eff} , $\log(g)$, and $[\text{Fe}/\text{H}]$) from the LAMOST survey. Training took 25 hr on one GPU node and was used to infer the parameters for all 220 million Gaia XP stars in roughly 36 hr on 2–4 GPU nodes (X. Zhang & G. M. Green 2024, private communication), making this type of approach very computationally efficient in the era of LSST and Roman.

One potential challenge of data-driven approaches is that the quality of the inference is predicated upon capturing a broad range of stellar types in the spectroscopic training data set: If a certain type is not represented in training, those stars will be mismodeled in the final catalog. Two of the most critical stellar types to capture in the era of LSST and Roman are stars with low initial stellar masses (e.g., M dwarfs) and those that have evolved off the main sequence (e.g., RGB, horizontal branch, and AGB stars). Roman will detect low-mass dwarfs out to a few kiloparsecs, representing a critical foreground population for anchoring the distribution of dust at larger distances. Likewise, strong coverage of the post main sequence will allow future dust maps to model highly reddened giants at distances far beyond the distances probed in this work, enabling the construction of 3D dust maps beyond the Galactic center. The data-driven models from X. Zhang et al. (2023) have poor coverage of low-mass M dwarf solutions, as well as sparser coverage of post-main-sequence evolution than the MIST models we employ here. Therefore, incorporating new and upcoming spectroscopic surveys like SDSS-V’s Milky Way Mapper (A. Almeida et al. 2023) will prove critical for accurately modeling these stellar types in the context of future data-driven machine learning approaches.

Finally, future 3D dust mapping frameworks will need to implement more sophisticated modeling of the 3D variation in the extinction curve. Improved modeling of the intrinsic colors of stars (across a broader range of spectral types) will only go so far if the modeling of the extinction curve is too simplistic. Recall from Section 4.1.1 that the shape of the extinction curve is parameterized by the total-to-selective extinction ratio, R_V . In this work, we adopt a tight prior on R_V , with a mean $R_V = 3.32$ and a standard deviation of $\sigma_{R_V} = 0.18$ based on E. F. Schlafly et al. (2016). We also assume that the extinction curve is independent of stellar type. However, the dust extinction curve depends on the underlying stellar spectrum and is known to vary substantially throughout the Milky Way (X. Zhang & G. Green 2025). Failure to capture these effects could, for example, create systematics in the modeling of A_V that would be nearly indistinguishable from changes in a star’s effective temperature T_{eff} . SDSS-V’s Milky Way Mapper is currently targeting a much larger number of stars at high extinction. These measurements will prove critical for studying variations in the extinction curve across the Galaxy at large, which can be folded into future inference frameworks in the era of LSST and Roman.

7. Code and Data Availability

The software to reproduce the per-star and line-of-sight inference is publicly available on Zenodo at doi:[10.5281/zenodo.16813633](https://doi.org/10.5281/zenodo.16813633). The stellar inference pipeline is based on

the publicly available BRUTUS software package (v0.8.3), which is publicly available on Zenodo (J. Speagle et al. 2025).

For the per-star inference, we release the 2.5th, 16th, 50th, 84th, and 97.5th percentiles of the samples for the stellar properties, distance, extinction, and the total-to-selective extinction ratio of each star in our high-quality catalog (709 million stars; see Section 4.2.1), which have been computed from all 1000 posterior samples prior to thinning. We also release five random samples of distance, extinction, total-to-selective extinction ratio, and model indices, the latter of which can be translated to samples of stellar type given the underlying BRUTUS model grid. In addition to being available on the Dataverse at doi:10.7910/DVN/K88GFI, the stellar catalog is archived in AstroDataLab (decaps_dr2_stellar_inference),²⁷ so it is accessible via TAP-accessible clients, including the astroquery Python package.

For the line-of-sight inference, we release the mean 3D reddening map (computed from all 100 samples) and five random samples (in units of E_{B-V} in mag) as the core data product, which is publicly available for download on the Dataverse at doi:10.7910/DVN/J9JCKO. E_{B-V} is derived from the underlying stellar inference on A_V and R_V , where $E(B-V) = \frac{A_V}{R_V}$, with a mean of $R_V = 3.32$ based on our prior on the variation in the extinction curve from E. F. Schlafly et al. (2016). This core data product is released in HEALPix $N_{\text{side}} = 8192$ format, where we provide the reddening in 120 logarithmically spaced distance bins. Recall from Section 4.2.4 that we infill roughly 1% of pixels in our footprint due to an insufficient number of stars. All five released samples have been infilled, alongside the mean map. We release a number of quality flags, including whether the line-of-sight fit converged in a given pixel, whether the pixel was infilled, and the minimum and maximum reliable distance modulus in that pixel (see Appendix E). We also provide the number of stars used to inform the line-of-sight fit for each $N_{\text{side}} = 8192$ pixel.

In addition to being available on the Dataverse, the map can also be queried via the Python package dustmaps using the DECaPSQuery class, which is the recommended way of performing extinction corrections. We also provide the functionality to query with memory mapping (DECaPSQueryLite) so the entire map does not have to be read into memory. Via dustmaps, users can also combine the DECaPS query with the existing Bayestar query to perform extinction corrections across the entire disk $|b| < 10^\circ$, and an example is provided online in the dustmaps documentation.

8. Conclusion

We present a deep, high-angular-resolution 3D dust map over the range $239^\circ < l < 6^\circ$ and $|b| < 10^\circ$. We start by inferring the distance, extinction, and stellar type of almost 1 billion stars using optical and infrared photometry from the DECaPS2, VVV, 2MASS, and unWISE surveys. Unlike most solar-neighborhood-based maps, we do not require Gaia parallax measurements, though we incorporate these distance constraints when available (roughly half the sample). We then group the stars into pixels, and fit the distance and extinction measurements of stars in each pixel to infer the distribution of dust along 51 million lines of sight in the southern Galactic

plane, totaling over 6 billion voxels. Our main conclusions are as follows:

1. Thanks to the increased stellar density provided by the DECaPS2 survey, we produce a 3D dust map with an angular resolution of $\text{FWHM} = 1'$, which is roughly an order of magnitude finer than any existing 3D dust map and on par with the angular resolution of the Herschel 2D dust emission maps.
2. The flexibility of our pipeline to model stars using photometry alone allows us to probe significantly deeper into the Galactic plane than most current maps, targeting regions inaccessible to Gaia. We resolve cloud complexes toward the nearby Sagittarius-Carina and Scutum-Centaurus arm as well as more distant complexes lying between 6 and 10 kpc from the Sun. However, we find that most of the cloud complexes lie in the interarm regions (between Sagittarius-Carina and Scutum-Centaurus), suggesting the need for substantial revision to spiral arm models in the fourth quadrant, where constraints are currently limited by a lack of maser parallax measurements.
3. Our map fills in the one-third of the Galactic plane absent from the Bayestar19 3D dust map (G. M. Green et al. 2019). By combining Bayestar19 with our new map, we enable extinction corrections over the entire disk within $|b| < 10^\circ$.

Our map serves as a valuable proof of concept for the future of 3D dust mapping in the era of LSST and Roman, which will provide deep optical and infrared photometry for tens of billions of stars. In this work, we focused on pushing the angular resolution frontier, in contrast to, for example, the G. Edenhofer et al. (2024) map, which pushed the limits of distance resolution. As we transition from voxel counts of several billion in this current work to the trillion that may be possible with LSST and Roman, simply storing and handling maps of this size will become challenging, and we may need to move beyond voxelization. Improved storage and data handling—alongside the development of faster stellar inference frameworks, improved modeling of low-mass and post-main-sequence stars, and more sophisticated treatments of the dust extinction curve—will prove critical for fully realizing the potential of 3D dust mapping in the coming decade.

Acknowledgments

C.Z. acknowledges insightful discussions with Doug Marshall, Mark Reid, Tom Dame, Michael Rugel, Philipp Frank, and Torsten Enßlin. C.Z. thanks Doug Marshall for facilitating the comparison with his 3D dust map in Figure 14. C.Z., D.P.F., and A.G. acknowledge support by NASA ADAP grant No. 80NSSC21K0634, “Knitting Together the Milky Way: An Integrated Model of the Galaxy’s Stars, Gas, and Dust.” A.K.S. acknowledges support by a National Science Foundation Graduate Research Fellowship (DGE-1745303) and that support for this work was provided by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51564.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. A.K.S. acknowledges Sophia Sánchez-Maes for helpful discussions and much support. J.S.S. was supported by funding from

²⁷ https://datalab.noirlab.edu/data-explorer?showTable=decaps_dr2_stellar_inference

the Dunlap Institute, an NSERC Banting Postdoctoral Fellowship, NSERC Discovery grant No. RGPIN-2023-04849, and a University of Toronto Connaught New Researcher Award. J.S.S. would also like to thank Rebecca Bleich for not stealing too many french fries during the process of preparing and revising this manuscript. The authors thank the anonymous referee for their prompt and constructive review of our manuscript that improved the quality of this work.

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.

This research has used data, tools, or materials developed as part of the EXPLORE project that has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No. 101004214.

A portion of this work was enabled by the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

Software: `astropy` (Astropy Collaboration et al. 2022), `brutus` (J. S. Speagle et al. 2025), `glue` (T. Robitaille et al. 2019), `numpy` (C. R. Harris et al. 2020), `scipy` (P. Virtanen et al. 2020), `healpy` (K. M. Górski et al. 2005; A. Zonca et al. 2019), `dustmaps` (G. M. Green 2018).

Appendix A Artifacts

Here we highlight several artifacts in our 3D dust map that potential users should be cognizant of. Many of these artifacts preferentially manifest in certain parts of the sky and/or over a certain range in distances. Likewise, some artifacts may only be visible if the map is stretched in a particularly extreme way. A few of these artifacts are captured by our accompanying bitmask, including for example the known limitations of the map outside our minimum and maximum reliable distance range.

In Figure 15, we show a gallery of known artifacts, where each panel highlights a different artifact. In panel (a), we show the limitations of our map at very close distances toward a local molecular cloud, using the Pipe nebula ($d = 152$ pc; C. Zucker et al. 2021) as an example. We integrate our map out to $d = 200$ pc (beyond the nominal distance of the Pipe nebula) to show that while we recover the gross morphology of the nebula, the map is very patchy at close distances. The patchiness stems from a combination of our M dwarf cut (removing stars foreground to the cloud; see Section 4.2.1), our high-angular resolution (see Section 4.2.2), and our lack of a spatial regularization scheme to correlate distance slices across neighboring pixels. This artifact is typically captured by our minimum reliable distance bitmask.

In panel (b), we show a speckling that appears in our map over a small fraction of the sky when integrated out to large distances. As an example, we show a piece of the Herschel filament originally seen in Figure 10 but instead of integrating out to $d = 7$ kpc (where no artifacts are seen), we integrate out to the edge of the map. We are unable to explain the root cause of this artifact, but note that it is typically captured by our maximum reliable distance bitmask.

In panel (c), we highlight the artifact associated with the VVV boundary. We show an 8° longitude strip of the Galactic plane integrated out to $d = 2$ kpc, where the VVV boundary is clearly visible at $b = +1.75^\circ$. This artifact is also visible over a broader distance range, manifesting between $d = 1\text{--}5$ kpc. This artifact stems from the fact that not only do we detect more highly reddened stars with VVV compared to 2MASS, but the inclusion of VVV photometry alongside 2MASS photometry (in regions of overlap) can modify the underlying distance–extinction posterior, which can cause clouds to be placed at slightly different distances above and below the boundary.

In panel (d), we show a striping artifact that appears at the edges of the DECaPS map (near $b \pm 8^\circ\text{--}10^\circ$) when integrated out to $d \rightarrow \infty$. This artifact stems from chip gaps in the underlying DECaPS survey, causing stars within the gaps to have fewer photometric detections, which again modifies the underlying distance–extinction posteriors, and propagates to the inferred amount of reddening within the gaps.

Finally, in panel (e), we show an artifact that manifests in all dense regions across the plane, where the cores of dense clouds seem to be “missing” in the sense that they appear to be in “shadow” (at much lower extinction) than their envelopes. As an example, we show a cloud envelope (integrated out to $d = 3$ kpc) at $E_{B-V} \approx 3\text{--}4$ mag, where the core of the cloud lies at lower extinction ($E_{B-V} \approx 1$ mag). This artifact stems from the fact that the cloud is so dense that it is fully extinguishing, rather than simply reddening, the light from stars background to the clouds, so the core of the cloud is undetectable in 3D dust maps.

In Figure 16, we better illustrate the DECaPS map’s inability to recover the total extinction in very dense regions. Specifically, we show the ratio of the reddening of our projected DECaPS map (integrated out to 8.5 kpc, roughly the distance of the Galactic center) over the reddening obtained from the 2D infrared star-based dust map of F. Surot et al. (2020). We convert the F. Surot et al. (2020) map in units of E_{J-K_s} to E_{B-V} assuming $E_{J-K_s} = 0.57 \times E_{B-V}$ based on our reddening vector in Table 2. Unlike our DECaPS map, F. Surot et al. (2020) leverages VVV infrared photometry toward RC and RGB stars and does not require a detection in the optical. Therefore, the F. Surot et al. (2020) map is able to detect more stars in dense regions and recover more of the total extinction. However outside $|b| \lesssim 0.5^\circ$, we generally find strong agreement with the F. Surot et al. (2020) 2D dust map.

Gallery of Artifacts




Figure 15. Gallery of known artifacts in the DECaPS 3D dust map. Each panel highlights a different artifact, where the stretch and the distance integration range (summarized in the top-left corner of each panel) were chosen to highlight each artifact in its most extreme form. See Appendix A for a description of the artifacts.




Figure 16. Ratio of the reddening of the DECaPS map (integrated out to 8.5 kpc) over the reddening from the 2D F. Surot et al. (2020) dust map. We find good agreement except for in the midplane, where we underpredict the reddening with respect to 2D maps based solely on infrared photometry.

Appendix B Computational Details

The generation of this 3D dust map took significant computational resources. In total, we utilized roughly 3.5 million CPU hours on the FASRC Cannon cluster, supported by Harvard’s FAS Division of Science Research Computing. The per-star inference took $\approx 887,000$ CPU hours, primarily with 9 GB per process. Across all 793 million stars, we averaged roughly 4 s/star. The line-of-sight inference took ≈ 2.6 million CPU hours, ranging from 4 to 8 GB per process. Across all 51 million pixels, we averaged roughly 182 s per pixel. Infilling the 3D dust map took ≈ 320 CPU hours, with ≈ 90 GB per process.

Appendix C Inferred R_V Distribution

Recall that we infer the total-to-selective extinction ratio, R_V , toward every star as part of the methodology described in Section 4.1.1. In Figure 17, we show the distribution of R_V , tabulated across the full sample of 709 stars used in the generation of the 3D dust map. For each star, we draw a random R_V sample. As is apparent in Figure 17, we caution that R_V is not strongly constrained by our data and closely aligns with the adopted mean R_V of $\mu_{R_V} = 3.32$ and standard deviation of $\sigma_{R_V} = 0.18$ of our prior.




Figure 17. Normalized probability density function (PDF) of the distribution of R_V inferred per star for the high-quality stellar sample (709 million stars). For each star, we draw a random sample. The distribution closely follows our adopted prior, limiting potential implications of our work for understanding the variation in the extinction curve across the Milky Way.

Appendix D Predicting Gaia G -band Magnitudes

While we do not incorporate Gaia photometry into our stellar modeling pipeline described in Section 4.1.1, the Gaia photometry provides an additional means to verify the fidelity of our inferred stellar parameters. For each star with a Gaia detection (roughly half the sample see; see Figure 4), we predict the Gaia G -band magnitudes based on the inferred stellar type constrained using the photometry from other wide-




Figure 18. Offset between the predicted and observed Gaia G -band magnitudes, computed as an average per $N_{\text{side}} = 1024$ over the full DECaPS2 footprint for stars with a Gaia detection. While we find good agreement overall, we tend to systematically predict slightly smaller (brighter) Gaia G -band magnitudes than observed, at the level of ≈ 0.05 mag.

field surveys (see Section 2). In Figure 18, we show the difference between the predicted and observed Gaia G -band magnitude, averaged in $N_{\text{side}} = 1024$ pixels over the full DECaPS2 footprint. We overall find good agreement, but underpredict the observed Gaia G -band magnitudes by ≈ 0.05 mag, favoring a slightly brighter Gaia G -band detection. There are a number of potential causes for this discrepancy. For example, we adopt a different reddening vector to produce the predicted G -band magnitudes because the Gaia passbands are not modeled in E. F. Schlafly et al. (2016). Moreover, the intrinsic Gaia colors in our current model grid are based on the Gaia DR2 photometric system, but we compare to magnitudes reported in the Gaia DR3 photometric system in Figure 18 (M. Riello et al. 2021). We defer a detailed investigation of this small systematic offset to future work.

Appendix E Minimum and Maximum Reliable Distance

To compute an estimate of the variation in the minimum and maximum reliable distance of the 3D dust map across the sky, we utilize the set of 2D binned posteriors on distance and reddening (Section 4.1.2) used in the line-of-sight fits (Section 4.2.3). The reliability of our inference is directly tied to how the stars are distributed along the line of sight toward a given patch of sky—if there is an insufficient number of stars constraining the line-of-sight fit at a given distance, we cannot adequately constrain the amount of reddening at that distance. To quantify this effect for each $N_{\text{side}} = 8192$ pixel, we sum the set of binned posteriors on distance and reddening (of shape $N_{\text{stars}} \times N_{E_{\text{bins}}} \times N_{\mu_{\text{bins}}}$) over the N_{stars} and $N_{E_{\text{bins}}}$ axes to obtain the probability distribution in distance modulus along that line of sight. We then compute the cumulative probability distribution in distance modulus stepping both forward (starting at $\mu = 4$ mag) and backward in distance (starting at $\mu = 19$ mag) to characterize the minimum and maximum reliable distance modulus, respectively.

We define the minimum and maximum reliable distance modulus as the first distance bin (computed forward in distance for the minimum and backward in distance for the maximum) to exceed some probability threshold. We tested a range of probability thresholds ranging from $p = 0.025$ (cumulative probability of a small fraction of a single star) to $p = 10$ (cumulative probability of 10 stars) over a range of test beds, including those shown in the gallery of artifacts in Figure 15. While the selection of a probability threshold is subjective, our goal was to select thresholds that minimized the known artifacts across the test beds, including striping, speckling, patchiness, and undersampling of the reddening in very dense regions. We find an optimal value of $p = 0.5$ for the minimum reliable distance modulus and $p = 5$ for the maximum reliable distance modulus. In many, but not all, lines of sight, we find that the $p = 0.5$ minimum cumulative probability threshold roughly corresponds to the distance of the first foreground star.

In Figure 19, we show the spatial variation in the minimum (top) and maximum (bottom) reliable distance across our footprint. In Figure 20, we show normalized histograms of the minimum reliable distance (left), maximum reliable distance (middle), and corresponding reddening at the maximum reliable distance (right, computed from the mean map) over the full footprint. We find a median and 1σ spread in the minimum reliable distance modulus of $\mu_{\text{min}} = 9.6^{+0.8}_{-0.9}$ ($d_{\text{min}} = 0.8^{+0.3}_{-0.3}$ kpc). We find a median and 1σ spread in the maximum reliable distance modulus of $\mu_{\text{max}} = 14.9^{+0.4}_{-0.8}$ ($d_{\text{max}} = 9.4^{+1.8}_{-2.8}$ kpc). Since the maximum reliable reddening—defined as the reddening of the map at the maximum reliable distance—is a function of both the sensitivity limit of the map and of natural variation in reddening over the plane (i.e., higher latitudes will have less reddening), we take 2σ above the median as the adopted maximum reliable extinction (spread spanning $+1\sigma$ to $+3\sigma$ above the median), finding $E_{B-V_{\text{max}}} = 3.7^{+1.5}_{-2.4}$ mag ($A_{V_{\text{max}}} = 12.3^{+5.0}_{-7.9}$ mag) for the mean 3D dust map (see Section 5.2).




Figure 19. Top: plane-of-sky variation in the minimum reliable distance, computed for each $N_{\text{side}} = 8192$ pixel as described in Appendix E. Bottom: plane-of-sky variation in the maximum reliable distance.




Figure 20. Normalized probability density functions of the distribution of the minimum reliable distance (left), maximum reliable distance (middle) and reddening at the maximum reliable distance (right), computed over the 51 million pixels in the DECaPS footprint.

ORCID iDs

Catherine Zucker <https://orcid.org/0000-0002-2250-730X>

Andrew K. Saydjari <https://orcid.org/0000-0002-6561-9002>

Joshua S. Speagle

(沈佳士) <https://orcid.org/0000-0003-2573-9832>

Edward F. Schlafly <https://orcid.org/0000-0002-3569-7421>

Gregory M. Green <https://orcid.org/0000-0001-5417-2260>

Robert Benjamin <https://orcid.org/0000-0002-8109-2642>

Joshua Peek <https://orcid.org/0000-0003-4797-7030>

Gordian Edenhofer <https://orcid.org/0000-0003-3122-4894>

Alyssa Goodman <https://orcid.org/0000-0003-1312-0477>

Michael A. Kuhn <https://orcid.org/0000-0002-0631-7514>

Douglas P. Finkbeiner <https://orcid.org/0000-0003-2808-275X>

References

- Almeida, A., Anderson, S. F., Argudo-Fernández, M., et al. 2023, *ApJS*, **267**, 44
- Alonso-García, J., Minniti, D., Catelan, M., et al. 2017, *ApJL*, **849**, L13
- Alonso-García, J., Saito, R. K., Hempel, M., et al. 2018, *A&A*, **619**, A4
- Anders, F., Khalatyan, A., Chiappini, C., et al. 2019, *A&A*, **628**, A94
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, **935**, 167
- Bailer-Jones, C. A. L., Rybizki, J., Foesneau, M., Mantelet, G., & Andrae, R. 2018, *AJ*, **156**, 58
- Bland-Hawthorn, J., & Gerhard, O. 2016, *ARA&A*, **54**, 529
- Boggess, N. W., Mather, J. C., Weiss, R., et al. 1992, *ApJ*, **397**, 420
- Burstein, D., & Heiles, C. 1978, *ApJ*, **225**, 40
- Burstein, D., & Heiles, C. 1982, *AJ*, **87**, 1165
- Cargile, P. A., Conroy, C., Johnson, B. D., et al. 2020, *ApJ*, **900**, 28
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
- Chen, B. Q., Huang, Y., Yuan, H. B., et al. 2019, *MNRAS*, **483**, 4277
- Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, **823**, 102
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, Explanatory Supplement to the AllWISE Data Release Products, Explanatory Supplement to the AllWISE Data Release Products, University of California
- Czekaj, M. A., Robin, A. C., Figueras, F., Luri, X., & Haywood, M. 2014, *A&A*, **564**, A102
- Dharmawardena, T. E., Bailer-Jones, C. A. L., Foesneau, M., et al. 2024, *MNRAS*, **532**, 3480
- Draine, B. T. 2003, *ARA&A*, **41**, 241
- Drew, J. E., Greimel, R., Irwin, M. J., et al. 2005, *MNRAS*, **362**, 753
- Drimmel, R., & Spergel, D. N. 2001, *ApJ*, **556**, 181
- Edenhofer, G., Zucker, C., Frank, P., et al. 2024, *A&A*, **685**, A82
- Ehlerová, S., & Palouš, J. 2013, *A&A*, **550**, A23
- Fabrizius, C., Luri, X., Arenou, F., et al. 2021, *A&A*, **649**, A5
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, **524**, 867

- Gaia Collaboration 2022, *yCat*, **1**, 355
- Gaia Collaboration, Prusti, T., de Bruijne, J., H., J., et al. 2016, *A&A*, **595**, A1
- Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, *A&A*, **674**, A1
- Gonzalez, O. A., Rejkuba, M., Zoccali, M., et al. 2012, *A&A*, **543**, A13
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Green, G. M. 2018, *JOSS*, **3**, 695
- Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *ApJ*, **887**, 93
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2014, *ApJ*, **783**, 114
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, *ApJ*, **810**, 25
- Green, G. M., Schlafly, E. F., Finkbeiner, D., et al. 2018, *MNRAS*, **478**, 651
- Green, G. M., Rix, H.-W., Tschesche, L., et al. 2021, *ApJ*, **907**, 57
- Griffin, M. J., Abergel, A., Abreu, A., et al. 2010, *A&A*, **518**, L3
- Han, J. L. 2017, *ARA&A*, **55**, 111
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Indebetouw, R., Mathis, J. S., Babler, B. L., et al. 2005, *ApJ*, **619**, 931
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Juric, M., 2012 LSD: Large Survey Database Framework, Astrophysics Source Code Library, ascl:1209.003
- Knollmüller, J., & Enßlin, T. A. 2019, arXiv:1901.11033
- Kroupa, P. 2001, *MNRAS*, **322**, 231
- Kuhn, M. A., Benjamin, R. A., Zucker, C., et al. 2021, *A&A*, **651**, L10
- Lallement, R., Babusiaux, C., Vergely, J. L., et al. 2019, *A&A*, **625**, A135
- Lallement, R., Capitanio, L., Ruiz-Dern, L., et al. 2018, *A&A*, **616**, A132
- Lallement, R., Vergely, J. L., Babusiaux, C., & Cox, N. L. J. 2022, *A&A*, **661**, A147
- Lang, D. 2014, *AJ*, **147**, 108
- Lee, J. C., Whitmore, B. C., Thilker, D. A., et al. 2022, *ApJS*, **258**, 10
- Leike, R. H., & Enßlin, T. A. 2019, *A&A*, **631**, A32
- Leike, R. H., Glatzle, M., & Enßlin, T. A. 2020, *A&A*, **639**, A138
- Lenz, D., Hensley, B. S., & Doré, O. 2017, *ApJ*, **846**, 38
- Lindgren, L., Bastian, U., Biermann, M., et al. 2021, *A&A*, **649**, A4
- Lindgren, L., Klioner, S. A., Hernández, J., et al. 2021, *A&A*, **649**, A2
- LSST 2024, *Survey Cadence Optimization Committee's Phase 3 Recommendations* PSTN-056, Vera C. Rubin Observatory <https://pstn-056.lsst.io/>
- Marshall, D., Montillaud, J., Cambresy, L., & Cornu, D. 2025, *A&A*, Submitted
- Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., & Picaud, S. 2006, *A&A*, **453**, 635
- Meisner, A. M., & Finkbeiner, D. P. 2015, *ApJ*, **798**, 88
- Minniti, D., Lucas, P. W., Emerson, J. P., et al. 2010, *NewA*, **15**, 433
- Molinari, S., Schisano, E., Elia, D., et al. 2016, *A&A*, **591**, A149
- NASA 2024, Wide Field Imager (WFI) Technical Information v1.3, NASA https://roman.gsfc.nasa.gov/science/WFI_technical.html
- Nataf, D. M., Gould, A., Fouqué, P., et al. 2013, *ApJ*, **769**, 88
- Neugebauer, G., Habing, H. J., van Duinen, R., et al. 1984, *ApJL*, **278**, L1
- Paladini, R., Zucker, C., Benjamin, R., et al. 2023, arXiv:2307.07642
- Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2014, *A&A*, **571**, A11
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, *A&A*, **594**, A10
- Reid, M. J., Menten, K. M., Brunthaler, A., et al. 2019, *ApJ*, **885**, 131
- Rezaei Kh., S., Beuther, H., Benjamin, R. A., et al. 2024, *A&A*, **692**, A255
- Riello, M., De Angeli, F., Evans, D. W., et al. 2021, *A&A*, **649**, A3
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, **409**, 523
- Robitaille, T., Beaumont, C., Qian, P., Borkin, M., & Goodman, A. 2019, glueviz v0.15.2: Multidimensional Data Exploration, doi:10.5281/zenodo.3385920
- Rybizki, J., Demleitner, M., Fouesneau, M., et al. 2018, *PASP*, **130**, 074101
- Saito, R. K., Hempel, M., Alonso-García, J., et al. 2024, *A&A*, **689**, A148
- Sale, S. E., Drew, J. E., Barentsen, G., et al. 2014, *MNRAS*, **443**, 2907
- Sanderson, R. E., Hickox, R., Hirata, C. M., et al. 2024, arXiv:2404.14342
- Saydjari, A. K., & Finkbeiner, D. P. 2022, *ApJ*, **933**, 155
- Saydjari, A. K., M. Uzsoy, A. S., Zucker, C., Peek, J. E. G., & Finkbeiner, D. P. 2023a, *ApJ*, **954**, 141
- Saydjari, A. K., Schlafly, E. F., Lang, D., et al. 2023b, *ApJS*, **264**, 28
- Schechter, P. L., Mateo, M., & Saha, A. 1993, *PASP*, **105**, 1342
- Schlafly, E. F., 2021 crowdsourc: Crowded Field Photometry Pipeline, Astrophysics Source Code Library, ascl:2106.004
- Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, **737**, 103
- Schlafly, E. F., Green, G. M., Lang, D., et al. 2018, *ApJS*, **234**, 39
- Schlafly, E. F., Meisner, A. M., & Green, G. M. 2019, *ApJS*, **240**, 30
- Schlafly, E. F., Meisner, A. M., Stutz, A. M., et al. 2016, *ApJ*, **821**, 78
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
- Sellgren, K. 1984, *ApJ*, **277**, 623
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163
- Speagle, J., Beane, A., & Zucker, C. 2025, joshspeagle/brutus: brutus v0.8.3, doi:10.5281/zenodo.14915000
- Speagle, J. S. 2019, arXiv:1909.12313
- Speagle, J. S., Zucker, C., Bonaca, A., et al. 2024, *ApJ*, **970**, 121
- Speagle, J. S., Zucker, C., Bonaca, A., et al. 2025, arXiv:2503.02227
- Stetson, P. B. 1987, *PASP*, **99**, 191
- Surot, F., Valenti, E., Gonzalez, O. A., et al. 2020, *A&A*, **644**, A140
- Tchernyshyov, K., Peek, J. E. G., & Zasowski, G. 2018, *AJ*, **156**, 248
- Vergely, J. L., Lallement, R., & Cox, N. L. J. 2022, *A&A*, **664**, A174
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Wang, S., & Chen, X. 2019, *ApJ*, **877**, 116
- Weingartner, J. C., & Draine, B. T. 2001, *ApJS*, **134**, 263
- Xue, X.-X., Rix, H.-W., Ma, Z., et al. 2015, *ApJ*, **809**, 144
- Zhang, M., & Kainulainen, J. 2019, *A&A*, **632**, A85
- Zhang, M., & Kainulainen, J. 2022, *MNRAS*, **517**, 5180
- Zhang, X., & Green, G. 2025, *Sci*, **387**, 1209
- Zhang, X., Green, G. M., & Rix, H.-W. 2023, *MNRAS*, **524**, 1855
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298
- Zucker, C., Goodman, A., Alves, J., et al. 2021, *ApJ*, **919**, 35
- Zucker, C., Speagle, J. S., Schlafly, E. F., et al. 2019, *ApJ*, **879**, 125