



Mapping the Milky Way in 5D with 170 Million Stars

Joshua S. Speagle (沈佳士)^{1,2,3,4}, Catherine Zucker^{4,5}, Ana Bonaca^{4,6}, Phillip A. Cargile⁴, Benjamin D. Johnson⁴, Angus Beane⁴, Charlie Conroy⁴, Douglas P. Finkbeiner^{4,7}, Gregory M. Green⁸, Harshil M. Kamdar⁴, Rohan Naidu⁴, Hans-Walter Rix⁸, Edward F. Schlafly⁹, Aaron Dotter⁴, Gwendolyn Eadie^{1,2}, Daniel J. Eisenstein⁴, Alyssa A. Goodman^{4,10}, Jiwon Jesse Han⁴, Andrew K. Saydjari^{4,7}, Yuan-Sen Ting (丁源森)^{11,12,13,14,15}, and Ioana A. Zelko⁴

¹ Department of Statistical Sciences, University of Toronto, Toronto, ON M5S 3G3, Canada; j.speagle@utoronto.ca

² David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada

³ Dunlap Institute for Astronomy & Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada

⁴ Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA

⁵ Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218, USA

⁶ Carnegie Observatories, 813 Santa Barbara St., Pasadena, CA 91101, USA

⁷ Department of Physics, Harvard University, 17 Oxford St., Cambridge, MA 02138, USA

⁸ Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

⁹ Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, CA 94550, USA

¹⁰ Radcliffe Institute for Advanced Study, Harvard University, 10 Garden St., Cambridge, MA 02138, USA

¹¹ Institute for Advanced Study, Princeton, NJ 08540, USA

¹² Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

¹³ Observatories of the Carnegie Institution of Washington, 813 Santa Barbara St., Pasadena, CA 91101, USA

¹⁴ Research School of Astronomy and Astrophysics, Australian National University, Cotter Road, ACT 2611, Canberra, Australia

¹⁵ Research School of Computer Science, Australian National University, Acton, ACT 2601, Australia

Received 2021 May 12; revised 2022 February 15; accepted 2022 February 15; published 2024 July 24

Abstract

We present *Augustus*, a catalog of distance, extinction, and stellar parameter estimates for 170 million stars from $14 \text{ mag} < r < 20 \text{ mag}$ and with $|b| > 10^\circ$ drawing on a combination of optical to near-infrared photometry from Pan-STARRS, 2MASS, UKIDSS, and unWISE along with parallax measurements from Gaia DR2 and 3D dust extinction maps. After applying quality cuts, we find 125 million objects have “high-quality” posteriors with statistical distance uncertainties of $\lesssim 10\%$ for objects with well-constrained stellar types. This is a substantial improvement over the distance estimates derived from Gaia parallaxes alone and in line with the recent results from Anders et al. We find the fits are able to reproduce the dereddened Gaia color–magnitude diagram accurately, which serves as a useful consistency check of our results. We show that we are able to detect large, kinematically coherent substructures in our data clearly relative to the input priors, including the Monoceros Ring and the Sagittarius Stream, attesting to the quality of the catalog. Our results are publicly available at doi:[10.7910/DVN/WYMSXV](https://doi.org/10.7910/DVN/WYMSXV). An accompanying interactive visualization can be found at <http://allsky.s3-website.us-east-2.amazonaws.com>.

Unified Astronomy Thesaurus concepts: Stellar distance (1595); the Milky Way (1054); Sky surveys (1464); Photometry (1234); Parallax (1197)

Materials only available in the *online version of record*: interactive figures

1. Introduction

A central challenge in astronomy is converting the projected 2D positions of sources on the sky into 3D maps that we can use to infer properties about the Universe. This is especially true when studying the Milky Way, where recent observational advances have opened possibilities for 3D mapping across our Galaxy. But many new discoveries depend on the fidelity of such 3D mapping. Recent works have exploited full phase-space data to uncover the remnants of a major merger ~ 10 Gyr ago (e.g., Belokurov et al. 2018; Helmi et al. 2018; Koppelman et al. 2018; Naidu et al. 2021) and a phase-space “spiral” (e.g., Antoja et al. 2018), while in the halo accurate phase-space maps of stellar streams have begun to constrain the potential of the Galaxy (e.g., Johnston et al. 1999; Law & Majewski 2010; Bonaca & Hogg 2018).

These discoveries have benefited from simultaneous advances across multiple fronts. On the data side, large missions such as the ground-based Sloan Digital Sky Survey (SDSS; York et al. 2000) and the space-based Gaia mission (Gaia Collaboration et al. 2016) have published enormous public data sets. Together, these observational efforts promise to provide new, much sharper maps of the stellar components of the Galaxy using billions of individual sources. Simultaneously, advances in statistical modeling and computational power have enabled us to infer the 3D distribution of a large number of stars more robustly (e.g., Green et al. 2014; Bailer-Jones et al. 2018) along with additional properties such as ages and abundances (e.g., Ness et al. 2015; Anders et al. 2019; Leung & Bovy 2019a; Xiang et al. 2019). Finally, advances in numerical simulations and Galactic dynamics have enabled us to interpret these results in much more detail (see Rix & Bovy 2013; Sellwood 2014; Helmi 2020, and references therein).

As most sources (~99%) seen in large photometric surveys do not have measured spectra, much of the work associated with deriving 3D maps to billions of stars has relied on modeling coarse spectral energy distributions (SEDs)



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

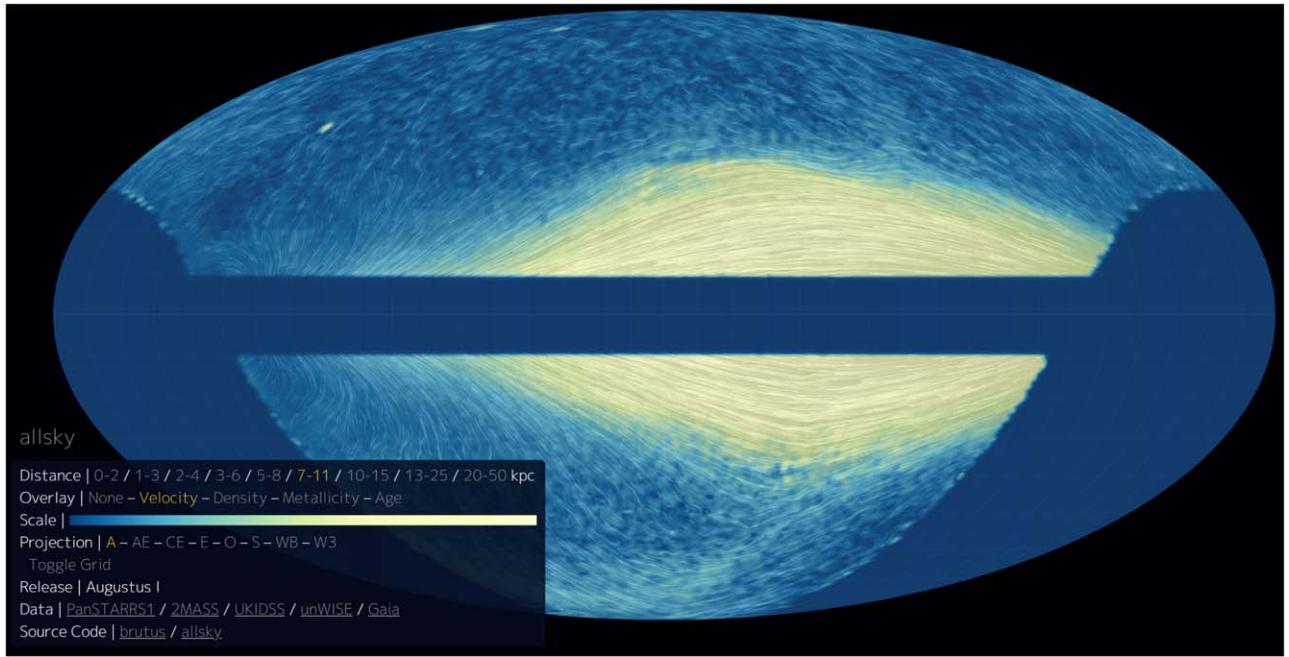


Figure 1. A screenshot of an interactive visualization of 3D distance and 2D velocity structure from the Augustus catalog that can be accessed online at <http://allsky.s3-website.us-east-2.amazonaws.com>. The background color is based on the overall tangential speed of the sources in a given distance bin, while the white streamlines follow the tangential velocities of the same sources in the given coordinate projection. The interface to change these properties (distance bin, background properties, and projection) is shown in the bottom left and can be opened/minimized by clicking on the “allsky” text.

comprised of flux densities estimated across a range of broadband and narrowband photometric filters (e.g., Anders et al. 2019; Green et al. 2019). More recently, Gaia DR2 (Gaia Collaboration et al. 2018) has also provided astrometric parallax measurements for many of these sources, giving independent constraints on the distances.

Efforts in this area range from 3D dust mapping (e.g., Rezaei Kh. et al. 2018; Green et al. 2019; Lallement et al. 2019; Leike & Enßlin 2019) to stellar parameter estimation (e.g., Ness et al. 2015; Anders et al. 2019; Cargile et al. 2020). In Speagle et al. (2024), we described new methods implemented in the public, open-source PYTHON package BRUTUS¹⁶ that further contribute to these efforts by allowing for quick and robust estimation of stellar properties, distances, and reddenings to stars with photometric and astrometric data. In this paper, we present Augustus, a “proof-of-concept” application of BRUTUS to estimate distances, reddenings, and various stellar properties to 170 million sources brighter than $r < 20$ mag with Galactic latitudes $|b| > 10^\circ$.

The outline of the paper is as follows. In Section 2, we describe the data sets, quality cuts, and selections used to select the 170 million objects in this work. In Section 3, we summarize the approach taken to model and fit the 170 million sources described with BRUTUS. In Section 4, we describe the catalogs produced by this modeling. In Section 5, we discuss results demonstrating the quality of the output data, including “blind” recovery of the Gaia color–magnitude diagram (CMD) and the detection of known large-scale Galactic substructure. We conclude in Section 6. A detailed description of the data products provided as part of this work can be found in the Appendix. An interactive figure highlighting many of the features of our output catalog can be found at <http://allsky.s3-website.us-east-2.amazonaws.com>.

[website.us-east-2.amazonaws.com](http://allsky.s3-website.us-east-2.amazonaws.com); a screenshot is shown in Figure 1 and described in more detail in Section 5.5.

Throughout the paper, individual parameters (scalars) are notated using standard italicized math fonts (θ) while vectors and matrices are notated using boldface (θ). Collections of parameters are notated using sets ($\theta = \{\theta_i\}_{i=1}^n$). Vectors should be assumed to be in column form (i.e., of shape $n \times 1$) unless explicitly stated otherwise. All magnitudes reported are in the native units provided by their corresponding data sets.

2. Data

Our analysis is based on the combination of several surveys:

1. the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS; Chambers et al. 2016),
2. the Two Micron All Sky Survey (2MASS; Skrutskie et al. 2006),
3. the United Kingdom Infrared Telescope Infrared Deep Sky Survey (UKIDSS; Lawrence et al. 2007),
4. the “unofficial” Wide-field Infrared Survey Explorer (unWISE) catalog (Wright et al. 2010; Schlafly et al. 2019), and
5. the Gaia survey (Gaia Collaboration et al. 2016).

In Sections 2.1–2.5, we describe each of the various data sets. In Section 2.6, we describe how the data sets are combined into a final catalog. The overall sky and wavelength coverage from these combination of surveys are shown in Figures 2 and 3, respectively.

2.1. Pan-STARRS

The Pan-STARRS survey¹⁷ is a multiepoch, deep, broadband optical survey of the northern sky visible from Haleakala

¹⁶ <https://github.com/joshspeagle/brutus>

¹⁷ <https://panstarrs.stsci.edu/>

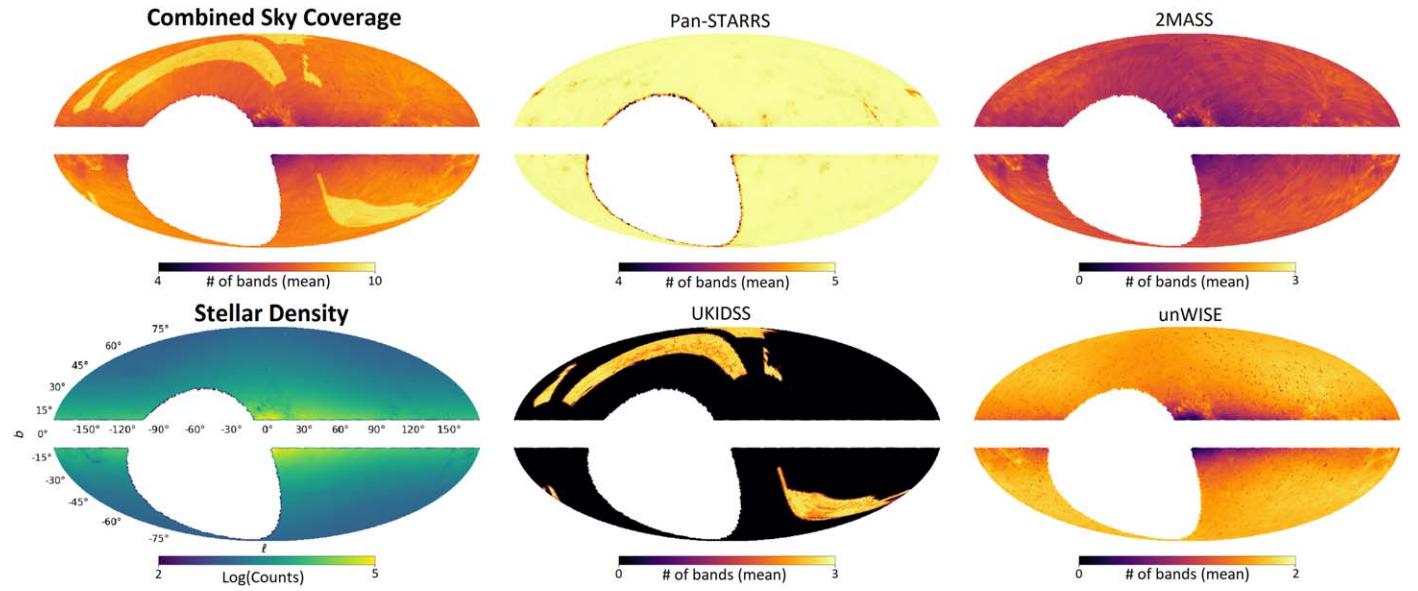


Figure 2. Projected maps showing the coverage of the data used in the work at a `healpix` resolution of `nside` = 64 plotted as a function of Galactic longitude ℓ and latitude b , centered on $(\ell, b) = (0^\circ, 0^\circ)$. The top left panel shows the mean number of bands from all surveys combined, while the bottom left panel shows the overall (log-)number of stars in each pixel. The mean number of bands from the individual surveys are shown in the remaining panels, with Pan-STARRS (Section 2.1) in the top middle panel, 2MASS (Section 2.2) in the top right panel, UKIDSS (Section 2.3) in the bottom middle panel, and unWISE (Section 2.4) in the bottom right panel. While toward the Galactic center we lose a substantial amount of coverage (transition from orange/yellow to purple in the upper left panel) due to crowding and dust extinction (transition from blue to yellow in the lower left panel), at high Galactic latitudes and in the Galactic outskirts we have 8–10 bands of optical to near-infrared (NIR) coverage. The uniform coverage of the Pan-STARRS data is due to the sample selection, which requires ≥ 4 bands of Pan-STARRS photometry. Due to the depth of the UKIDSS Large Area Survey (LAS) data, in regions that overlap with the survey area we have almost 10 bands of wavelength coverage from the optical through the IR. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_sky_coverage.html. An interactive version of Figures 2 and 3 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting the wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top right provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

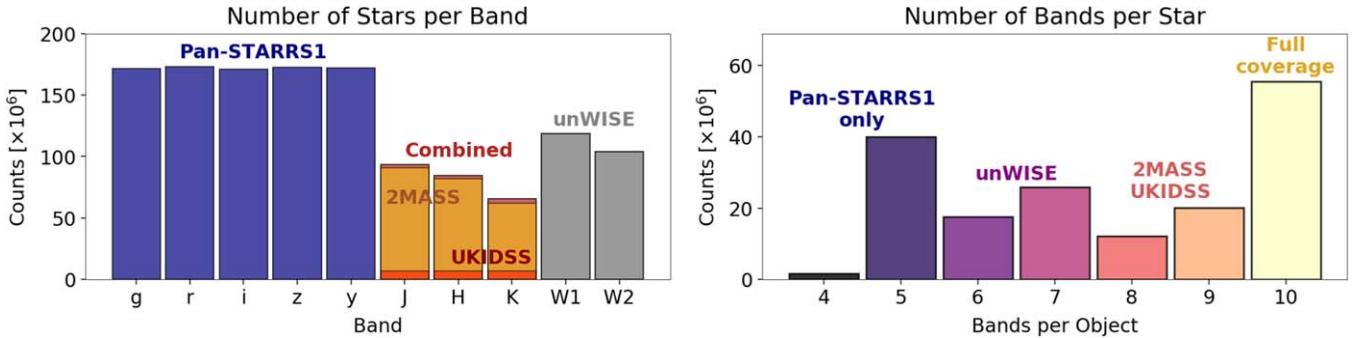


Figure 3. Histograms illustrating the wavelength coverage of the data used in this work that serves to complement Figure 2. The left panel shows the number of stars available in each band, color coded by survey, illustrating the uniform Pan-STARRS selection and amount of NIR and IR coverage available through 2MASS and unWISE, respectively. Compared with Figure 2, we can see that although the UKIDSS data are quite deep, the total amount of objects with UKIDSS photometry is quite small due to its high-latitude targeting. We also see that a substantial amount of objects are detected in the unWISE catalog. The right panel shows the number of bands per star, highlighting that the majority of the sample has full coverage across all 10 possible bands. Regions where we roughly lose coverage across various surveys are labeled, with the next largest peak of objects having only ~ 5 bands of coverage (mostly in the optical) closer to the Galactic plane. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_sky_coverage.html. An interactive version of Figures 2 and 3 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top right provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

in Hawaii (i.e., $\delta > -30^\circ$). It observed in five photometric bands (*grizy*) spanning 0.4–1 μm with typical single-epoch 5σ point-source exposure depths of $g = 22.0$ mag, $r = 21.8$ mag, $i = 21.5$ mag, $z = 20.9$ mag, and $y = 19.7$ mag in the AB system (Oke & Gunn 1983).

The photometry used in this work is based on combined single-epoch photometry obtained as part of the Pan-STARRS1 3π

Steradian Survey DR1 (Chambers et al. 2016). Photometry and astrometry were derived from combined images as described in Magnier et al. (2020). We remove galaxies by requiring that the difference between the point-spread function model photometry and aperture photometry is < 0.1 mag across at least four bands. We remove objects below the Pan-STARRS saturation limit in each band, and also objects with magnitude errors > 0.2 mag.

2.2. 2MASS

2MASS¹⁸ is a NIR survey of the entire sky in three photometric bands (JHK_s) spanning 1–2.3 μm with typical 10σ point-source exposure depths of $J = 15.8$ mag, $H = 15.1$ mag, and $K_s = 14.3$ mag in the Vega system.¹⁹ We utilize data from the 2MASS “high-reliability” catalog,²⁰ which minimizes contamination and confusion by neighboring point and/or extended sources. We also require errors to be <0.2 mag and that no photometry quality (ph_qual), read quality (rd_qual), or galaxy contamination (gal_contam) flags are set.

2.3. UKIDSS

The UKIDSS project²¹ is defined in Lawrence et al. (2007). UKIDSS uses the UKIRT Wide Field Camera (Casali et al. 2007). The photometric system is described in Hewett et al. (2006), and the calibration is described in Hodgkin et al. (2009). The science archive is described in Hambly et al. (2008).

We use data from the UKIDSS LAS second data release (DR2; Dye et al. 2006; Warren et al. 2007a, 2007b). UKIDSS LAS imaged 4000 degree² in three fields (an equatorial block, a northern block, and a southern stripe) that were a subset of the SDSS footprint (York et al. 2000) in four photometric bands ($YJHK$) spanning 1–2 μm with typical 5σ point-source exposure depths of $Y = 20.5$ mag, $J = 20.0$ mag, $H = 18.8$ mag, and $K = 18.4$ mag in the Vega system. We require the errors to be <0.2 mag, the probability of being a star $\text{pstar} > 0.9$, and that no processing/warning errorbit flags ($[\text{J}_1/\text{h}/\text{k}] \text{ppErrBits}$) are set.

2.4. unWISE

The unWISE catalog²² (Schlafly et al. 2019) is a collection of two billion sources observed over the entire sky in the IR over two bands (W_1 and W_2) from 3 to 5 μm with the WISE satellite as part of the WISE and Near Earth Object WISE (NEOWISE) missions. Compared to the existing AllWISE catalog (Cutri et al. 2013), the unWISE catalog is 0.7 mag deeper due to its use of additional images from the extended mission along with improved source modeling in crowded regions using the CROWDSOURCE code (Schlafly et al. 2018). This allows it to reach 50% completeness point-source depths of $W_1 = 17.9$ mag and $W_2 = 16.7$ mag in the Vega system.

We select all objects which are flagged as primary (no duplicate sources), contain no bitwise quality flags (flag_s_unwise), and that have errors of <0.2 mag. We also further constrain the fractional flux (fracflux) at an object’s given position, which measures the fraction of contamination of light from the source due to neighboring objects, to be >0.5 (i.e., a majority of the light at a given position is contributed by the source being modeled).

¹⁸ <https://old.ipac.caltech.edu/2mass/>

¹⁹ While references to the “Vega system” suggest a single alternate system, it instead represents a variety of independent photometric calibrations to differing models of Vega. This introduces additional systematics when attempting to combine photometry across different surveys, which will be discussed later.

²⁰ The cuts used in this selection are described at <https://old.ipac.caltech.edu/2mass/releases/allsky/doc/sec22.html>.

²¹ <http://www.ukidss.org/surveys/surveys.html>

²² <https://catalog.unwise.me/>

2.5. Gaia

We use data taken from the Gaia second data release (DR2; Gaia Collaboration et al. 2018), which provides photometry (G , BP , and RP) and astrometry (proper motion and parallax) measurements for over one billion stars, along with radial velocity measurements for a small fraction of nearby sources. The astrometric catalog has a 99.875% completeness point-source depth of $G \approx 21$ mag in the Vega system (Gaia Collaboration et al. 2018; Lindegren et al. 2018). The typical astrometric uncertainty is around 0.7 mas for the faintest stars and 0.04 mas at the bright limit.

In this work, we only incorporate Gaia parallax measurements and their uncertainties when modeling individual sources. We only use the photometry and proper motions as additional checks to validate our results, as will be discussed in Section 5. We impose the same quality cuts as recommended in Equation (11) of Lindegren et al. (2018), which requires sources to have:

1. $G \leq 21$,
2. $\text{visibility_periods_used} \geq 6$, and
3. $\text{astrometric_sigma5d_max} \leq 1.2 \text{ mas} \times \gamma(G)$,

where $\gamma(G) = \max[1, 10^{0.2(G-18)}]$. We do not impose or record astrometric quality information from any other quantities (e.g., RUWE).

As discussed in Lindegren et al. (2018), there are numerous systematics present in the parallax measurements provided as part of Gaia DR2. The ones we consider here are overall zero-point offsets in the parallax measurements as well as possible underestimates of the provided errors. Following work by Schönrich et al. (2019), Leung & Bovy (2019b), and Khan et al. (2019), among others, we add 0.054 mas to all parallaxes and increase the measurement errors by adding 0.043 mas in quadrature with the reported uncertainties.

2.6. Assembling the Augustus Catalog

We crossmatched all sources in Pan-STARRS, 2MASS, UKIDSS, unWISE, and Gaia DR2 after applying the cuts described above within a radius of 1'', with the closest source being selected in the presence of multiple matches. This operation was performed using the Large Survey Database (Juric 2011) hosted on the Cannon computing cluster at Harvard University and required Pan-STARRS to be the “primary” catalog to which all others were matched. We impose a minimum error threshold of 0.005 mag in all bands after removing any additional survey-imposed error floors; we also add in our own (larger) error floors as described in Section 3. In addition, we purposely mask out 2MASS photometry whenever UKIDSS photometry in the same band is available. This helps to avoid adding additional noise from the 2MASS observations, which are substantially shallower than the UKIDSS observations in the same footprint, and also prevents us “double-counting” systematics.²³

In addition to the cuts described above, we imposed four additional cuts:

1. ≥ 4 bands of photometry in Pan-STARRS,
2. a parallax measurement in Gaia DR2,

²³ If we estimate uncertainties to be dominated by a systematic error Δ , observations from $n \sim$ identical bands with errors floors of Δ in each effectively makes the error floor Δ/\sqrt{n} .

3. $r < 20$ mag in Pan-STARRS, and
4. a Galactic latitude of $|b| > 10^\circ$.

The first cut guarantees that we have approximately uniform coverage in Pan-STARRS and guarantees we have enough photometry (≥ 4 bands) to be able to run BRUTUS, and ensures our spatial coverage matches that of the 3D dust prior from Green et al. (2019). The second requirement guarantees that all sources will have parallax measurements (and Gaia photometry), making it straightforward to compare to previous work such as Anders et al. (2019). The third requirement limits the size of the sample and helps us remain within the Gaia $G \lesssim 21$ mag completeness limit, making our sample roughly r -band magnitude limited. Finally, the fourth cut helps to limit the size of the sample further while avoiding the intense crowding/dust extinction near the Galactic plane.

The final combined catalog, which we refer to as *Augustus*, has roughly 170 million sources. The distribution of sources on the sky is shown in Figure 2, while the wavelength coverage for stars in our sample is shown in Figure 3. The lack of data in the south is due to the required coverage in Pan-STARRS, which only allows for 3π steradians of coverage. In general, we find a plurality of sources (~ 60 million) have 10 bands of optical-through-IR photometry across Pan-STARRS, 2MASS/UKIDSS, and unWISE, and that only 40 million ($\sim 25\%$) have ≤ 5 bands of coverage.

3. Modeling

The modeling approach for constructing this catalog, BRUTUS, is described in Speagle et al. (2024). We provide a brief summary below.

3.1. Statistical Framework

BRUTUS uses Bayesian inference to model the posterior probability $\mathcal{P}(\theta, \phi)$ of a set of intrinsic stellar parameters θ (e.g., initial mass M_{init}) and extrinsic stellar parameters (e.g., distance d) as the product of three different components:

$$\mathcal{P}(\theta, \phi) \propto \mathcal{L}_{\text{phot}}(\theta, \phi) \mathcal{L}_{\text{astr}}(\phi) \pi(\theta, \phi). \quad (1)$$

The first term is the photometric likelihood between our model of the flux densities $\mathbf{F}(\theta, \phi) \equiv \{F_i(\theta, \phi)\}_{i=1}^{i=b}$ and the observed flux densities $\hat{\mathbf{F}} \equiv \{\hat{F}_i\}_{i=1}^{i=b}$ and their associated errors $\hat{\sigma} \equiv \{\hat{\sigma}_i\}_{i=1}^{i=b}$ across b bands. We assume the data follows a normal distribution in each band such that:

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

The second term is the astrometric likelihood, which compares the predicted parallax $\varpi(\phi)$ to the observed value $\hat{\varpi}$ and the associated error $\hat{\sigma}_\varpi$.²⁴ We also assume the data follow a normal distribution such that:

$$\mathcal{L}_{\text{astr}}(\phi) \equiv \frac{1}{\sqrt{2\pi\hat{\sigma}_\varpi^2}} \exp \left[-\frac{1}{2} \frac{(\varpi(\phi) - \hat{\varpi})^2}{\hat{\sigma}_\varpi^2} \right]. \quad (3)$$

The final term is the Galactic prior, which describes our prior belief over the 3D distribution of stars, dust, and their associated properties. We use the same default prior outlined in Speagle et al. (2024). This includes a thin disk, thick disk,

and halo component whose size/shape are based on Bland-Hawthorn & Gerhard (2016) and Xue et al. (2015) with simple, spatially independent distributions of initial metallicities $[\text{Fe}/\text{H}]_{\text{init}}$ and ages t_{age} as described in Speagle et al. (2024).²⁵ The 3D distribution of dust attenuation A_V is taken to follow the 3D dust map from Green et al. (2019), with variations in the dust curve (as parameterized by R_V) taken to be spatially independent following that from Schlafly et al. (2016).

3.2. Stellar Modeling

The stellar models used in this work are the Modules for Experiments in Stellar Astrophysics (Paxton et al. 2011, 2013, 2015, 2018, 2019) Isochrone & Stellar Tracks (MIST; Choi et al. 2016). In particular, we utilize the nonrotating models from MIST v1.2. These are defined in terms of initial mass M_{init} , initial metallicity $[\text{Fe}/\text{H}]_{\text{init}}$, and equivalent evolutionary point (EEP; Dotter 2016), which correspond to a unique age $t_{\text{age}}(\text{EEP}|M_{\text{init}}, [\text{Fe}/\text{H}]_{\text{init}})$ for a given M_{init} and $[\text{Fe}/\text{H}]_{\text{init}}$.

We use three separate pieces to predict the underlying stellar spectrum $F_\nu(\lambda|\theta, \phi)$. The first is the C3K stellar atmosphere models described in Cargile et al. (2020) and Speagle et al. (2024), which are computed as a function of effective temperature T_{eff} , log surface gravity $\log g$, surface metallicity $[\text{Fe}/\text{H}]_{\text{surf}}$, and surface alpha abundance variation $[\alpha/\text{Fe}]_{\text{surf}}$. While v1.2 of the MIST models provide predictions for T_{eff} , $\log g$, and $[\text{Fe}/\text{H}]_{\text{surf}}$, they do not model any variations in $[\alpha/\text{Fe}]$; as a result, we set $[\alpha/\text{Fe}]_{\text{surf}} = 0$ by default. The second is a set of “empirical corrections” to the MIST models based on isochrone modeling of nearby open clusters described in Speagle et al. (2024). These are implemented as adjustments to the effective temperature T_{eff} and stellar radius $\log R_\star$ (and by extension the surface gravity $\log g$ and bolometric luminosity $\log L_{\text{bol}}$) as a function of M_{init} , with small modifications as a function of EEP and $[\text{Fe}/\text{H}]_{\text{init}}$. The third is a set of dust extinction curves (i.e., reddening laws) from Fitzpatrick (2004) to account for the effect of dust extinction. These are defined as a function of extinction in the V -band A_V and the “differential extinction” $R_V \equiv A_V/(A_B - A_V)$ based on the ratio of A_V to the difference in extinction $A_B - A_V$ between the B and V bands. Altogether, these give a framework for generating spectra as a function of $\theta = \{M_{\text{init}}, [\text{Fe}/\text{H}]_{\text{init}}, \text{EEP}\}$ and $\phi = \{d, A_V, R_V\}$ with $\theta_* = \{t_{\text{age}}, T_{\text{eff}}, \log g, \log R_\star, \log L_{\text{bol}}\}$ generated as intermediate values, where d is again the distance to the source.

Finally, to generate the model flux density in a given filter i , we integrate the underlying stellar spectrum $F_\nu(\lambda|\theta, \phi)$ through a filter curve with transmission $T_i(\lambda)$:

$$F_i(\theta, \phi) = \frac{\int F_\nu(\lambda|\theta, \phi) T_i(\lambda) \lambda^{-1} d\lambda}{\int S_\nu(\lambda|\theta, \phi) T_i(\lambda) \lambda^{-1} d\lambda}, \quad (4)$$

where $S_\nu(\lambda)$ is the source spectrum used to calibrate the observations. This is the chosen spectrum of Vega in the Vega system and a constant in the AB system. To avoid having to compute integrals “on the fly,” we precompute photometry over a large grid of T_{eff} , $[\text{Fe}/\text{H}]_{\text{surf}}$, $\log g$, $[\alpha/\text{Fe}]$, A_V , and R_V values in each band, and interpolate over the resulting photometric predictions in each band using a neural network

²⁴ Note that while other astrometric measurements such as proper motions are measured by Gaia, we do not model them here.

²⁵ No bulge or bar component is currently included but will be added in future work.

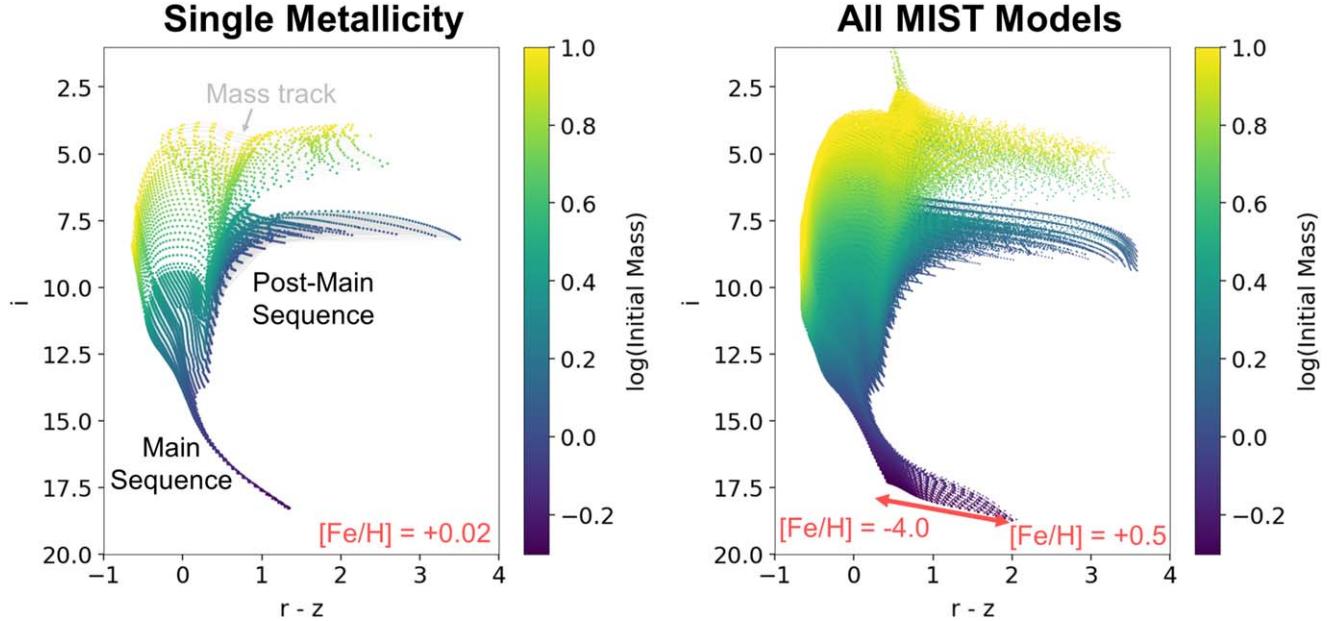


Figure 4. CMDs for the MIST models used in this work as a function of Pan-STARRS i -band magnitude at $d = 1$ kpc and Pan-STARRS $r - z$ color. The left panel shows the models used with $[\text{Fe/H}]_{\text{init}} = +0.02$, with the rough locations of the main-sequence and post-main-sequence evolutionary phases indicated. The underlying mass tracks are shown as light gray lines, with the actual models used shown as points and colored by $\log M_{\text{init}}$. The right panel shows the entire collection $\sim 750,000$ models defined over $[\text{Fe/H}]_{\text{init}} = -4.0$ to $+0.5$ that are used in this work.

as described in Speagle et al. (2024). Cross-validation and hold-out testing suggest that the difference between photometry predicted using the neural network versus direct integration is $\lesssim 0.01$ mag over a large majority of the parameter space.

3.3. Application to Data

BRUTUS exploits the nature of the statistical problem to derive continuous estimates of the extrinsic stellar parameters $\phi = \{d, A_V, R_V\}$ over a grid of intrinsic stellar parameters $\theta = \{M_{\text{init}}, [\text{Fe/H}]_{\text{init}}, \text{EEP}\}$. While interpolating over an input grid of stellar models allows for smoother probabilistic estimation of the underlying parameters (Cargile et al. 2020), this process in general is substantially slower than using precomputed grids when the number of parameters being inferred is small ($\lesssim 4$). BRUTUS uses grids to exploit this speedup. This leads to “gridding effects” that will be visible in subsequent plots shown in this work.

In brief, BRUTUS fits each object in three steps.

1. *Magnitude step.* Compute a “quick approximation” of the solution in magnitudes.
2. *Flux density step.* Improve the magnitude solution after converting the data back to a flux density.
3. *Prior step.* Incorporate information from the prior (and the parallax) using Monte Carlo sampling.

These steps are then parallelized across all models in the grid, with the final posterior estimated using Monte Carlo integration and resampling. The entire process takes only a few seconds for a typical source with a mildly informative parallax measurement. The grid of stellar models we use in this work (grid_mist_v8) are defined in Table 2, shown in Figure 4, and available online through the BRUTUS GitHub page.

As discussed in Speagle et al. (2024) and elsewhere (e.g., Choi et al. 2016), there are known systematics offsets between

the MIST models used in this work and the photometric data to which it is being fit. To account for some of this additional uncertainty not captured by the empirical corrections described earlier, we apply zeroth-order photometric offsets to the observed flux densities and increase the effective errors by adding in error floors in quadrature. A summary of the offsets and error floors applied in this work can be found in Table 1.

We run BRUTUS using this setup over the 170 million sources in the Augustus catalog using the default hyperparameters enabled in v0.7.5 (Speagle & Zucker 2020). All computations were performed on the Cannon research computing cluster at Harvard University. Including overheads, the final runtime was $\sim 700,000$ CPU hours.

4. Catalogs

Using the posterior samples for each object, we postprocess the data into two output catalogs:

1. a “point” catalog²⁶ containing various information about each object and summary statistics describing the stellar parameters, and
2. a “samples” catalog²⁷ containing a subset of 25 posterior samples for each object.

We will use results from the former catalog when highlighting results in this paper; the latter is meant to be used as a supplement for users interested in additional error modeling. Detailed descriptions of both catalogs and examples of their usage can be found in the Appendix.

Note that after performing most of the computation, we discovered that BRUTUS v0.7.5 contained a bug (fixed in more recent versions of the code) that used the wrong sign when sampling from correlations between A_V and R_V with distance. We have confirmed this has a negligible impact on the overall

²⁶ doi:10.7910/DVN/WYMSXV.

²⁷ doi:10.7910/DVN/530UYQ.

Table 1
A Summary of the Photometric Offsets

Filter	Offset	Error Floor
Pan-STARRS		
<i>g</i>	1.01	0.02
<i>r</i>	0.97	0.02
<i>i</i>	0.97	0.02
<i>z</i>	0.96	0.02
<i>y</i>	0.97	0.02
2MASS		
<i>J</i>	0.99	0.03
<i>H</i>	1.04	0.03
<i>K_s</i>	1.04	0.03
UKIDSS		
<i>J</i>	0.99	0.03
<i>H</i>	1.04	0.03
<i>K</i>	1.04	0.03
unWISE		
W1	1.02	0.04
W2	1.03	0.04

Note. Photometric offsets that are multiplied to the observed flux densities and the error floors (as a fraction of the flux density) that are added in quadrature to the observational uncertainties.

posterior distributions and marginal distributions for all parameters and therefore should not impact the quality of the catalog; however, it does affect quantities computed directly from the samples which depend on these quantities (e.g., reddened photometry).

In addition, we found that survey artifacts from the UKIDSS footprint (deeper NIR photometry changed the distribution of stellar parameter estimated) were prominent when projecting results onto the plane of the sky or Galaxy. As a result, we reran all objects in the UKIDSS LAS footprint without UKIDSS photometry (i.e., using 2MASS instead when available) and with the same version of the code (v0.7.5) for consistency; the catalogs for this subset of objects are also provided online and described in the [Appendix](#). While this does degrade the quality of the stellar parameter estimates, it makes the resulting maps more homogeneous. As a result, we opt to use the “no UKIDSS” versions of our results when highlighting results in this work unless otherwise explicitly stated.

The distribution of a random posterior sample of the distances for each star in *Augustus* is shown in Figure 5. We see that the distribution peaks around a few kiloparsecs, with a sharp decline toward larger distances and a shallower decline toward smaller ones. The former behavior can be understood by our $r = 20$ mag faint magnitude limit in Pan-STARRS, which makes us primarily sensitive to giants at larger distances. The latter behavior is due to a combination of two effects. The first is the increasing differential volume which goes as d^2 , which increases the raw number of sources available between d and $d + \Delta d$ at larger distances. This counteracts the decreasing number density of stars as we move away from the Galactic center and out of the Galactic plane. The second is the $r \sim 14$ mag saturation limit in Pan-STARRS, which makes us increasingly incomplete at nearby distances.

As part of the catalog we generate two quality flags:

Table 2
Grid of Parameters for the MIST Models Used in This Work

Minimum	Maximum	Spacing
Initial Mass (M_{init})		
$0.5 M_{\odot}$	$2.8 M_{\odot}$	$0.02 M_{\odot}$
$2.8 M_{\odot}$	$3.0 M_{\odot}$	$0.1 M_{\odot}$
$3.0 M_{\odot}$	$8.0 M_{\odot}$	$0.25 M_{\odot}$
$8.0 M_{\odot}$	$10.0 M_{\odot}$	$0.5 M_{\odot}$
Initial Metallicity ([Fe/H] _{init})		
-4.0	+0.5	0.06
Equivalent Evolutionary Point		
202	454	12
454	808	6

Note. See Section 3.2 for additional details.

1. FLAG_FIT, which diagnoses problems in the best-fit model SED, including the predicted parallax, and
2. FLAG_GRID, which diagnoses when the output posterior appears to be artificially truncated by the input model parameter grid.

The details of these flags are discussed in the [Appendix](#). There are roughly 14 million sources (~8%) which have FLAG_FIT = TRUE set, 38 million (~22%) with FLAG_GRID = TRUE set, and 47 million (~27%) with either flag set. We consider the set of roughly 125 million sources with neither flag set to have “reliable,” high-quality posteriors that are sufficient for analysis. We utilize this subset of stars in *Augustus* in all subsequent analyses and will henceforth refer to them as the *Augustus-Gold* subset. Objects with FLAG_FIT = FALSE are deemed “acceptable” fits, which will henceforth be referred to as the *Augustus-Silver* subset.

In Figure 5, we show the impact each of these flags has on the distribution of stellar distances. Overall, we find that most poor fits tend to happen preferentially at either small distances ($\lesssim 5$ kpc) or extremely large ones ($\gtrsim 15$ kpc). Internal investigation reveals this can be due to a variety of failure modes, some of which are outlined below.

1. *Bad photometry.* One outlying band will lead to extremely poor fits. This is more common near the Pan-STARRS $r \sim 14$ mag saturation limit.
2. *Failed crossmatching.* Multiple objects within the same 1'' radius can be inappropriately matched, leading to “mixed” SEDs that are difficult to model.
3. *Blending effects.* In crowded regions, significant portions of the flux at a given position may be contributed by nearby objects, which can impact the measured flux densities for any particular source. This becomes stronger at lower $|b|$ values.
4. *Quasar/galaxy contamination.* As discussed in Green et al. (2019), quasars and other point sources can often contaminate these samples, especially at fainter magnitudes. Since these have very different SEDs compared to our stellar models, they often are poorly fit.
5. *Unresolved binaries.* A nonnegligible fraction of sources in *Augustus* are expected to be in unresolved binaries

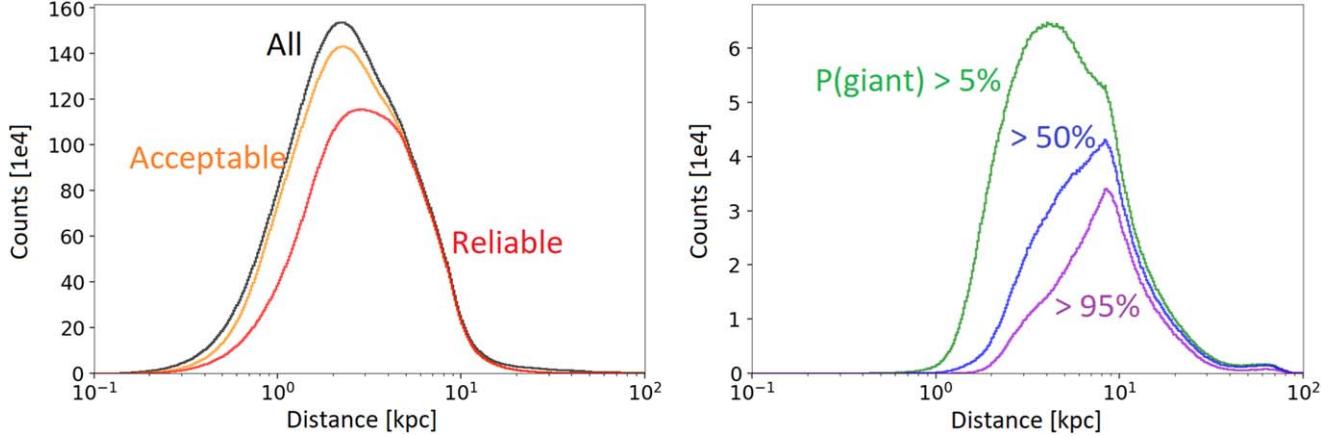


Figure 5. A distribution of the distances (taken from a random posterior sample) for the 170 million objects in our catalog. The left panel shows the number of objects for the entire sample (black), all the sources with “acceptable” fits (the Augustus-Silver subset; orange), and all sources with “reliable” posteriors (the Augustus-Gold subset; red). There are around 125 million sources with reliable posteriors that have distances up to tens of kiloparsecs. The right panel shows the subset of sources for which the probability of being a giant (defined as $\log g < 3.5$) is $>5\%$ (green), $>50\%$ (blue), and $>95\%$ (purple). These panels illustrate there are potentially millions of photometrically classified giants in the sample, although only a few million at very high confidence. See Section 4 for more information on the cuts applied here.

(see, e.g., Belokurov et al. 2020). These are not modeled in this work.

6. *Missing models.* Since our grid only goes down to $M_{\text{init}} = 0.5 M_{\odot}$ and only includes stellar models after they reach the main sequence (EEP = 202), nearby sources that have $M_{\text{init}} \lesssim 0.5 M_{\odot}$ or EEP < 202 are not part of our grid and therefore will be mismodeled. This can also occur if sources fall between our grid points, which has coarser spacing on the high-mass end.
7. *Strong prior disbelief.* As shown in Speagle et al. (2024), BRUTUS can fail to locate the correct solution if it is sufficiently disfavored by the prior. This leads to mismodeling of the SED.
8. *Imprecise photometric parallaxes.* For nearby sources with extremely high signal-to-noise ratio (S/N) astrometric parallax measurements from Gaia DR2, the model grid in BRUTUS may be too coarse to estimate distances with the accuracy needed to match the observed parallax measurements within their measurement uncertainties.
9. *Heavily extinguished.* The default modeling used in this work assumes that $0 \leq A_V \leq 6$. In nearby regions with large A_V (and possible large variations in R_V), the models will fail to reproduce the heavily extinguished SEDs.

By contrast, we find that the vast majority of cases where `FLAG_GRID = TRUE` occur when the initial mass M_{init} hits the lower edge of the grid. Imposing this flag therefore imposes a de facto cut on initial mass, limiting the catalog to mostly sources above $M_{\text{init}} \gtrsim 0.55 M_{\odot}$. This also explains why imposing this cut almost exclusively removes sources at smaller distances, where we are more sensitive to lower-mass objects.

In the right panel of Figure 5, we highlight subsets of Augustus-Gold as a function of the probability that a source is classified as a “giant,” which we define as the probability that it has a $\log g < 3.5$. The vast majority of stars in our sample (~ 150 million) are classified as dwarfs with $P(\text{giant}) \sim 0$, such that even allowing stars that have only $P(\text{giant}) > 5\%$ only includes roughly nine million objects. Imposing even stricter criteria such as $P(\text{giant}) > 50\%$ or P

(giant) $> 95\%$ leaves around five million and 3.5 million sources, respectively. Given that BRUTUS is inherently biased against classifying sources as giants (Speagle et al. 2024), we find these to be likely underestimates of the true number of giants in our catalog. Regardless, a sample of >3 million giants at high latitude is already several orders of magnitude larger than targeted spectroscopic surveys such as the Hectochelle in the Halo at High Resolution (H3) survey (Conroy et al. 2019).

In Figure 6, we show the statistical uncertainties in the estimated distances within Augustus-Gold as a function of Pan-STARRS r -band magnitude, parallax S/N $\varpi_{\text{S/N}}$, and median estimated distance. We find typical statistical uncertainties $\sim 3\%-5\%$ near $r \sim 14$ mag that degrade smoothly to $8\%-10\%$ at $r \sim 20$ mag and $\varpi_{\text{S/N}} \lesssim 1$. Note that these uncertainties do not include systematic uncertainties related to issues with the underlying stellar models, the assumed Galactic priors, etc. While in general estimating overall uncertainties from 1σ and 2σ scatter in the distance realizations give consistent answers, these diverge strongly for median distances $d_{\text{med}} \gtrsim 10$ kpc, at which point confusion between dwarf and giant solutions lead to multimodal distance estimates and largely inflated uncertainties in the tails. These estimates, especially on the brighter end, are similar to recent work from Anders et al. (2019) and Bailer-Jones et al. (2021) and an improvement over the the purely geometric uncertainties from Bailer-Jones et al. (2018, 2021; see Section 5.3).

The distribution of random posterior samples from all the stars in Augustus-Gold is shown in Figure 7. As expected for objects with $|b| > 10^\circ$, the vast majority have $A_V \lesssim 2$ mag with $>50\%$ of the sample having $A_V < 0.4$ mag. The metallicity distribution of the sample peaks around $[\text{Fe}/\text{H}]_{\text{init}} \sim -0.5$, similar to that of the thin disk in our prior, although there is a substantial tail out to metallicities as low as $[\text{Fe}/\text{H}]_{\text{init}} \sim -2.5$. As expected, most sources have subsolar initial masses, with 95% having $0.55 M_{\odot} \lesssim M_{\text{init}} \lesssim 1.05 M_{\odot}$, and are located on the main sequence with EEP < 454 . We do, however, observe a substantial tail of stars up to and beyond the main-sequence turnoff (EEP > 454).

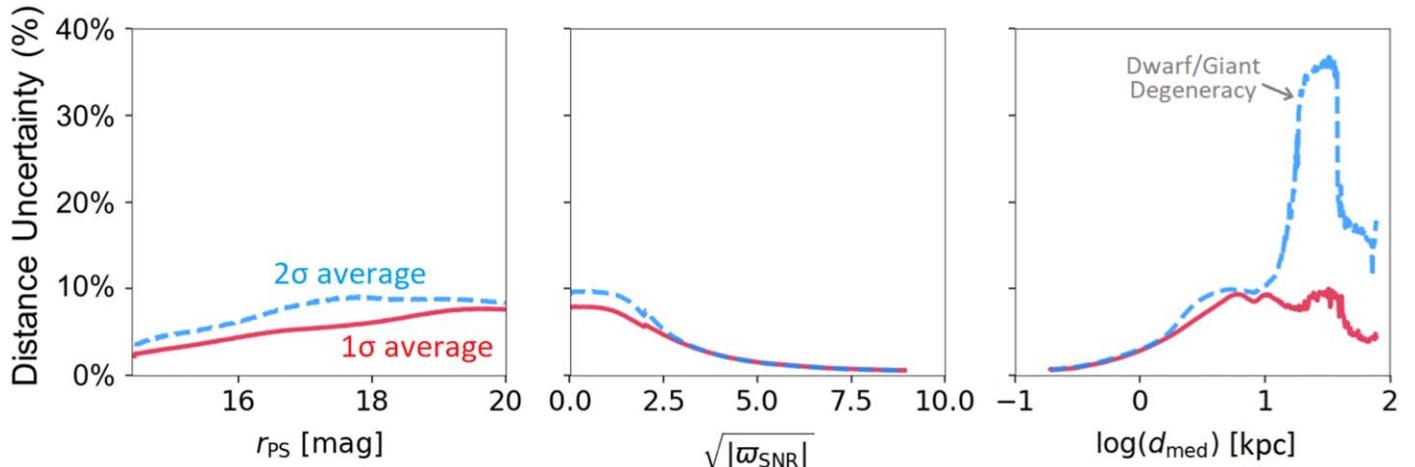


Figure 6. Estimates of the statistical distance uncertainties (i.e., excluding systematic uncertainties) as a function of Pan-STARRS r -band magnitude (r_{PS} ; left), parallax S/N (middle), and median estimated distance (right). These are estimated using the distribution of random distance realizations around the median distance estimates, with results from half the width of the 68% credible interval (“ 1σ average;” solid red) and a quarter the width of the 95% credible (“ 2σ average;” dashed blue). We find typical uncertainties of 8%–10% at the faintest magnitudes and lowest parallax S/Ns, with uncertainties derived from the 95% confidence intervals (CIs) larger than those derived from the 68% CIs. These differences can become particularly pronounced at larger distances (right panel) due to possible degeneracies between nearby dwarf (main-sequence) and faraway giant (post-main-sequence) stellar evolutionary solutions.

5. Results and Discussion

We now wish to highlight some preliminary results illustrating the quality of the data from the 125 million stars in Augustus-Gold.

5.1. Reproducing the Gaia Color-Magnitude Diagram

Given known systematics in the theoretical stellar models used in this and other works (Choi et al. 2016; Anders et al. 2019; Speagle et al. 2024), one way to examine the reliability of the results is to examine the ability of the stellar models to reproduce the empirical CMD. In order to accomplish this, we purposefully did not use any of the observed Gaia DR2 photometry when computing the predictions. Since these bands are so much broader than the underlying Pan-STARRS bands that they overlap with, they can serve as a (limited) posterior predictive check on the overall quality of the fits derived using the Pan-STARRS, 2MASS, UKIDSS, and unWISE data. In other words, we can test what the photometry in the Gaia bands should be (the “predictive”) based on the model constraints imposed from the other fitted bands (the “posterior”).

Using the BRUTUS package, we take a random posterior sample of the distance d , extinction A_V , and differential reddening R_V from each object and use it to compute the corresponding distance modulus μ and extinctions A_G , A_{BP} , and A_{RP} in the Gaia bands using the filter curves from Maíz Apellániz & Weiler (2018) as described in Speagle et al. (2024). We then use these to “dedistanc” and deredden the observed photometry. We then compare this “empirical” CMD to the “intrinsic” CMD predicted directly from the corresponding model. The results of this exercise for all sources in Augustus-Gold along with a subset of 12.5 million sources with full photometric coverage and reasonably constrained distances (<30% 2σ errors) are shown in Figure 8. We find excellent overall agreement between the intrinsic CMD predicted by the models and the “empirical” CMD derived from the data, with uncertainties mostly scattered in the direction of the reddening vector. Note that this scatter is expected given the incorrect sign of the covariances between A_V and R_V with distance discussed at the beginning of

Section 4. Some exceptions to the excellent overall agreement include noticeable gridding effects in post-main-sequence stellar evolutionary phases and higher masses as well as clear discrepancies properly modeling the horizontal giant branch.

5.2. Comparison to STARHORSE

We also compare our results to previous work. In particular, Anders et al. (2019; henceforth A19) use a similar approach to derive distances to 270 million sources with Gaia $G < 18$ mag. As in this work, their approach also involved using a set of theoretical isochrones applied to similar photometric and astrometric data sets. The main differences between the two studies are as follows.

1. A19 utilize the PARSEC v1.2 models (Bressan et al. 2012; Chen et al. 2014; Tang et al. 2014) rather than the (color-corrected) MIST models employed here.
2. A19 use the empirical dust curve from Schlafly et al. (2016) to model A_V variations but no R_V variations. Here we use the dust curve from Fitzpatrick (2004) and models variations in both A_V and R_V .
3. A19 use a different form for the underlying Galactic prior that includes an additional bulge component and different thin disk, thick disk, and halo properties.
4. A19 do not apply a 3D dust prior to supply additional constraints on A_V .
5. A19 use a version of the STARHORSE code (Santiago et al. 2016; Queiroz et al. 2018; A19) to fit a grid over stellar models, d , and A_V and evaluate the prior over these data points. BRUTUS samples d , A_V , and R_V values and attempts to integrate over the Galactic prior while doing so.
6. A19 fit photometric data from Gaia, Pan-STARRS, 2MASS, and AllWISE. This work excludes fitting photometric Gaia data and uses unWISE instead of AllWISE.
7. A19 apply different Gaia DR2 parallax zero-point corrections and photometric offsets/errors compared to this work.

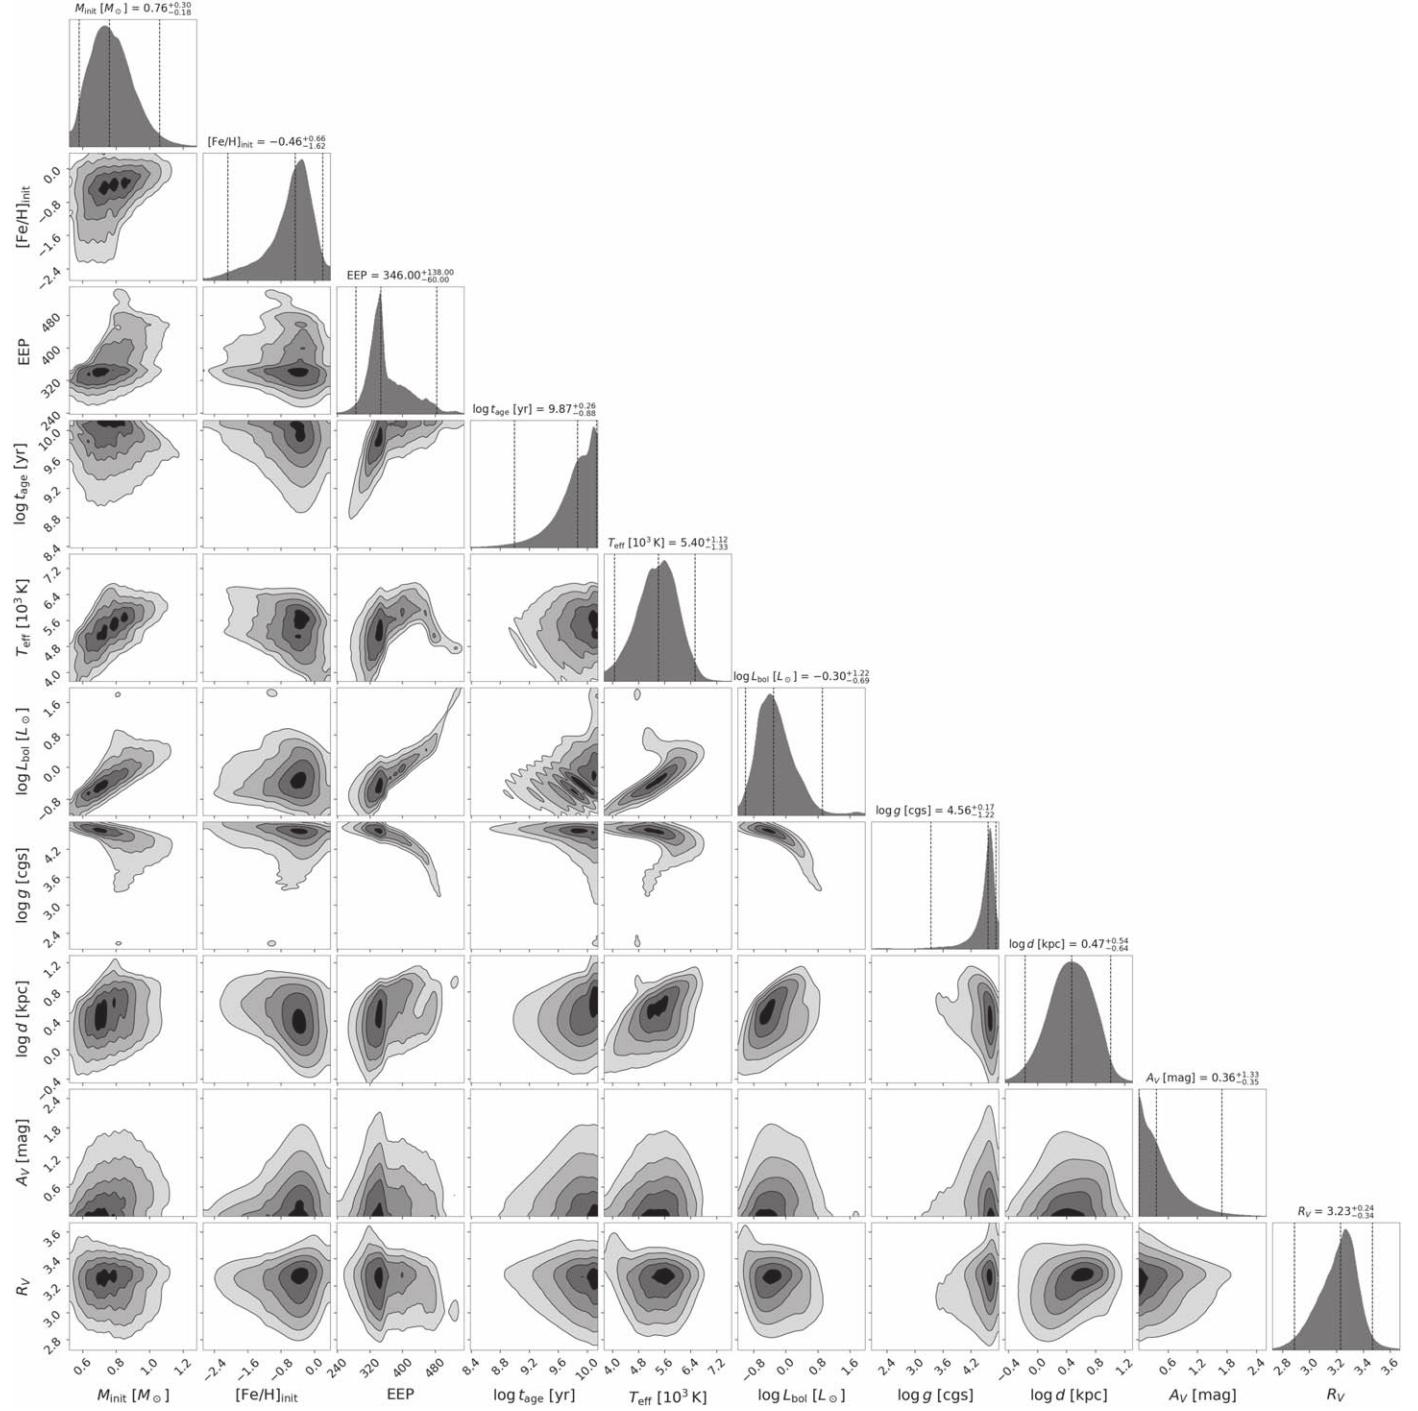


Figure 7. A corner plot showing the collective 1D and 2D posterior distributions for the parameters constrained for each of the 125 million Augustus-Gold stars. These parameters are (from left to right): the initial mass (M_{init}), initial metallicity ($[\text{Fe}/\text{H}]_{\text{init}}$), EEP, age ($\log t_{\text{age}}$), effective temperature (T_{eff}), bolometric luminosity ($\log L_{\text{bol}}$), surface gravity ($\log g$), distance $\log d$, extinction (A_V), and “differential” extinction (R_V). The titles of each column show the median and the interval encompassing 95% of the sample. As expected, the majority of the sample comprises sources with low extinction ($A_V \lesssim 1$ mag) and initial masses ranging from $0.55 M_{\odot} < M_{\text{init}} < 1.2 M_{\odot}$, with the lower limit imposed by the lower M_{init} bound on our underlying grid of stellar models. Gridding effects can be seen in a few panels.

As A19 perform significant vetting of their associated catalog across a wide range of surveys, we want to confirm that we are able to recover similar results for sources that overlap between the two catalogs. After crossmatching sources based on their Gaia object ID and only selecting “high-quality” sources with $\text{SH_GAIAPLFLAG} = 000$ and $\text{SH_OUTFLAG} = 00000$, we find roughly 26 million sources that overlap between the two catalogs. Part of the reason for this small overlap is that A19

only go down to $G = 18$ mag rather than the $r = 20$ mag used in this work; the other main reason is that this work excludes all sources with $|b| < 10^{\circ}$ (where the vast majority of stars actually lie) as well as sources in the Southern Hemisphere.

Figure 9 shows a comparison between the two data sets for a few parameters of interest including d , $\log g$, T_{eff} , and A_V . Overall, we find the estimated distances between the two data sets are extremely consistent with each other, although

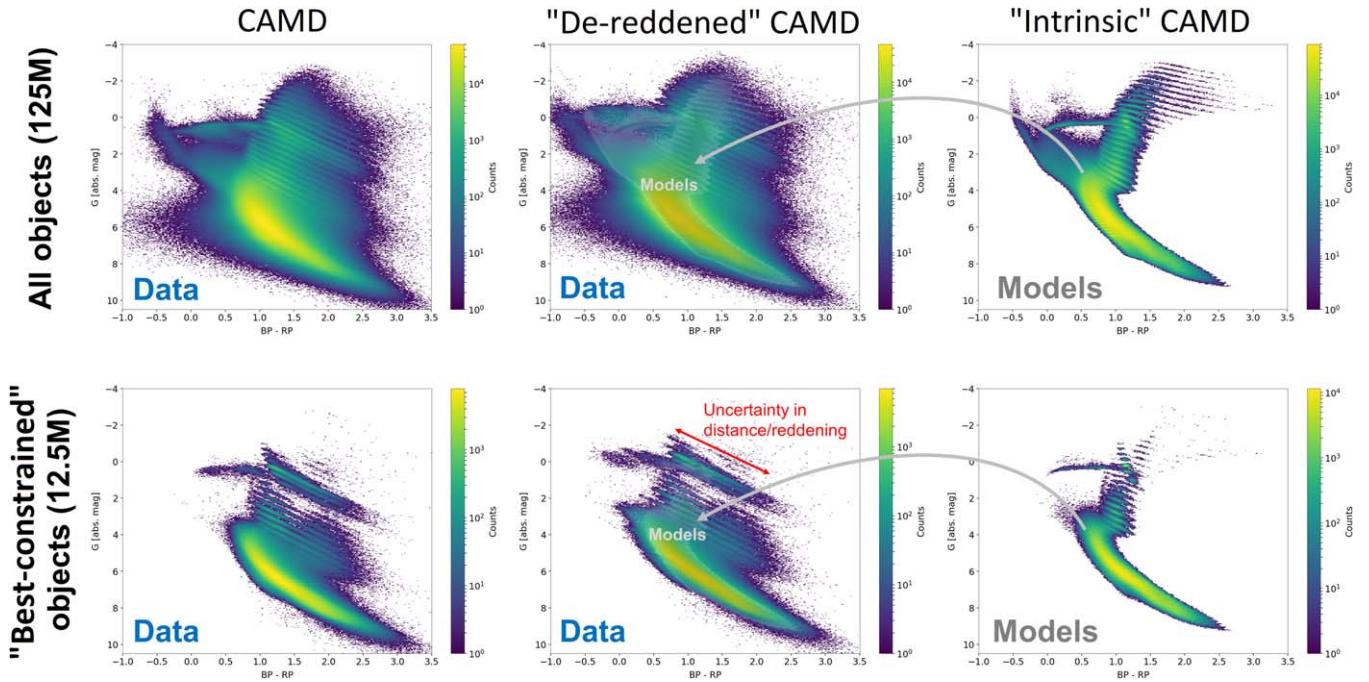


Figure 8. The Gaia G vs. $BP - RP$ color–absolute magnitude diagram (CAMD) for all sources in Augustus–Gold (125 million objects; top) and “best-constrained” sources (12.5 million objects; bottom) which have 10 bands of Pan-STARRS, 2MASS, UKIDSS, and/or unWISE photometry and 95% distance credible intervals that are $<30\%$ of the median distance (i.e., $|(d_{97.5} - d_{2.5})/d_{50}| < 0.3$). The left panels show the CAMD after shifting sources to $d = 10$ pc using a random sample drawn from the stellar posterior. The middle panels show the “de-reddened” CAMD using the A_V and R_V values from the same random posterior sample and the predicted linear reddening vector from the stellar parameters associated with them. The right panels show the predicted CAMD computed directly from the models; this is also overplotted as the light gray shaded region in the middle panel for ease of comparison. As the Gaia photometry was not used when deriving the stellar posteriors but whose wavelength coverage overlaps with the Pan-STARRS data, this serves as a useful but limited check on the internal self-consistency and overall quality of the results. We find excellent overall agreement between the intrinsic CAMD predicted by the models and the “empirical” CAMD derived from the data, with uncertainties mostly scattered in the direction of the reddening vector. Some exceptions to this include noticeable gridding effects in evolved stellar evolutionary phases (thin overdense regions), issues modeling the horizontal giant branch (upper left regions), and dwarf/giant misclassification (middle right regions).

Augustus–Gold generally prefers sources nearer than a few kiloparsecs to be slightly closer. This is likely due to small differences in the underlying Galactic prior. We also find strong agreement between the predicted log g values (outside of the gridding effects due to coarse sampling of post-main-sequence evolutionary phases in this work), although there is a clear excess of sources that are classified as low-mass dwarfs in Augustus–Gold relative to A19. The reason for this discrepancy can be seen when examining T_{eff} , which shows that while values are consistent across both data sets below $T_{\text{eff}} \lesssim 6000$ K, sources in Augustus–Gold with estimates of $T_{\text{eff}} \sim 6000$ K are much more likely to have associated estimates of $T_{\text{eff}} \gtrsim 8000$ K in A19. To reproduce the observed SED, these intrinsically bluer sources need to have more reddening from foreground dust, leading to an expected increase in higher associated values of A_V for some sources in A19 relative to Augustus–Gold. Note that while Augustus–Gold strictly enforces $A_V \geq 0$ mag and A19 do not, we find the impact of this choice does not appear to hamper comparisons other than around $A_V \sim 0$ significantly. We expect that this will slightly bias the distances and intrinsic colors for stars behind very little dust, but helps to avoid scenarios that can arise when the reddening vector is allowed to compensate for systematic color offsets in the model by exploring nonphysical solutions.

5.3. Comparison to Bailer-Jones et al. (2018, 2021)

In addition to A19, we also compare our distances to those derived purely from Gaia data based on both DR2 (Bailer-

Jones et al. 2018, hereafter BJ18) and EDR3 (Bailer-Jones et al. 2021, hereafter BJ21), which was released between constructing the initial catalog and writing up this manuscript. While BJ18 utilize only parallax information when deriving their distances, BJ21 include both parallax-only estimates and ones that also include contributions from an empirical model that incorporates Gaia photometry. After crossmatching with both catalogs, we find roughly 126 million sources in common between those samples and the Augustus–Gold sample.

In Figure 10, we compare the distance estimates from BJ18 and BJ21 to those from Augustus–Gold. Overall, we find excellent agreement between Augustus–Gold, BJ18, and BJ21 within a few kiloparsecs, which further improves when considering the BJ21 estimates derived including Gaia photometry. Given that the BJ18 and BJ21 estimates were derived using substantially different prior assumptions from our model and, in the case of BJ21, mutually exclusive photometric data sets (as no Gaia photometry was used to derive any stellar properties reported in this work), this agreement lends further confidence to the overall accuracy of our distance estimates.

5.4. Galactic Substructure seen in Augustus

As discussed in Speagle et al. (2024), inference from photometry alone is strongly influenced by the underlying Galactic priors and even including strong constraints from parallaxes can still lead to biases in the inferred stellar properties without tight constraints on A_V (i.e., some knowledge of the intrinsic SED). As such, it is important to

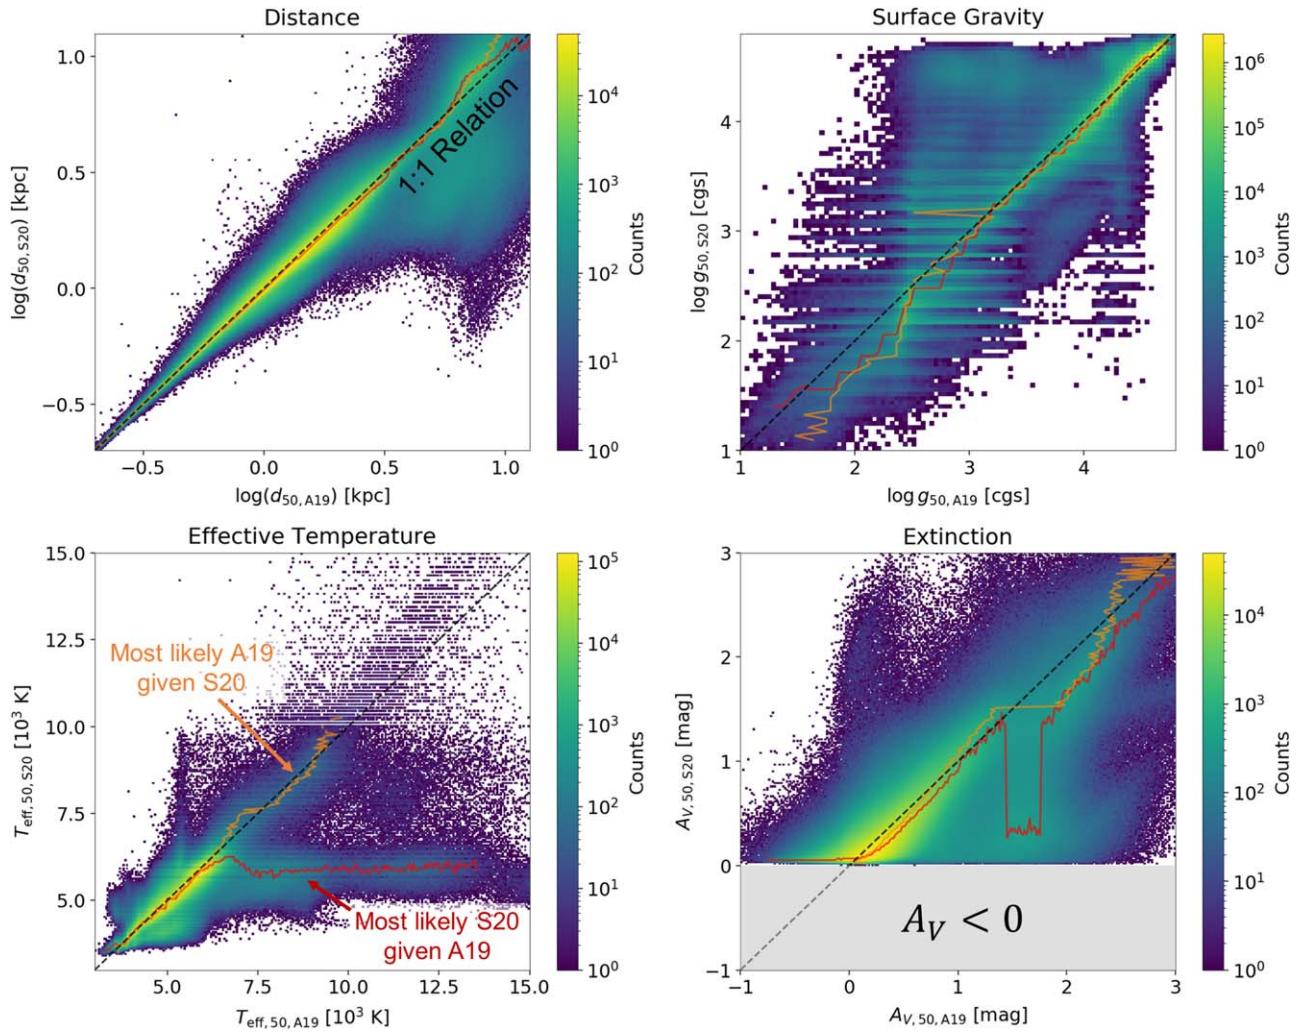


Figure 9. Comparisons between the median distance (top left), $\log g$ (top right), T_{eff} (bottom left), and A_V (bottom right) in this work (Augustus-Gold; “S20”) and A19 for 26 million sources present in both catalogs. In each panel, the most common (i.e., likely) associated value from Augustus-Gold given a value from A19 (i.e., y given x) is shown as a solid red line, and the most common associated value from A19 given Augustus-Gold (i.e., x given y) is shown as a solid orange line. The 1:1 relationship is also overplotted as a dashed black curve. We see that the estimated distances from both catalogs are extremely consistent with each other across most distances, although this work generally prefers source below a few kiloparsecs to be slightly closer. The $\log g$ values are also consistent with each other (outside of gridding effects), although there is a clear excess of sources that are classified as low-mass dwarfs in Augustus-Gold relative to A19. The reason for this discrepancy can be seen when examining the estimated T_{eff} , which shows that while the values are consistent across both data sets below $T_{\text{eff}} \lesssim 6000$ K, A19 prefer to make sources substantially hotter than Augustus-Gold for $T_{\text{eff}} \gtrsim 6000$ K. This leads to a corresponding increase in the number of sources with $A_V \gtrsim 1$ mag, where the higher reddening combined with the intrinsically bluer (hotter) colors end up giving similar SEDs as intrinsically redder (cooler) sources with less reddening. Note that the naming convention “S20” is based on the fact that the original analysis was included in the lead author’s PhD Thesis, which was accepted in 2020 and can be found online at <https://nrs.harvard.edu/URN-3:HULINSTREPOS:37365889>.

investigate just how much information we are able to recover relative to the prior.

It has been demonstrated in Green et al. (2015), Chen et al. (2019), A19, and other works that there is enough information in stellar photometry to recover distance and extinction estimates to stars with enough precision to construct detailed, accurate 3D dust maps. A19 show that these estimates may also be detailed enough to begin resolving large-scale features such as the Galactic bar. Given that our map targets high-latitude regions and goes substantially deeper than that work, we want to investigate whether we too can (in principle) recover large-scale substructure purely from astro-photometry (i.e. using astrometry and photometry) alone.

In Figure 11, we plot the mean proper motion directions for all the sources in Augustus-Gold with $n_{\text{side}} = 64$ healpix resolution, split into (overlapping) median distance bins ranging from $d_{50} = 0$ kpc to $d_{50} = 50$ kpc and colored by

counts in each pixel. Overall, we see clear evidence of large-scale features in our maps, including evidence of the Sagittarius Stream (bottom right) and the Monoceros Ring (middle right; Newberg et al. 2002; Jurić et al. 2008; Purcell et al. 2011; Gómez et al. 2013; Laporte et al. 2018a, 2018b). We also observe issues where systematics clearly play a role in the inferred stellar properties, especially near the Galactic center and the Galactic plane as well as in regions of substantial foreground extinction (where we either miss stars entirely or likely somewhat mismodel them).

To get a sense for how significant these features are, we need to compare them against what we expect given our Galactic prior. In Figure 12, we show the exact same plot except this time colored by the expected number of counts in each pixel, normalized so that each distance bin contains the same total number of stars. In this version, we see no evidence for any substructure in density alone (since our prior includes no

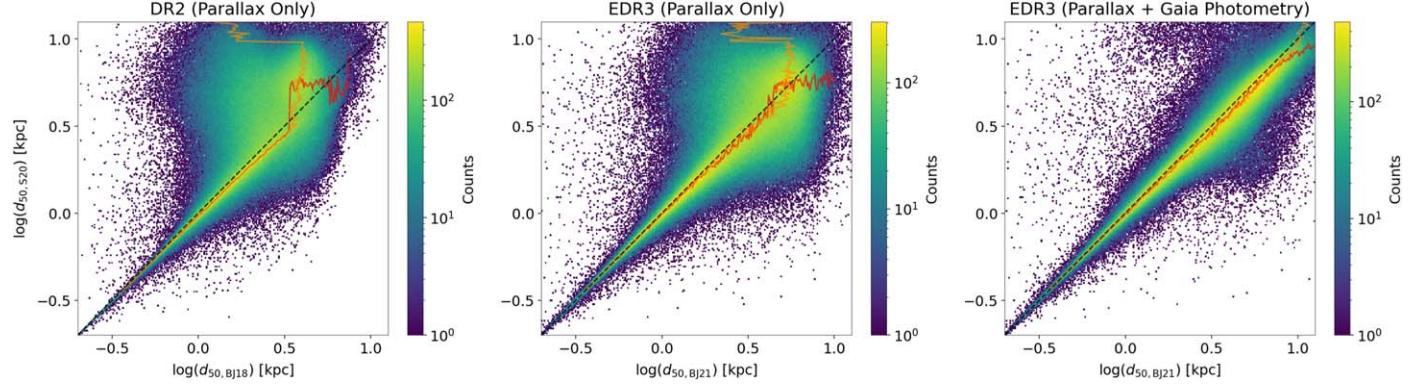


Figure 10. As Figure 9, but now comparing the median distances from this work with those derived from Gaia DR2 parallaxes only from BJ18 (left) and from Gaia EDR3 parallaxes from BJ21 without (middle) and with (right) Gaia photometry. The most common (i.e., likely) associated value in Augustus-Gold given a value from BJ18 or BJ21 (i.e., y given x) is shown as a solid red line, and the most common associated value from BJ18 or BJ21 given Augustus-Gold (i.e., x given y) is shown as a solid orange line. The 1:1 relationship is overplotted as a dashed black curve. Estimated distances from all three catalogs are in excellent agreement with those derived here for stars within a few kiloparsecs using only the parallax and agree out to further distances when also considering the BJ21 estimates that also incorporate information from photometry. Disagreements at larger distances compared to the parallax-only estimates generally arise due to stronger distance constraints from photometry overwhelming the distance estimates for objects with low parallax S/Ns where the Galactic prior tends to dominate the inference. As with A19, this work prefers sources to be slightly closer than in BJ18 and BJ21.

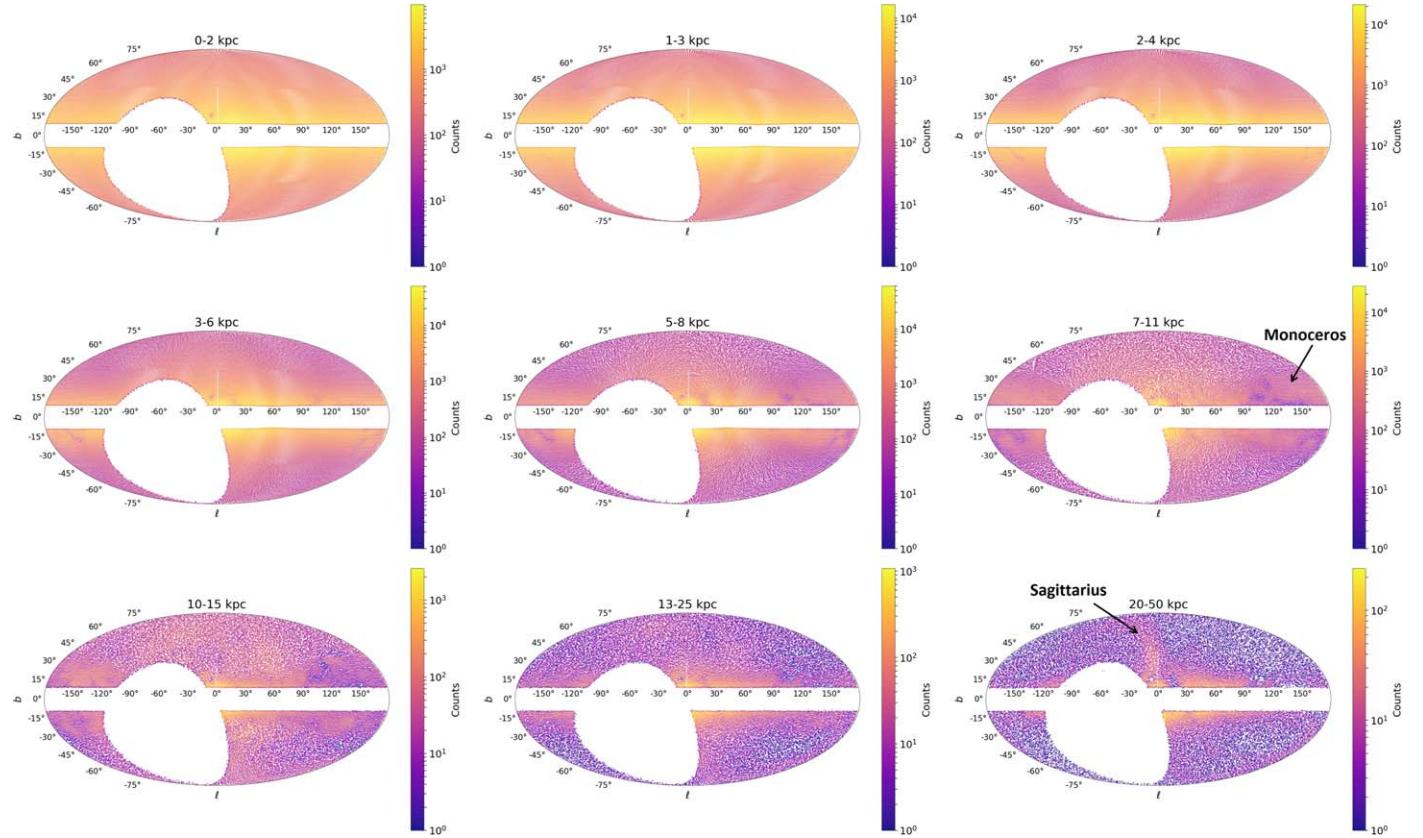


Figure 11. The mean number density of stars in (overlapping) distance bins ranging from $d = 0$ –50 kpc at a healpix resolution of $nside = 64$ (see Figure 2). The bulk motion of stars within each healpix pixel (corrected for solar reflex motion) are indicated by small arrows and estimated using the mean proper motion. We can see clear evidence of large-scale features in our maps, including evidence of the Sagittarius Stream (bottom right) and the Monoceros Ring (middle right) along with “missing regions” where substantial foreground dust extinction removes stars from our catalog. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_multipanel_toggle.html. An interactive version of Figures 11–14 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

kinematic information). This makes sense, since our prior was, by construction, a smooth model of the Galaxy that did not account for any small-scale structure.

In Figure 13 we compare the ratio of the observed number of counts to the expected number of counts. As expected, we have an underdensity of sources in the direction of the Galactic

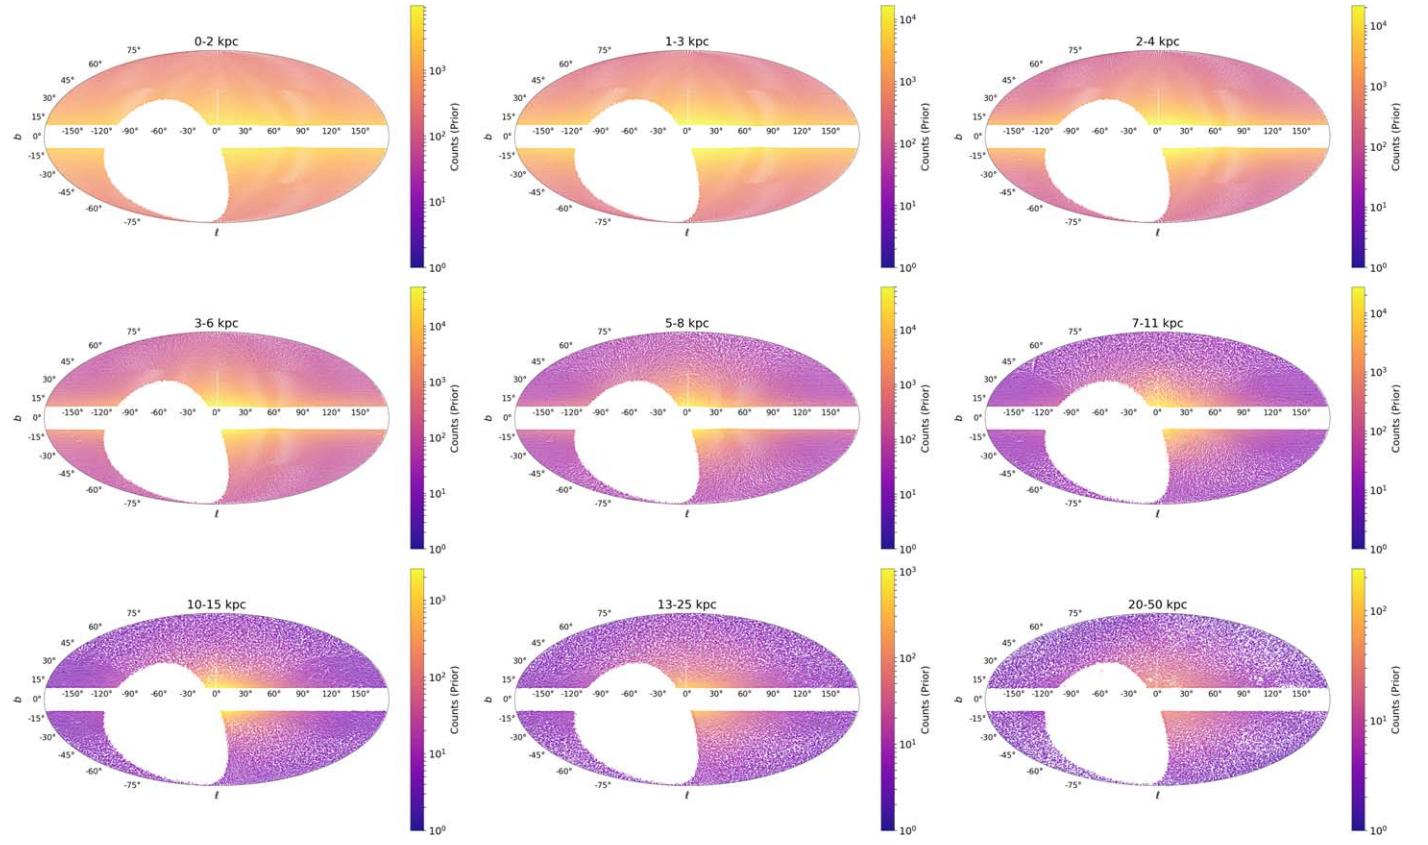


Figure 12. As Figure 11, but now showing the expected distribution of sources from the Galactic prior. These have been normalized so that the total number of sources in each distance bin are the same in both figures. Since the prior contains only smooth components, it does not show any evidence for additional substructure compared to Figure 11. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_multipanel_toggle.html. An interactive version of Figures 11–14 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top right provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

center, where there is a substantial amount of dust extinction and we are likely incomplete. We can clearly see the overdensity corresponding to the Monoceros Ring appears as early as $d \sim 3$ kpc, peaks at $d \sim 9$ kpc, and extends out to $d \sim 13$ kpc. While there are strong asymmetries as a function of ℓ and b , it is difficult to interpret these differences since these regions have dust clouds along the line of sight.

Beyond $d \sim 15$ kpc, we see clear evidence for the Sagittarius Stream, which has densities at high latitudes similar to those near the plane. These densities are in fact so large that they dominate the normalization of the expected number of counts, leading to these structures being considered “normal” with ratios ~ 1 while the rest of the halo is considered “underdense.”

In Figure 14, we now highlight differences between the predicted mean $[\text{Fe}/\text{H}]_{\text{init}}$ from the Galactic prior (which is spatially independent within each component of our prior) and the mean $[\text{Fe}/\text{H}]_{\text{init},50}$ derived from the data. Deviations can be seen in all distance bins, although the size of the deviations increases as a function of distance. Here we see clear evidence for issues related to survey coverage at nearby distances with visible survey stripes and at far distances near the Galactic plane. There are also clear systematics in $[\text{Fe}/\text{H}]_{\text{init}}$ estimation correlated with foreground dust extinction, with regions with large A_V having discrepant $[\text{Fe}/\text{H}]_{\text{init}}$ relative to background sources. That said, we do see strong, correlated evidence for lower-than-expected $[\text{Fe}/\text{H}]_{\text{init}}$ values for stars in the Monoceros Ring and higher-than-expected $[\text{Fe}/\text{H}]_{\text{init}}$ values for stars

in the Sagittarius Stream. These findings suggest that we should be able to use photometric metallicities to explore the kinematic and chemical origins of these structures, such as to follow up on kinematic and chemical separations with other structures from the Monoceros Ring (e.g., Laporte et al. 2020).

In Figure 15 we plot the associated mean tangential velocities in each pixel derived from the median distances in *Augustus-Gold* and measured proper motions from Gaia, corrected for solar reflex motion using ASTROPY (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018). The overall structure in velocities observed at $d \lesssim 6$ kpc agrees with what we would expect from geometry, where sources orbiting in the disk are moving directly along/opposite our line of sight. Kinematically coherent large-scale structure associated with the Monoceros Ring and Sagittarius Stream are clearly visible. We also clearly see the presence of known large open/globular clusters, which show up as “outliers” in tangential velocity in a given distance bin relative to the underlying background. Both of these results are encouraging.

Finally, we examine our detection of the Sagittarius Stream in more detail. We transform the coordinates from our sources from Galactic coordinates (ℓ, b) to coordinates aligned with the orbital plane of the Sagittarius Stream (Λ, β) from Law & Majewski (2010). In Figure 16 we try to compare these results directly with those from simulations taken from Law & Majewski (2010) by using the density contrast for sources with $|\beta| < 10$ compared with those with $10 < |\beta| < 30$ as a tracer of

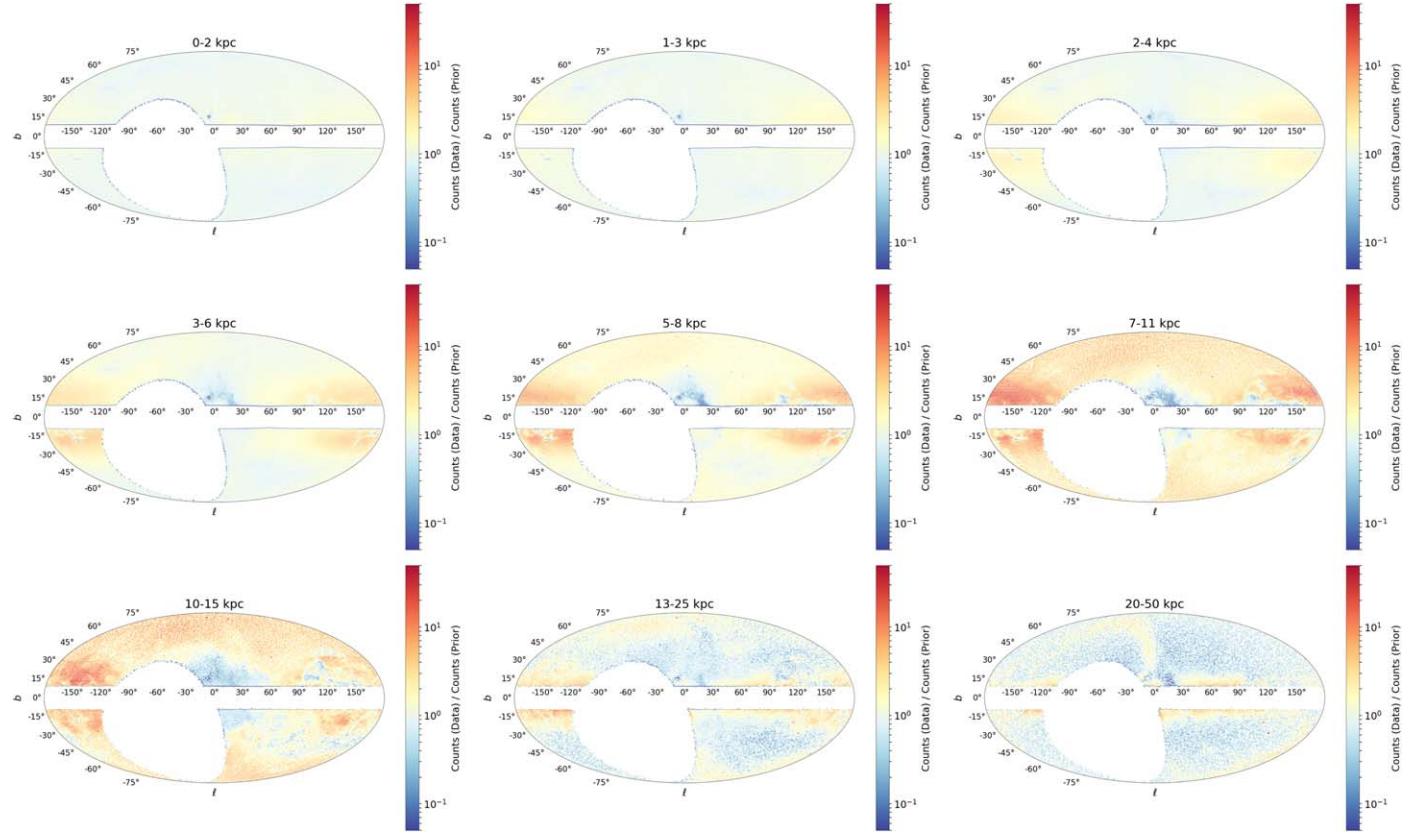


Figure 13. As Figure 11, but now showing the ratio between the data and the prior. This highlights deviations from the assumed Galactic model (i.e., “the background”) in order to emphasize substructure present in the data. The Monoceros Ring in the Galactic anticenter can be clearly seen as a $\gtrsim 10$ times overdensity relative to the background from $d \sim 5$ –15 kpc. We also observe broad flaring in the disk, which is not modeled in our prior. At $d > 20$ kpc, the Sagittarius Stream is also clearly visible, with number densities also $\gtrsim 10$ times higher than the background. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_multipanel_toggle.html. An interactive version of Figures 11–14 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top right provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

the Sagittarius Stream as a function of d and Λ . We find our results to be qualitatively consistent regardless of whether we use all sources or limit ourselves to only those with P (giant) $> 95\%$. This broad correspondence in overall structure between the astro-photometric (i.e. astrometry and photometry) distances in Augustus-Gold and the results from the simulations lends confidence that our distances and stellar classifications are reliable even out to large distances.

5.5. Visualizing 5D Substructure with ALLSKY

As a final tool to aid exploration and characterization of the data presented in this work, we modify the public, open-source code EARTH²⁸ used to visualize wind and ocean currents on the surface of the Earth to handle the similar types of structure present in 2D velocities in a given set of 3D distance bins. Our public, open-source code ALLSKY²⁹ is able to illustrate velocity motion as moving “streamlines” following simple nonlinear trajectories and allows users to explore various underlying properties of the data (velocity, number density, metallicity, etc.) in various projections (Atlantis, orthographic, equirectangular, etc.) for each of the nine distance bins shown in this work. A screenshot illustrating the code is shown in Figure 1.

The full interactive visualization can be accessed online at <http://allsky.s3-website.us-east-2.amazonaws.com> and the data can be downloaded from the ALLSKY GitHub repository.

5.6. Additional Remarks

While the results described in this section highlight some of the successes of BRUTUS and the overall quality of the Augustus-Gold subsample, we also want to take some time to mention limitations as well as future directions for improvement explicitly. These fall under a few broad categories.

1. *Limited spatial coverage and depth.* For practical reasons, Augustus is limited to only covering the northern sky that overlaps with the 3D dust prior from Green et al. (2019), does not include the large amount of objects at $|b| \lesssim 10^\circ$, and only extends down to $r < 20$. We aim to break away from these limitations in future work by removing our reliance on a previously estimated 3D dust prior, improving the computational efficiency of the underlying BRUTUS code, using larger all-sky data sets such as Gaia EDR3 as a base to expand off of, and using a broader set of surveys to provide deeper optical and (near-)IR photometry in both the north and the south.
2. *Overly simplistic priors.* The current set of priors (see Speagle et al. 2024 for additional details) includes simple

²⁸ <https://earth.nullschool.net/>

²⁹ <https://github.com/joshspeagle/allsky>

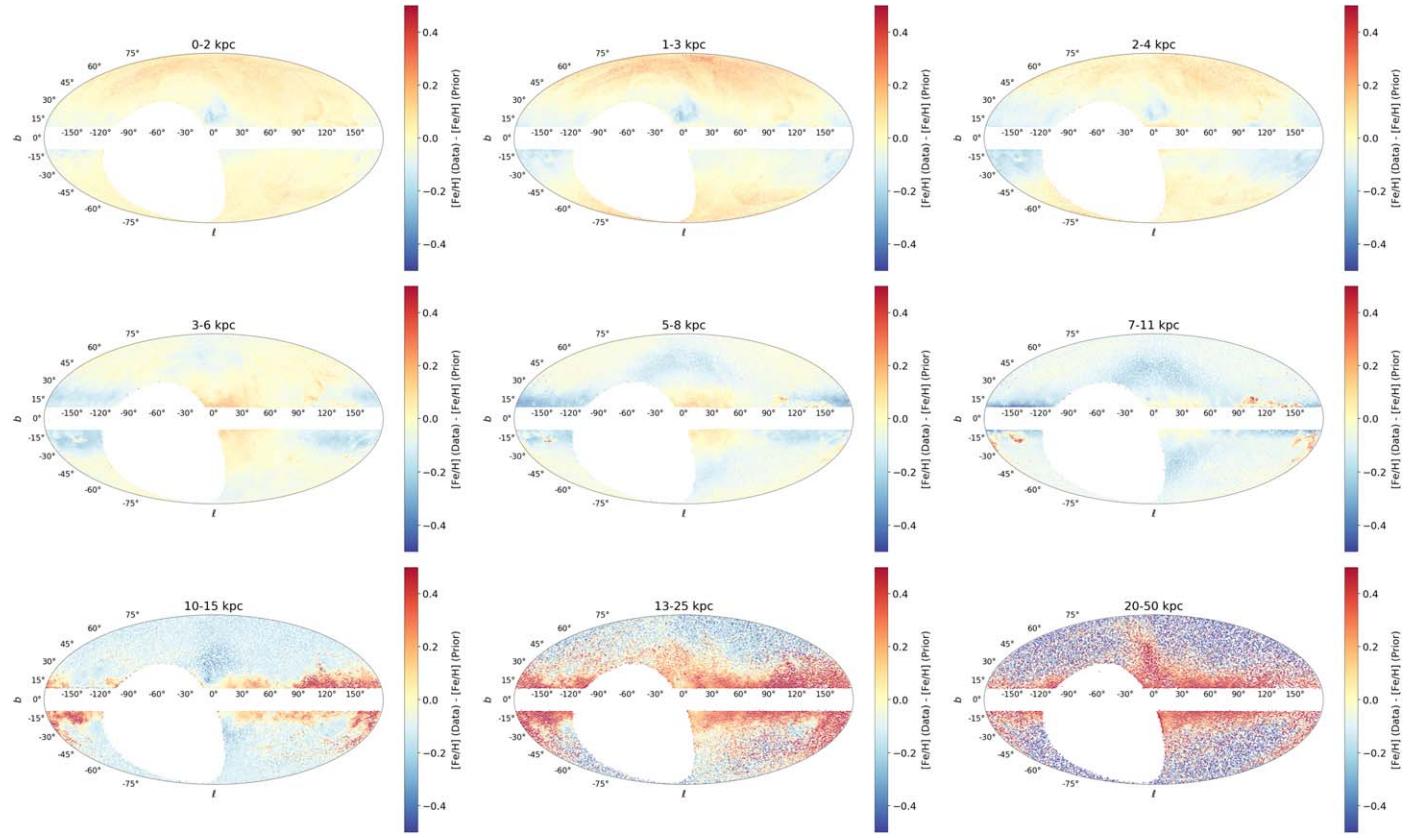


Figure 14. As Figure 13, but now showing the difference between the mean $[\text{Fe}/\text{H}]_{\text{init}}$ from the data and the prior. Deviations can be seen in all distance bins, although the size of the deviations increases as a function of distance. Clear systematics in the $[\text{Fe}/\text{H}]_{\text{init}}$ estimation aligned with foreground dust extinction features and objects near the Galactic plane are visible. We see evidence for lower-than-expected $[\text{Fe}/\text{H}]_{\text{init}}$ values (compared to the prior) for stars in the Monoceros Ring and higher-than-expected $[\text{Fe}/\text{H}]_{\text{init}}$ values for stars in the Sagittarius Stream. An interactive version of this figure is available in the article and also at https://faun.rc.fas.harvard.edu/czucker/Paper_Figures/brutus_multipanel_toggle.html. An interactive version of Figures 11–14 is available. In this interactive version, the projected sky coverage maps are shown at the top while histograms highlighting wavelength coverage are shown at the bottom. Buttons at the top middle allow the user the ability to see how the coverage changes when including/excluding data from Pan-STARRS1, 2MASS, UKIDSS, and unWISE. Buttons at the top right provide pan, zoom, save, and reset functionality; hover the mouse over each button to see what it does. An interactive version of this figure is available.

models for a three-component model that includes a thin disk, thick disk, and halo. Not only are all components overly simplistic smooth models (with no substructure, warps, etc.) with spatially independent metallicity and age constraints, but there is no prior for a bulge or bar. We aim to improve on these in future work.

3. *Low A_V limits.* While the Augustus catalog targeted the halo, there are still regions where extremely high- A_V values occur. Stars in these regions are inherently mismodeled due to the default maximum value of $A_V=6$ imposed in the catalog. This will be raised to much higher values in future work.
4. *Mass limits.* The theoretical MIST isochrones we are using have large systematic biases in the predicted photometric colors below $M \lesssim 0.5 M_{\odot}$, even with some of the empirical corrections implemented in Speagle et al. (2024). This primarily affects our ability to model M dwarfs/giants. We hope to use improved isochrone models in future work to extend our modeling to lower masses.
5. *Model–data mismatch.* Offsets between the predicted photometry vary systematically across the CMD for our given set of isochrones. This imposes systematic limits on parameter recovery above the statistical errors present in the photometric measurements. We hope the incorporation of new data-driven stellar models such as

those from Green et al. (2021) will help to alleviate these issues.

6. *Gridding effects.* As shown most clearly in Figure 9, the current grid does sample a range of surface gravities and effective temperatures but still contains visible gridding effects that could impact inference. We hope to alleviate this in future work through approaches that can apply iterative adaptive refinements.
7. *Nonnormal errors.* It is possible that the assumption of strictly normal uncertainties is not valid for the assumed photometric uncertainties in both magnitude and flux density, leading to incorrect (likely underestimated) statistical uncertainties in the derived properties. While systematic effects from survey photometric pipelines may lead to empirical error distributions with broader tails (e.g., such as the complex photometric processing for data in Gaia DR2; Gaia Collaboration et al. 2018), in general most of the reported uncertainties in the derived properties are dominated by the effects of model–data mismatch. Similarly, while the distribution of statistical parallax uncertainties in Gaia DR2 appears to follow a normal distribution out to several standard deviations, the applied zero-point corrections can lead to systematic offsets that are distinctly nonnormal in nature (Lindegren et al. 2018).

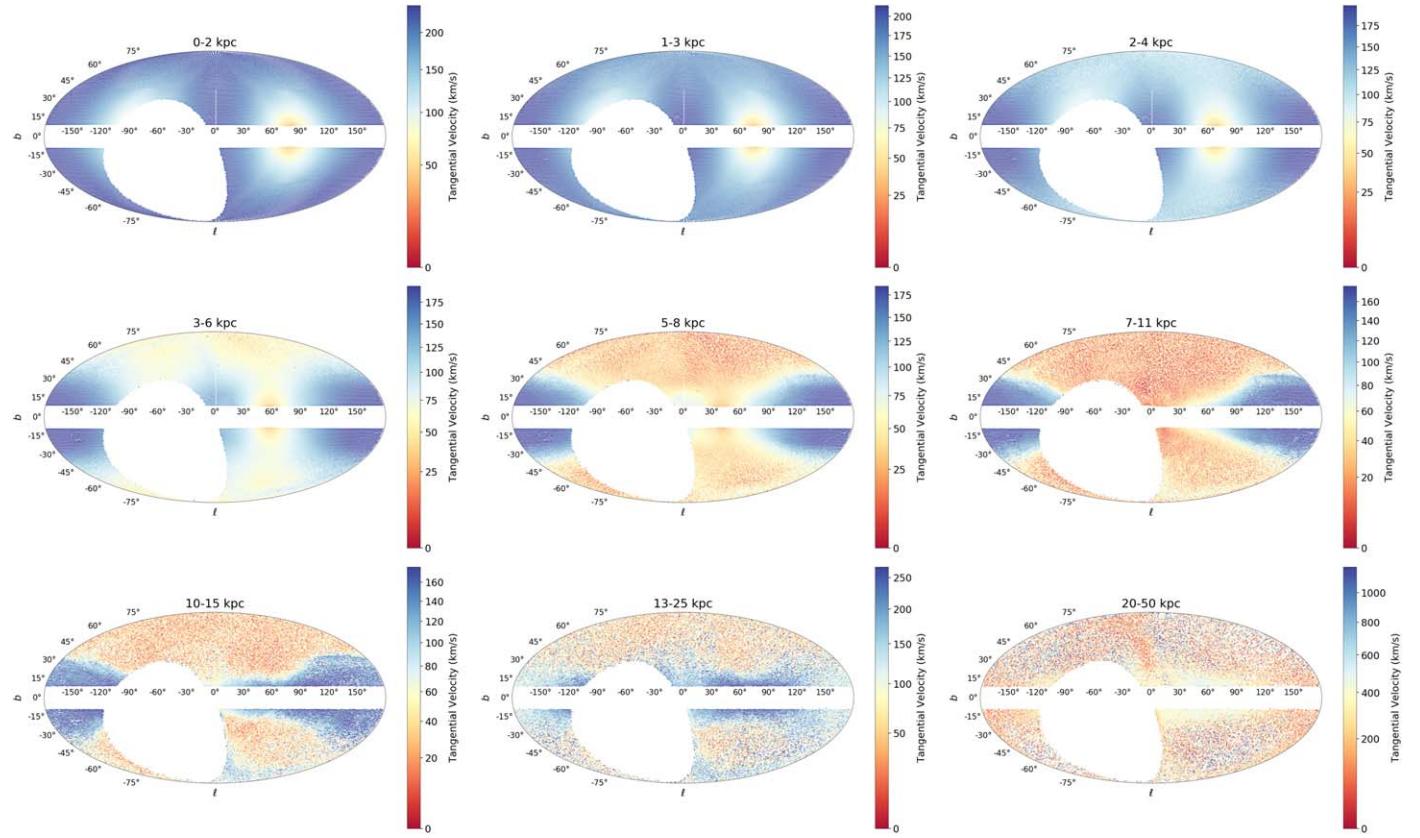


Figure 15. As Figure 11, but now showing the mean tangential velocity (in units of km s^{-1}) corrected for solar reflex motion. The overall structure in the observed velocities at $d \lesssim 6$ kpc is due to pointing directly along/opposite the expected motion of stars orbiting in the disk. Kinematically coherent structures associated with known clusters (which show up as “outliers” relative to the background motion of nearby sources) as well as more extended structures are easily visible.

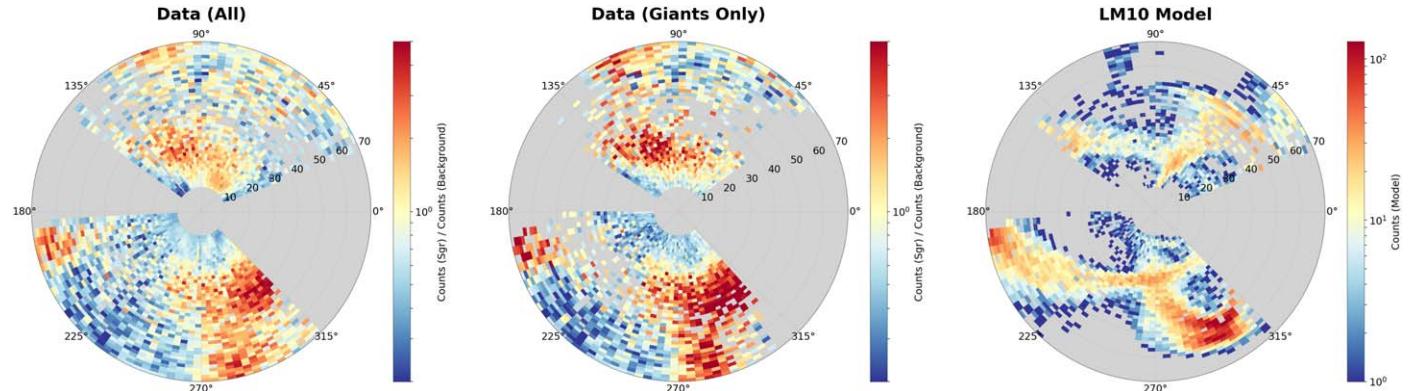


Figure 16. The ratio of the observed number density for objects with $|\beta| < 10$ (“Sagittarius”) and with $10 < |\beta| < 40$ (“Background”) as a function of distance and Λ_{Sgr} , where (Λ, β) are coordinates in the Sagittarius orbital plane. The results for all stars with median distances $10 \text{ kpc} < d_{50} < 70 \text{ kpc}$ and the subset with $P(\text{giant}) > 95\%$ are shown in the left and middle panels, respectively. The stellar density from stars taken from the Law & Majewski (2010) simulations with $|\beta| < 10$ are shown in the right panel. The broad correspondence in overall structure between the astro-photometric (i.e. astrometry and photometry) distances derived in this work and the results from the simulations lend confidence that our distances and stellar classifications are reliable even out to large distances.

8. No kinematic constraints. As can be seen in some panels in Figure 15, the predicted tangential velocities for some sources indicate clear mismodeling of the distances. Especially given the recent improved parallaxes and proper motions from Gaia EDR3, in future work we hope to incorporate additional kinematic constraints to resolve some of these degeneracies better.

Although the primary purpose of this catalog is to provide distance and reddening estimates, we also want to highlight possible deficiencies in secondary derived quantities explicitly,

namely photometric metallicities. In Figure 17, we show the metallicities recovery for a subset of $n \sim 5100$ stars from the H3 survey (Conroy et al. 2019). In brief, the H3 survey is a high-latitude ($|b| > 30^\circ$), high-resolution ($R = 32,000$) spectroscopic survey of the distant ($d \gtrsim 2$ kpc) Galaxy. Targets are selected purely on their Gaia parallax ($\varpi < 0.4\text{--}0.5$ mas), brightness ($15 < r_{\text{PS1}} < 18$), and accessibility to the 6.5 m MMT in Arizona, USA (decl. $> -20^\circ$). The survey measures radial velocities to 0.1 km s^{-1} precision, surface abundances ($[\text{Fe}/\text{H}]_{\text{surf}}$ and $[\alpha/\text{Fe}]_{\text{surf}}$) to 0.1 dex precision, and spectro-photometric distances to 10% precision using MINESWEEPER

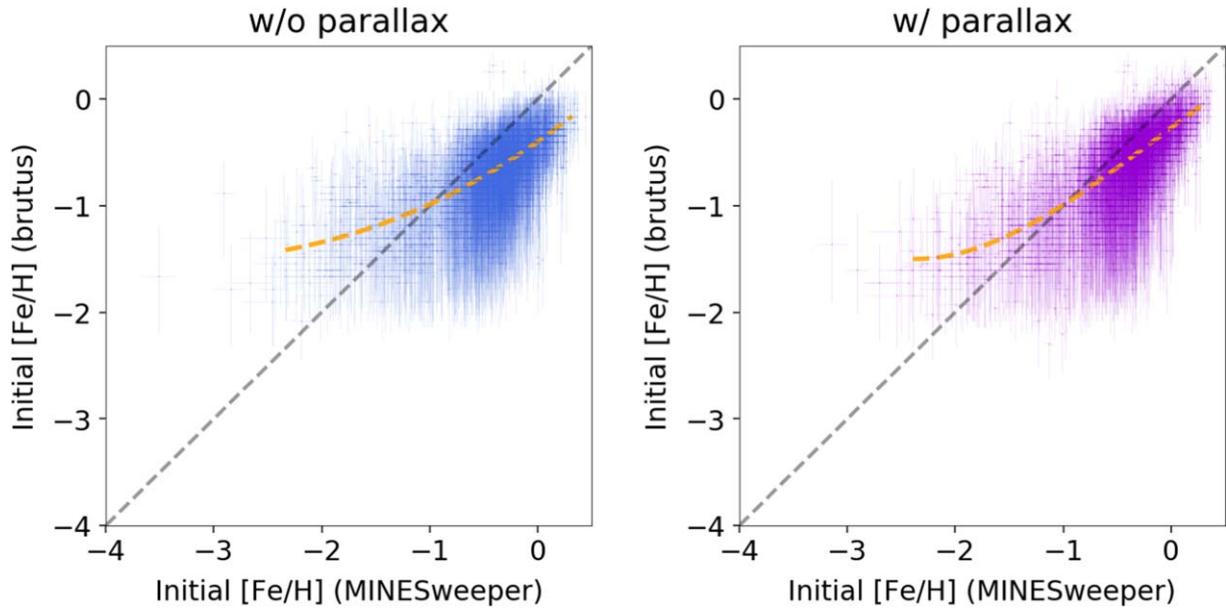


Figure 17. A comparison of the metallicities derived from BRUTUS (used in this work with photometry only) and MINESWEeper (including both photometry and spectroscopy) for a sample of $n \sim 5100$ objects shown from the H3 survey using similar isochrones and photometry without (left, blue) and with (right, purple) Gaia parallax constraints. 1σ errors from both sources are plotted for each point and the one-to-one relation is shown with a dashed gray line along with a sliding median (orange dashed line). In both cases, the metallicities derived from photometry only are found to be substantially biased, although they broadly follow the same trend as those derived using both spectra and photometry. As discussed in Section 5.6, this is due to a fundamental degeneracy where changes in the estimated reddening can be accounted for with corresponding changes to the underlying stellar properties, which lead to increased reliance on the metallicity prior from our Galactic model. These results suggest our metallicities, along with other derived quantities without good external constraints, should be used with caution.

(Cargile et al. 2020). In addition to being a representative, low-reddening subsample of sources, all derived quantities were estimated using the same underlying MIST isochrones (excluding the empirical corrections and photometric offsets derived in this work). This makes the estimated metallicity comparisons both independent (derived using different code bases and with/without spectra) while still remaining internally consistent (using similar photometry and underlying stellar models).

As seen in Figure 17, the parameter recovery is substantially biased and prior dominated, even with reasonable S/N photometry and parallax measurements. As discussed at length in Speagle et al. (2024), this is because there is an intrinsic degeneracy in reddened stellar colors that only can be broken if the dust attenuation is known to high precision. Without this, it is possible to shift the brightness and colors of stars by adjusting dust attenuation (which affects the observed star colors) along with metallicity and other stellar properties (which affect the intrinsic star colors). As a result, estimates are generally dragged toward our prior means for both the thin/thick disk and halo populations, which show up as a bias for both high- and low-metallicity objects. Since distance estimates do not break these color degeneracies, this effect is present when modeling objects with or without parallax constraints, emphasizing the prior is actually dominating much of the inference. While the internal tests and comparisons presented here imply this bias does not substantially impact our distance estimates, it does mean that our metallicity estimates, and other prior-dominated derived quantities, should be used with caution in any downstream analyses.

6. Conclusion

As large-area surveys such as SDSS (York et al. 2000), Pan-STARRS (Chambers et al. 2016), Gaia (Gaia Collaboration

et al. 2016), and the Legacy Survey of Space and Time (Ivezic et al. 2019) continue to or soon promise to provide measurements of billions of stars, the challenge of transforming observations of the projected 2D positions of sources on the sky into full 3D maps becomes ever more pressing when trying to study the Milky Way. In this work, we presented results applying BRUTUS (Speagle et al. 2024)—a public, open-source PYTHON package that uses a combination of statistical approaches to infer stellar properties, distances, and extinctions for sources using photometry and astrometry—to a catalog of 170 million sources (Augustus) at high Galactic latitudes ($|b| > 10$ mag) down to $r < 20$ mag with data from Pan-STARRS, Gaia, 2MASS, UKIDSS, and unWISE.

We find 125 million objects (Augustus-Gold) have good fits and reliable posteriors with estimated statistical distance uncertainties $\sim 3\%-5\%$ at $r = 14$ mag to $\sim 8\%-10\%$ at $r = 20$ mag. We show that our results are able to predict the “empirical,” dereddened Gaia CMD based on astro-photometric (i.e. astrometry and photometry) modeling in other bands, and that the derived stellar parameters are in excellent agreement with similar results derived in A19. We then illustrate the quality of the data by highlighting its ability to recover large- and small-scale Galactic substructure such as the Monoceros Ring at $d \sim 10$ kpc and the Sagittarius Stream at $d \sim 25$ kpc in density, metallicity, and kinematics relative to expectations from the underlying Galactic prior. Finally, we present an interactive visualization (ALLSKY) that is able to highlight limited 5D distance and tangential velocity structure present in our data.

Catalogs summarizing our results are publicly available at the Harvard Dataverse³⁰ and summarized in the Appendix. Overall, we hope that our results serve as a useful value-added

³⁰ https://dataverse.harvard.edu/dataverse/brutus_augustus

catalog that highlight the power of combined astro-photometric (i.e. astrometry and photometry) constraints to estimate stellar properties.

Acknowledgments

Contributions. The author list is divided up into three groups:

1. a list of primary authors who made direct contributions to the construction, computation, and analysis of the catalog (J.S.S. to B.D.J.),
2. an alphabetized secondary list of authors who made direct or indirect contributions to the project and/or data products (A.B. to E.F.S.), and
3. an alphabetized tertiary list of authors who provided useful feedback during the development process and/or on the paper itself (A.D. to I.A.Z.).

Personal. J.S.S. would like to thank Rebecca Bleich for her truly incredible support—mental, physical, emotional, and spiritual—during these difficult times. Without it, this paper (along with a great many other things) would likely never have seen the light of day. J.S.S. would also like to thank Jan Rybizki, Seth Gossage, and Nayantara Mudur for insightful discussions that improved the quality of this work and Blakesley Burkhart for providing feedback on earlier versions of the catalog.

Funding. J.S.S. and C.Z. were partially supported by the Harvard Data Science Initiative. H.M.K. acknowledges support from the DOE CSGF under grant No. DE-FG02-97ER25308. A.D. received support from the National Aeronautics and Space Administration (NASA) under Contract No. NNG16PJ26C issued through the WFIRST Science Investigation Teams Program. C.Z. and D.P.F. acknowledge support from NSF grant AST-1614941, “Exploring the Galaxy: 3-Dimensional Structure and Stellar Streams.” A.K.S. gratefully acknowledges support by a National Science Foundation Graduate Research Fellowship (DGE-1745303). Y.S.T. acknowledges financial support from the Australian Research Council through DECRA Fellowship DE220101520.

Data. The Pan-STARRS1 Surveys and the public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen’s University Belfast, the Center for Astrophysics, Harvard & Smithsonian, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

This publication makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by

the National Aeronautics and Space Administration and the National Science Foundation.

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

This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, and NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology. WISE and NEOWISE are funded by the National Aeronautics and Space Administration.

Computation. The bulk of the computation in this paper was done on the Cannon cluster, which is supported by the Research Computing Group in the FAS Division of Science at Harvard University.

Software: Astropy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018), NumPy (van der Walt et al. 2011), SciPy (Virtanen et al. 2020), matplotlib (Hunter 2007), healpy (Górski et al. 2005; Zonca et al. 2019), Gala (Price-Whelan 2017), galpy (Bovy 2015), corner (Foreman-Mackey 2016).

Appendix Data Products

The output Augustus stellar parameter catalogs can be found online through the Harvard Dataverse (see footnote 32). Two types of data products are made available:

1. a “point” catalog that contains information about each object along with a statistics summary describing the results (doi:[10.7910/DVN/WYMSXV](https://doi.org/10.7910/DVN/WYMSXV)) and
2. a “samples” catalog that contains 25 random posterior samples for each object (doi:[10.7910/DVN/530UYQ](https://doi.org/10.7910/DVN/530UYQ)).

A summary of the column names, the data format, and a brief description of each catalog can be found in Tables 3 and 4. Note that the “samples” catalogs is strictly meant to be supplementary to the “point” catalog and is matched to the latter row-wise.³¹ Catalogs are made available for sources modeled using all bands as well as excluding UKIDSS data for sources that have them (which have the _noukidss suffix). The _noukidss data products are available for download at the same Harvard Dataverse (see footnote 32) repository as their UKIDSS-included counterpart files. Additional information on the columns provided in these catalogs are described below.

Information on sources is provided through their corresponding Pan-STARRS ID (PS_ID) and Gaia DR2 ID (GAIA_ID) as well as by their 2D coordinates, in R.A. and decl. (SKY_COORDS) as well as Galactic longitude and latitude (GAL_COORDS). In addition, we also include astrometric measurements from Gaia including parallaxes (PARALLAX,

³¹ As described in Section 4, BRUTUS v0.7.5 contained a bug (fixed in more recent versions of the code) that used the wrong sign when sampling from correlations between A_V and R_V with distance d . While this has a negligible impact on the quality of the point catalog outside of the provided random draw, it does affect quantities computed directly from the “samples” catalog, which jointly depend on d and A_V or R_V (e.g., reddened photometry).

Table 3

Summary of the Augustus “Point” Catalog That Includes Object Information, a Statistics Summary, and a Description of the Results from the BRUTUS Fits

Name	Data Format	Description
<i>Object Information</i>		
PS_ID	64 bit uint	Pan-STARRS object ID
GAIA_ID	64 bit uint	Gaia DR2 object ID
SKY_COORDS	64 bit float (x2)	Sky coordinates (α, δ) in degrees
GAL_COORDS	64 bit float (x2)	Galactic coordinates (ℓ, b) in degrees
PARALLAX	32 bit float	Parallax from Gaia DR2 in milliarcseconds
PARALLAX_ERROR	32 bit float	Parallax error from Gaia DR2 in milliarcseconds
PROPER_MOTION	32 bit float (x2)	Proper motion in sky coordinates from Gaia DR2 in units of mas yr ⁻¹
PROPER_MOTION_ERROR	32 bit float (x2)	Proper motion error in sky coordinates from Gaia DR2 in units of mas yr ⁻¹
MAGNITUDES	32 bit float (x16)	Magnitudes from Gaia DR2, Pan-STARRS, 2MASS, UKIDSS, and unWISE
MAGNITUDES_ERROR	32 bit float (x16)	Magnitude errors from Gaia DR2, Pan-STARRS, 2MASS, UKIDSS, and unWISE
<i>Fit Information</i>		
LOG_EVID	32 bit float	Log evidence (base e) from the models used in the fit
BEST_CHI2	32 bit float	Best-fit χ^2 value (photometry and parallax) from the models used in the fit
NBANDS_IN_FIT	8 bit uint	Number of bands (photometry and parallax) included in the fit
FLAG_FIT	1 bit bool	Whether there was an issue with the fit (TRUE = yes)
FLAG_GRID	1 bit bool	Whether the posterior hits the edge of the grid (TRUE = yes)
<i>Stellar Properties</i>		
PROB_GIANT	16 bit float	Probability that $\log g < 3.5$ from the models used in the fit
INIT_MASS	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of M_{init} in units of M_{\odot}
INIT_FEH	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of [Fe/H] _{init}
EEP	16 bit int (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of EEP
LOG10_AGE	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of $\log t_{\text{age}}$ in years
LOG10_TEMP_EFF	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of $\log T_{\text{eff}}$ in Kelvin
LOG10_LBOL	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of $\log L_{\text{bol}}$ in units of L_{\odot}
LOG10_SURF_GRAV	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of $\log g$ in cgs units
DISTANCE	32 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of d in kiloparsecs
A_V	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of A_V in mag
R_V	16 bit float (x4)	2.5th, 50th, and 97.5th percentiles and a random sample of R_V

Note. See the [Appendix](#) for additional details. The table is available for download at doi:[10.7910/DVN/WYMSXV](https://doi.org/10.7910/DVN/WYMSXV).

Table 4

Summary of the Augustus “Samples” Catalog That Includes Random Posterior Samples from the BRUTUS Fits

Name	Data Format	Description
SAMPLES_MODEL_IDX	32 bit int (x25)	25 posterior samples of the model index in the input grid
SAMPLES_DISTANCE	32 bit float (x25)	25 posterior samples of d in kiloparsecs
SAMPLES_A_V	16 bit float (x25)	25 posterior samples of A_V in mag
SAMPLES_R_V	16 bit float (x25)	25 posterior samples of R_V in mag

Note. See the [Appendix](#) for additional details. The table is available for download at doi:[10.7910/DVN/530UYQ](https://doi.org/10.7910/DVN/530UYQ).

PARALLAX_ERROR) and proper motions (PROPER_MOTION, PROPER_MOTION_ERROR) as well as multiband photometry from Gaia, Pan-STARRS, 2MASS, UKIDSS, and unWISE (MAGNITUDES, MAGNITUDES_ERROR). Note that the parallaxes and magnitudes contain none of the offsets or additional systematic corrections/errors described in Section 2 outside of the 0.005 mag photometric error floor.

Information on the overall quality of the fit can be assessed in a few ways. One metric is the log evidence (LOG_EVID), defined as:

$$\ln \mathcal{Z} \equiv \ln \left(\int \mathcal{L}_{\text{phot}}(\boldsymbol{\theta}, \boldsymbol{\phi}) \mathcal{L}_{\text{astr}}(\boldsymbol{\phi}) \pi(\boldsymbol{\theta}, \boldsymbol{\phi}) d\boldsymbol{\theta} d\boldsymbol{\phi} \right). \quad (\text{A1})$$

We estimate this for each object by summing over the final (weighted) subset of models before applying the posterior resampling scheme described in Speagle et al. (2024). This provides information on the overall quality of the fit across all the models including the influence of the prior.

Another metric is simply the best-fit χ^2_{best} (BEST_CHI2) from all the models before applying the posterior resampling scheme. This provides information on the quality of the best possible fit ignoring the impact of the prior, thereby serving as a useful supplement to the log evidence.

Combined with the number of bands b used in the fit (NBANDS_IN_FIT), we use this information to derive a flag for rejecting sources that fail to achieve even a single

reasonable fit:

$$\text{FLAG_FIT} = \begin{cases} \text{TRUE} & \text{if } P(\chi^2 > \chi^2_{\text{best}} | b - 3) < 10^{-3}, \\ \text{FALSE} & \text{otherwise,} \end{cases} \quad (\text{A2})$$

where $P(\chi^2 > \chi^2_{\text{best}} | b - 3)$ is the probability of observing a χ^2 value larger than χ^2_{best} assuming $b - 3$ degrees of freedom. Note that we use $b - 3$ rather than b due to the fact that BRUTUS “optimizes” over three parameters (d , A_V , and R_V) before the posterior weighting and resampling step.

For every parameter we compute the 2.5th, 50th, and 97.5th percentiles (i.e., the median and the 2σ errors) for each parameter by rank ordering the final set of $n = 250$ posterior samples. We choose to report 2σ rather than 1σ errors here since they better reflect intuitive understanding of the uncertainty (i.e., objects are “unlikely” to be outside the errors bars) and better highlight possible degeneracies in the fits (e.g., between dwarf and giant solutions) when they occur. We provide these percentiles along with a random sample taken from the posterior for each parameter of the object in our model (see Table 3 for a full list). Since all the parameters from this random sample are correlated, it can be useful in certain contexts.

We use these percentiles to define a second flag `FLAG_GRID` that we set to `TRUE` if any of the 2.5th or 97.5th percentiles for each parameter that defines the grid of models (M_{init} , $[\text{Fe}/\text{H}]_{\text{init}}$, and EEP) are equal to the minimum or maximum possible value of that parameter, respectively, and `FALSE` otherwise. This flags posteriors that may be biased due to the hard edges present in our input model grid. This mainly flags sources with lower initial masses since our model grid only goes down to $M_{\text{init}} = 0.5 M_{\odot}$.

Finally, as part of the “point” catalog we also provide the probability `PROB_GIANT` that a source is a giant, which we define as:

$$\text{PROB_GIANT} \equiv P(\log g < 3.5), \quad (\text{A3})$$

following the definition used in the H3 survey (Conroy et al. 2019). We estimate this using the final set of posterior samples, which gives us a resolution of $1/n_{\text{samp}} = 1/250 = 0.4\%$ in probability.

In the “samples” catalog, we provide 25 samples from the posterior for the distance d (`SAMPLES_DISTANCE`), extinction A_V (`SAMPLES_A_V`), differential reddening R_V (`SAMPLES_R_V`), and “model index” (`SAMPLES_MODEL_IDX`). The model index can be used to grab the corresponding models from the input parameter grid, which can then be used to construct output predictions for associated quantities.

An example showing how to use these samples within BRUTUS is shown below:

```
import numpy as np
import h5py
from brutus import utils as butils
from brutus.filters import gaia, ps, tmass, ukidss, wise

# grab quality flags
cat = h5py.File('point_cat.h5', mode = 'r') # load h5 file
flag_fit, flag_grid = cat['FLAG_FIT'][100], cat['FLAG_GRID'][100] # first
100 elements
good = np.where(~flag_fit & ~flag_grid)[0] # no flags (good fits, good
posterior)
```

(Continued)

```
# load samples catalog
samples = h5py.File('samples_cat.h5', mode = 'r') # load h5 file
samples_idx = samples['SAMPLES_MODEL_IDX'][100][good] # first 100
elements + no flags

# load MIST grid
flts = gaia + ps[-2] + tmass + ukidss + wise[-2] # define filterset
mags, labels, _ = butils.load_models('grid_mist_v8.h5', filters = flts) #
read file

# get effective temperatures of corresponding models
logt = labels['logt'][samples_idx] # log(Teff)
teff = 10**logt # convert from log to linear

# compute percentiles (median, +/- 1 sigma, +/- 2 sigma)
teff_vals = np.percentile(teff, [2.5, 16, 50, 84, 97.5], axis=1)

# compute mean and standard deviation of predicted intrinsic Gaia G mag-
nitude at 1 kpc
G = mags[:, 0, 0][samples_idx]
G_mean, G_std = np.mean(G, axis=1), np.std(G, axis=1)
```

ORCID iDs

Joshua S. Speagle (沈佳士)  <https://orcid.org/0000-0003-2573-9832>
Catherine Zucker  <https://orcid.org/0000-0002-2250-730X>
Ana Bonaca  <https://orcid.org/0000-0002-7846-9787>
Phillip A. Cargile  <https://orcid.org/0000-0002-1617-8917>
Benjamin D. Johnson  <https://orcid.org/0000-0002-9280-7594>
Angus Beane  <https://orcid.org/0000-0002-8658-1453>
Charlie Conroy  <https://orcid.org/0000-0002-1590-8551>
Douglas P. Finkbeiner  <https://orcid.org/0000-0003-2808-275X>
Gregory M. Green  <https://orcid.org/0000-0001-5417-2260>
Harshil M. Kamdar  <https://orcid.org/0000-0001-5625-5342>
Rohan Naidu  <https://orcid.org/0000-0003-3997-5705>
Hans-Walter Rix  <https://orcid.org/0000-0003-4996-9069>
Edward F. Schlafly  <https://orcid.org/0000-0002-3569-7421>
Aaron Dotter  <https://orcid.org/0000-0002-4442-5700>
Gwendolyn Eadie  <https://orcid.org/0000-0003-3734-8177>
Daniel J. Eisenstein  <https://orcid.org/0000-0002-2929-3121>
Alyssa A. Goodman  <https://orcid.org/0000-0003-1312-0477>
Jiwon Jesse Han  <https://orcid.org/0000-0002-6800-5778>
Andrew K. Saydjarji  <https://orcid.org/0000-0002-6561-9002>
Yuan-Sen Ting (丁源森)  <https://orcid.org/0000-0001-5082-9536>
Ioana A. Zelko  <https://orcid.org/0000-0002-7588-976X>

References

Anders, F., Khalatyan, A., Chiappini, C., et al. 2019, [A&A](#), **628**, A94
Antoja, T., Helmi, A., Romero-Gómez, M., et al. 2018, [Natur](#), **561**, 360
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, [A&A](#), **558**, A33
Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Demleitner, M., & Andrae, R. 2021, [AJ](#), **161**, 147
Bailer-Jones, C. A. L., Rybizki, J., Fouesneau, M., Mantelet, G., & Andrae, R. 2018, [AJ](#), **156**, 58

Belokurov, V., Erkal, D., Evans, N. W., Koposov, S. E., & Deason, A. J. 2018, *MNRAS*, **478**, 611

Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, **496**, 1922

Bland-Hawthorn, J., & Gerhard, O. 2016, *ARA&A*, **54**, 529

Bonaca, A., & Hogg, D. W. 2018, *ApJ*, **867**, 101

Bovy, J. 2015, *ApJS*, **216**, 29

Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, **427**, 127

Cargile, P. A., Conroy, C., Johnson, B. D., et al. 2020, *ApJ*, **900**, 28

Casali, M., Adamson, A., Alves de Oliveira, C., et al. 2007, *A&A*, **467**, 777

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

Chen, Y., Girardi, L., Bressan, A., et al. 2014, *MNRAS*, **444**, 2525

Choi, J., Dotter, A., Conroy, C., et al. 2016, *ApJ*, **823**, 102

Conroy, C., Bonaca, A., Cargile, P., et al. 2019, *ApJ*, **883**, 107

Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, yCat, **II/328**

Dotter, A. 2016, *ApJS*, **222**, 8

Dye, S., Warren, S. J., Hambley, N. C., et al. 2006, *MNRAS*, **372**, 1227

Fitzpatrick, E. L. 2004, in ASP Conf. Ser. 309, Astrophysics of Dust, ed. A. N. Witt, G. C. Clayton, & B. T. Draine (San Francisco, CA: ASP), 33

Foreman-Mackey, D. 2016, *JOSS*, **1**, 24

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2016, *A&A*, **595**, A2

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1

Gómez, F. A., Minchev, I., O’Shea, B. W., et al. 2013, *MNRAS*, **429**, 159

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759

Green, G. M., Rix, H.-W., Tschesche, L., et al. 2021, *ApJ*, **907**, 57

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

Hambley, N. C., Collins, R. S., Cross, N. J. G., et al. 2008, *MNRAS*, **384**, 637

Helmi, A. 2020, *ARA&A*, **58**, 205

Helmi, A., Babusiaux, C., Koppelman, H. H., et al. 2018, *Natur*, **563**, 85

Hewett, P. C., Warren, S. J., Leggett, S. K., & Hodgkin, S. T. 2006, *MNRAS*, **367**, 454

Hodgkin, S. T., Irwin, M. J., Hewett, P. C., & Warren, S. J. 2009, *MNRAS*, **394**, 675

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

Ivezic, Z., Tyson, J. A., Abel, B., et al. 2019, *ApJ*, **873**, 111

Johnston, K. V., Zhao, H., Spergel, D. N., & Hernquist, L. 1999, *ApJL*, **512**, L109

Juric, M. 2011, AAS Meeting Abstracts, **217**, 433.19

Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008, *ApJ*, **673**, 864

Khan, S., Miglio, A., Mossner, B., et al. 2019, *A&A*, **628**, A35

Koppelman, H., Helmi, A., & Veljanoski, J. 2018, *ApJL*, **860**, L11

Lallement, R., Babusiaux, C., Vergely, J. L., et al. 2019, *A&A*, **625**, A135

Laporte, C. F. P., Belokurov, V., Koposov, S. E., Smith, M. C., & Hill, V. 2020, *MNRAS*, **492**, L61

Laporte, C. F. P., Gómez, F. A., Besla, G., Johnston, K. V., & Garavito-Camargo, N. 2018a, *MNRAS*, **473**, 1218

Laporte, C. F. P., Johnston, K. V., Gómez, F. A., Garavito-Camargo, N., & Besla, G. 2018b, *MNRAS*, **481**, 286

Law, D. R., & Majewski, S. R. 2010, *ApJ*, **714**, 229

Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, *MNRAS*, **379**, 1599

Leike, R. H., & Enßlin, T. A. 2019, *A&A*, **631**, A32

Leung, H. W., & Bovy, J. 2019a, *MNRAS*, **483**, 3255

Leung, H. W., & Bovy, J. 2019b, *MNRAS*, **489**, 2079

Lindgren, L., Hernández, J., Bombrun, A., et al. 2018, *A&A*, **616**, A2

Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., et al. 2020, *ApJS*, **251**, 6

Maíz Apellániz, J., & Weiler, M. 2018, *A&A*, **619**, A180

Naidu, R. P., Conroy, C., Bonaca, A., et al. 2021, *ApJ*, **923**, 92

Ness, M., Hogg, D. W., Rix, H. W., Ho, A. Y. Q., & Zasowski, G. 2015, *ApJ*, **808**, 16

Newberg, H. J., Yanny, B., Rockosi, C., et al. 2002, *ApJ*, **569**, 245

Oke, J. B., & Gunn, J. E. 1983, *ApJ*, **266**, 713

Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJS*, **192**, 3

Paxton, B., Cantiello, M., Arras, P., et al. 2013, *ApJS*, **208**, 4

Paxton, B., Marchant, P., Schwab, J., et al. 2015, *ApJS*, **220**, 15

Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, *ApJS*, **234**, 34

Paxton, B., Smolec, R., Schwab, J., et al. 2019, *ApJS*, **243**, 10

Price-Whelan, A. M. 2017, *JOSS*, **2**, 388

Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123

Purcell, C. W., Bullock, J. S., Tollerud, E. J., Rocha, M., & Chakrabarti, S. 2011, *Natur*, **477**, 301

Queiroz, A. B. A., Anders, F., Santiago, B. X., et al. 2018, *MNRAS*, **476**, 2556

Rezaei Kh., S., Bailer-Jones, C. A. L., Hogg, D. W., & Schultheis, M. 2018, *A&A*, **618**, A168

Rix, H.-W., & Bovy, J. 2013, *A&ARv*, **21**, 61

Santiago, B. X., Brauer, D. E., Anders, F., et al. 2016, *A&A*, **585**, A42

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

Schönrich, R., McMillan, P., & Eyer, L. 2019, *MNRAS*, **487**, 3568

Sellwood, J. A. 2014, *RvMP*, **86**, 1

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163

Speagle, J., & Zucker, C. 2020, *joshspeagle/brutus: brutus*, v0.7.5, Zenodo

Speagle, J. S., Zucker, C., Beane, A., et al. 2024, *ApJ*, submitted

Tang, J., Bressan, A., Rosenfield, P., et al. 2014, *MNRAS*, **445**, 4287

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22

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

Warren, S. J., Cross, N. J. G., Dye, S., et al. 2007a, arXiv:astro-ph/0703037

Warren, S. J., Cross, N. J. G., Dye, S., et al. 2007b, *MNRAS*, **375**, 213

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868

Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, *ApJS*, **245**, 34

Xue, X.-X., Rix, H.-W., Ma, Z., et al. 2015, *ApJ*, **809**, 144

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

Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298