Durham Research Online

Publisher's copyright statement:

Additional information:

A THREE-DIMENSIONAL MAP OF MILKY WAY DUST

Gregory M. Green[1], Edward F. Schlafly[2], Douglas P. Finkbeiner[1], Hans-Walter Rix[2], Nicolas Martin[2,3],
William Burgett[4], Peter W. Draper[5], Heather Flewelling[6], Klaus Hodapp[6], Nicholas Kaiser[6], Rolf Peter Kudritzki[6],
Eugene Magnier[6], Nigel Metcalfe[5], Paul Price[7], John Tonry[6], and Richard Wainscoat[6]

[1] Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA; ggreen@cfa.harvard.edu
[2] Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
[3] Observatoire astronomique de Strasbourg, Université de Strasbourg, CNRS, UMR 7550, 11 rue de l'Université, F-67000 Strasbourg, France
[4] GMTO Corporation, 251 S. Lake Ave., Suite 300, Pasadena, CA 91101, USA
[5] Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK
[6] Institute for Astronomy, University of Hawaii at Manoa, 2680 Woodlawn Dr., Honolulu, HI 96822, USA
[7] Princeton University Observatory, 4 Ivy Lane, Peyton Hall, Princeton University, Princeton, NJ 08544, USA

ABSTRACT

We present a three-dimensional map of interstellar dust reddening, covering three-quarters of the sky out to a distance of several kiloparsecs, based on Pan-STARRS 1 (PS1) and 2MASS photometry. The map reveals a wealth of detailed structure, from filaments to large cloud complexes. The map has a hybrid angular resolution, with most of the map at an angular resolution of $3\farcm4$–$13\farcm7$, and a maximum distance resolution of $\sim25\%$. The three-dimensional distribution of dust is determined in a fully probabilistic framework, yielding the uncertainty in the reddening distribution along each line of sight, as well as stellar distances, reddenings, and classifications for 800 million stars detected by PS1. We demonstrate the consistency of our reddening estimates with those of two-dimensional emission-based maps of dust reddening. In particular, we find agreement with the Planck $\tau_{353\mathrm{GHz}}$-based reddening map to within 0.05 mag in $E(B-V)$ to a depth of 0.5 mag, and explore systematics at reddenings less than $E(B-V) \approx 0.08$ mag. We validate our per-star reddening estimates by comparison with reddening estimates for stars with both Sloan Digital Sky Survey photometry and Sloan Extension for Galactic Understanding and Exploration spectral classifications, finding per-star agreement to within 0.1 mag out to a stellar $E(B-V)$ of 1 mag. We compare our map to two existing three-dimensional dust maps, by Marshall et al. and Lallement et al., demonstrating our finer angular resolution, and better distance resolution compared to the former within ~3 kpc. The map can be queried or downloaded at http://argonaut.skymaps.info. We expect the three-dimensional reddening map presented here to find a wide range of uses, among them correcting for reddening and extinction for objects embedded in the plane of the Galaxy, studies of Galactic structure, calibration of future emission-based dust maps, and determining distances to objects of known reddening.

Key words: dust, extinction – Galaxy: structure – methods: statistical

1. INTRODUCTION

The Milky Way is the only galaxy that can be observed in such close detail, yet most of the plane of the Galaxy is veiled by dust. While dust makes up just 1% of the mass of the interstellar medium (ISM), it absorbs all starlight on the order of 30% (Draine 2003). Interstellar dust absorbs and scatters light in the ultraviolet (UV), optical, and near-infrared (NIR), re-radiating it in the mid- to far-infrared (MIR and FIR, respectively). Understanding the spatial distribution of dust is therefore critical for UV, optical, and NIR astronomy, where dust is an extinguishing foreground, to extragalactic astronomy and cosmology, where it is a radiating foreground, as well as to star formation, where the dust itself is a primary object of study. Detailed studies of stellar populations and substructures in the Galaxy require accurate corrections for extinction and reddening due to dust. The plane of the Milky Way, containing the majority of the Galaxy's stellar content, is also the most heavily extinguished region. But dust is not only a nuisance to astronomers. Dust traces the interstellar medium, and a clearer picture of the spatial distribution of dust would aid in understanding the processes that shape the Galaxy, from star formation to how feedback from supernovae and stellar winds shape our Galaxy's ISM.

Dust can be mapped either through its extinction or through its emission. The spectral energy distribution and amplitude of FIR dust emission is most sensitive to dust column density, temperature, and grain-size distribution. By modeling these properties across the sky, a map of dust column density may be obtained, which can then be converted to extinction or reddening by calibration against astrophysical extinction or reddening standards. Such emission-based methods recover the angular distribution of dust, but not its distribution in distance. Burstein & Heiles (1982) produced an all-sky map of dust reddening based on H I emission and galaxy counts, assuming the gas and dust to be uniformly mixed. Schlegel et al. (1998) (hereafter "SFD") produced a widely used all-sky reddening map, using DIRBE and *IRAS* maps of FIR emission to model dust column density and temperature. More recently, the Planck Collaboration has released all-sky reddening maps based on a similar emission-modeling technique (Planck Collaboration et al. 2014).

A second class of dust maps is based on extinction or reddening estimates of sources distributed across the sky. Lada et al. (1994, NICE) compared the average $H - K$ colors and number counts of stars in target and control fields to map dust reddening in two dimensions. Lombardi & Alves (2001, NICER) and Lombardi (2009, NICEST) extended on this algorithm, which has now been applied to 2MASS photometry

to obtain reddening maps of a host of cloud complexes (Lombardi et al. 2006, 2008, 2010, 2011; Alves et al. 2014). Froebrich applied a similar *color excess* method to 2MASS photometry to produce a two-dimensional (2D), all-sky reddening map (Rowles & Froebrich 2009). Schlafly et al. (2010) used the location of main-sequence turn-off stars in color–color space to simultaneously map dust reddening and test different reddening laws. Peek & Graves (2010) used passively evolving galaxies as a standard color source in order to map reddening and correct the SFD map.

Because stars are distributed throughout the Galaxy, they can be used to trace the dust distribution in three dimensions. Most methods relying on this principle group stars into separate sightlines, and then determine stellar reddening as a function of distance along each sightline. The challenge with such methods is to simultaneously determine stellar type (and thus the intrisic stellar colors and luminosity), distance, and reddening on the basis of photometry alone. Marshall et al. (2006) developed a method that iteratively improves distance and reddening estimates to post-main sequence stars, updating the dust column in each distance bin with each iteration so that the intrinsic stellar colors match those predicted by the Besançon model of the Galactic stellar population (Robin et al. 2003). Marshall et al. (2006) applied this method to 2MASS photometry of the Galactic plane, producing a three-dimensional (3D) reddening map out to a distance of several kiloparsecs in the region $|b| < 10°$, $-100° \leqslant \ell \leqslant 100°$. A number of groups have pursued methods that determine maximum-likelihood parameters for the stars along each sightline, combining the estimates in each sightline into a distance-reddening relation. Berry et al. (2011) applied such a method to Sloan Digital Sky Survey (SDSS) photometry, producing a 3D reddening map and measuring variation in the dust extinction spectrum across the SDSS footprint. Chen et al. (2014) used optical XSTPS-GAC, NIR 2MASS and MIR *WISE* photometry to produce a 3D reddening map of the Galactic anticenter. Sale (2012) developed a probabilistic framework to simultaneously infer stellar parameters and the dust extinction distribution along each line of sight, and Sale et al. (2014) applied this method to IPHAS photometry of the northern Galactic plane. Hanson & Bailer-Jones (2014) developed a similar probabilistic framework for inferring stellar types, distances, and dust foreground properties, and applied this method to SDSS and UKIRT Infrared Deep Sky Survey photometry in regions of the high-Galactic-latitude sky. Lallement et al. (2013) took a somewhat different approach, using ~23,000 stellar parallaxes and reddening estimates from a number of sources to infer the 3D distribution of dust opacity out to a distance of 800–1000 pc in the plane of the Galaxy, and ~300 pc out of the plane. Green et al. (2014) developed a probabilistic method for determining dust reddening in 3D from stellar photometry, and Schlafly et al. (2014c) presented a map of dust reddening integrated to 4.5 kpc, covering three quarters of the sky, based on applying this method to optical Pan-STARRS 1 (PS1) stellar photometry. Using this same 3D mapping technique, Schlafly et al. (2014a) found that the Orion molecular complex appears to form part of a larger bubble structure. Schlafly et al. (2014b) determined the distance to a large number of molecular clouds by applying a related method to PS1 stellar photometry.

In this paper, we present a three-dimensional map of dust reddening over three quarters of the sky. Our map covers the Northern sky down to $\delta \approx -30°$, has an adaptive resolution of typical angular scale $3'.4$–$13'.7$, logarithmically spaced distance bins with a width of ~25%, and extends to $E(B-V) \sim 1.5$–2 mag. Our map is based on high-quality PS1 photometry for 800 million stars, ~200 million of which have matched 2MASS photometry. The reddening along each sightline is determined in a fully probabilistic manner, according to the method developed in Green et al. (2014). We strongly encourage readers who wish to understand our method for inferring the three-dimensional distribution of dust reddening to read that paper, as the present paper focuses primarily on results obtained using that method.

This paper is organized as follows. In Section 2, we briefly describe the method developed in Green et al. (2014), and describe the priors we place on the three-dimensional distribution of dust reddening. In Section 3, we describe the data sources which go into our map. In Section 4, we discuss results from the map, and in Section 5, we compare our map to existing 3D and 2D dust maps. We describe how the map can be accessed in Section 6.

2. METHOD

Our three-dimensional dust map is constructed pixel-by-pixel from independently determined line of sight reddening profiles. We group stars into sightlines with typical angular scales of $\sim 6'.8$, probabilistically inferring the reddening, distance, and stellar type of each star independently. We then assume that the distances and reddenings of the stars in a given line of sight lie along a single profile, with reddening increasing monotonically with distance. We sample from the posterior density of distance-reddening profiles, returning the uncertainty in the line of sight reddening in the pixel. By repeating this process in each angular pixel, we obtain the 3D distribution of dust throughout the PS1 survey volume (Kaiser et al. 2010). The details of our method are described in Green et al. (2014).

We have made three notable modifications to our method since the publication of Green et al. (2014). The first modification is the inclusion of stellar photometry from the 2MASS. We compile joint PS1+2MASS stellar templates, in a manner very similar to the process laid out for the PS1 templates used in Green et al. (2014). Appendix A contains details about the construction of our PS1+2MASS templates, while Appendix B details our estimate of the 2MASS survey depth.

The second modification to our method is to allow the reddening of individual stars to deviate slightly from the overall line of sight reddening profile in a pixel. Our basic model assumes that within an angular pixel, the dust density does not depend on angle, but is solely a function of distance. This is a good assumption at very fine angular scales, but it must obviously begin to break down at some angular scale, depending on the power spectrum of the dust density field. Because each star in an angular pixel is behind a slightly different column of dust, we allow the total column of dust in front of each star to vary slightly from the modeled column density. For a pixel with a small angular scale, in which the dust column should be more uniform across the pixel, the scatter of the stars around the modeled distance-reddening profile should be small. For pixels with larger angular scale, where angular variations in the dust column are correspondingly greater, one should allow the stars to deviate from the modeled distance-reddening profile to a greater extent. We therefore include an additional parameter for each star,

describing the fractional variation of the star from the modeled distance-reddening profile in its pixel. This modification to the model amounts to introducing a smoothing factor in the probability density of each star in distance-reddening space, and therefore contributes only negligibly to the overall computation time. We discuss this modification in more detail in Section 2.2.

A third modification we have made is to adopt the increased, bias-corrected stellar disk scale lengths and heights presented in Table 10 of Jurić et al. (2008) for our disk priors. We also adopt a less dense halo prior than Jurić et al. (2008), reducing the local halo density (relative to the local thin-disk stellar density) from $f_h = 0.0051$ to $f_h = 0.0006$, a similar value to that used by the Besançon model (Robin et al. 2003).

A brief summary of our model, including the changes introduced in Section 2.2, is given is Section 2.3.

Finally, we improve the convergence of our Markov Chain Monte Carlo (MCMC) sampling of the line of sight reddening profile by introducing a new type of proposal step, which we term the "swap" proposal. See Appendix C for details.

2.1. 3D Dust Priors

Green et al. (2014) describes how we use photometry of stars to determine the distribution of dust in the Galaxy. We begin by separating the sky into small sightlines. In a small region of the sky across which the dust column does not vary much angularly, more distant stars should be more heavily reddened than nearby stars, as all the stars lie along the same dust column. We therefore begin by calculating a probabilistic reddening and distance estimate for each star in the sightline. We then determine the amount of dust reddening in each distance bin along the sightline, under the constraint that the line of sight distance versus reddening profile has to be consistent with our distance and reddening inferences for all the stars in the sightline.

Green et al. (2014) states our map-making method generally, without choosing specific priors on the distribution of reddening in 3D. Here, we will briefly restate the formalism used in Green et al. (2014), and then define our 3D reddening priors.

Our goal of inferring the line of sight dust reddening on the basis of stellar photometry begins with inferring the distance and reddening of each star based on its photometry. Label each star in the line of sight by a number i, and denote the probability density for an individual star to lie at distance modulus, μ, and reddening, E, given photometry $\boldsymbol{m}$, by

$$p\left(\mu_i, E_i \middle| \boldsymbol{m}_i\right). \tag{1}$$

We precompute $p\left(\mu_i, E_i \middle| \boldsymbol{m}_i\right)$ using a kernel density estimate of MCMC samples. This technique, similar to that used in Hogg et al. (2010), is not guaranteed to converge to exactly the target density. It is possible, however, to substitute a mathematically correct, but somewhat more computationally expensive algorithm, *pseudo-marginal* MCMC, (Beaumont 2003; Andrieu & Roberts 2009), in which a new noisy estimate of a marginalized likelihood term is computed at each Markov chain step. However, we expect any possible bias introduced by precomputing Equation (1) once to be small due to the relatively large number of stars in each sightline. Compared to possible inaccuracies in our stellar model, for example, we expect the inexactness of our sampling method to be a small effect.

As explained in Green et al. (2014), we place a flat prior on E_i for each individual star when precomputing Equation (1), as the prior on reddening will come in when we combine information from all of the stars in a sightline. We now parameterize the line of sight reddening by a set of parameters, $\boldsymbol{\alpha}$, so that the cumulative reddening out to a distance modulus μ can be written as:

$$E(\mu; \boldsymbol{\alpha}). \tag{2}$$

The probability density of $\boldsymbol{\alpha}$ is then given by a product over line integrals through the individual stellar probability density functions, following $E(\mu; \boldsymbol{\alpha})$:

$$p(\boldsymbol{\alpha}|\{\boldsymbol{m}\}) \propto p(\boldsymbol{\alpha}) \prod_i \int d\mu_i \, p\left(\mu_i, E(\mu_i; \boldsymbol{\alpha}) \middle| \boldsymbol{m}_i\right). \tag{3}$$

We parameterize the reddening in each line of sight as a monotonically increasing, piecewise-linear function in distance modulus. Dividing up the line of sight into bins of equal width in distance modulus, we sample the increase in reddening in each bin, so that

$$\boldsymbol{\alpha} = \left\{ \Delta E_k \middle| k = 1, 2, \ldots, N_{\text{bins}} \right\}, \tag{4}$$

where k denotes the bin index, and N_{bins} is the number of distance bins. We place a log-normal prior on the reddening in each bin:

$$\Delta E_k = e^{\epsilon_k}, \text{ where } \epsilon_k \sim \mathcal{N}(\bar{\epsilon}_k, \sigma_\epsilon). \tag{5}$$

In each bin, we set $\bar{\epsilon}_k$ so that the mean reddening in each voxel matches what would be expected from the smooth disk component of Drimmel & Spergel (2001), up to a constant normalization that is the same for the entire map. In order to fully define the priors, there are therefore two global parameters that must be set:

1. $dE(B-V)/ds \mid_{s=0}$, the local normalization of the dust reddening per unit distance, and
2. σ_ϵ, the width of the log-normal prior on dust column density in each voxel.

We fit these two numbers so that our priors, in projection, produce two-dimensional dust maps with the same overall normalization and standard deviation as the SFD dust map. We find that $\sigma_\epsilon = 1.4$ and a local reddening per unit distance of $dE(B-V)/ds = 0.2 \text{ mag kpc}^{-1}$ roughly matches the mean and variance of the SFD dust map over large angular scales.

Figure 1 shows a random realization of our 3D dust priors, next to the SFD dust map. In each bin, we limit $\bar{\epsilon}$ to the range $-12 \leqslant \bar{\epsilon} \leqslant -4$. On the lower end, the limit helps our fit to converge by preventing the prior from becoming too stringent. On the upper end, the limit prevents our fit from inferring large amounts of dust in the absence of data. Regions where $\bar{\epsilon}$ has been limited to -4 are shaded in Figure 1. Our priors do not impose spatial correlations across lines of sight, and thus in the absence of data, produce cloud-free maps. The detailed cloud structures that emerge from our 3D dust modeling therefore derive entirely from the data, rather than from the priors.

It is worth noting that our priors diverge from what one should expect for the real distribution of dust in a number of ways. In order to render the problem of fitting the 3D distribution of dust tractable, we fit the dust column along each sightline separately. We know, however, that in real life, dust density has spatial correlations. Moreover, there is nothing
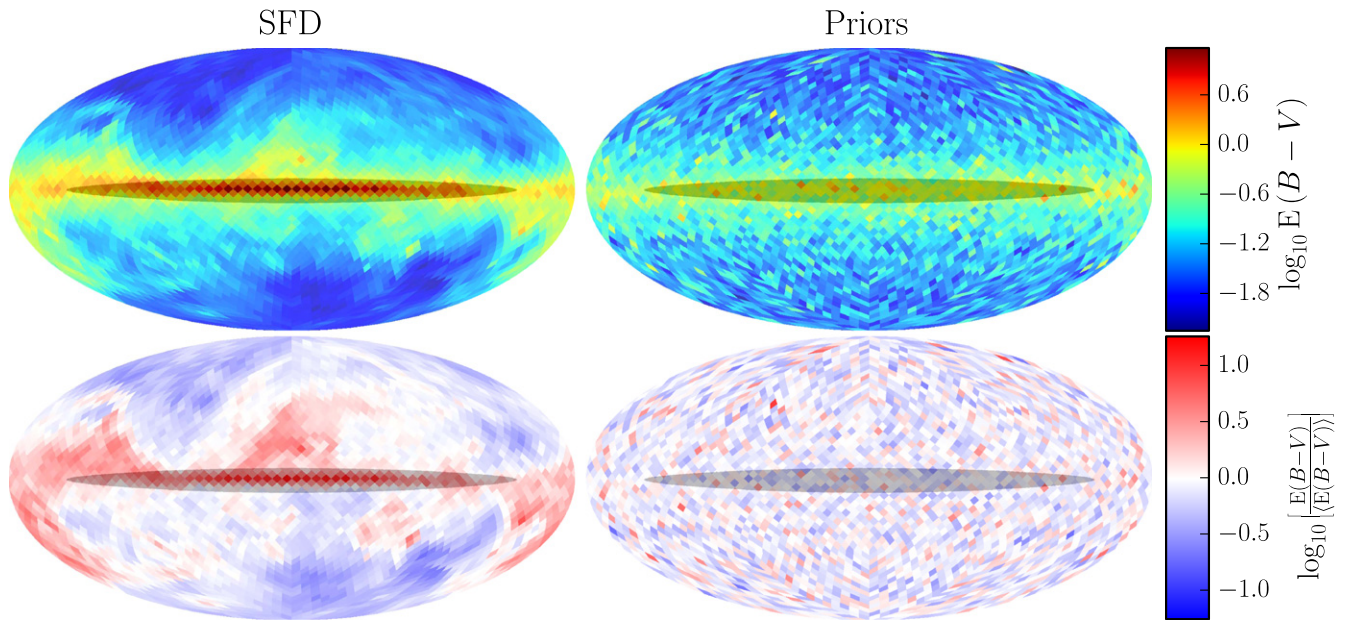
Figure 1. SFD dust map, compared to a random draw from the priors on the 3D dust reddening distribution, used in the construction of our map. The top two panels, from left to right, show the SFD reddening and the 2D reddening map that results from a draw from the priors, both on a log scale. The bottom two panels show the same maps, after dividing out the mean projected reddening in the priors. The priors are limited so that the mean expected reddening in any given distance bin does not exceed a pre-defined amount. This is done in order to avoid inferring large amounts of reddening in the absence of data. Regions that are affected by this clipping are shaded in gray. The priors do not include spatial correlations, as can be seen by comparing the bottom two panels.

special about the Sun's place in the Galaxy, but our dust map voxelizes the sky into pencil beams centered on the solar system. This of course makes practical sense, since angles are much easier to measure than distances in astronomy, but it is again an unreal feature of our voxelization. A more realistic model would treat the dust density as a continuous field, or voxelize the Galaxy in a way that treats the angular and radial directions equally, and would impose correlations between dust density in nearby points in space (see, e.g., Lallement et al. 2013; Sale & Magorrian 2014). This entails significantly more algorithmic and computational complexity than the method used here, and we defer such work to the future. Within the constraints of our present setup—independent sightlines with pencil-beam-like voxels—our priors attempt to reasonably trace the properties of the Galaxy, including the mean dust density in each voxel and the overall variance in dust column across the sky.

2.2. Scatter in Line of Sight Reddening Profiles

A basic assumption of our model, as described in Green et al. (2014), is that within a given pixel, the dust density varies only with distance, and not with angle. If the pixels are sufficiently small, this is a good assumption. We are limited, however, in how small we can make each pixel by the need to include enough stars in each pixel to probe the dust density at a range of distances. Increasing the angular resolution of the map decreases the number of stars in each pixel, effectively decreasing the distance resolution of the map. We have found that we obtain the best results for pixels containing a few hundred stars, and we vary the resolution of our pixels across the sky to obtain approximately the same number of stars in each pixel (see Figure 2).

Note that varying the pixel size across the sky in this way technically violates the principles of Bayesian inference. From a forward-modeling point of view, the distribution of dust
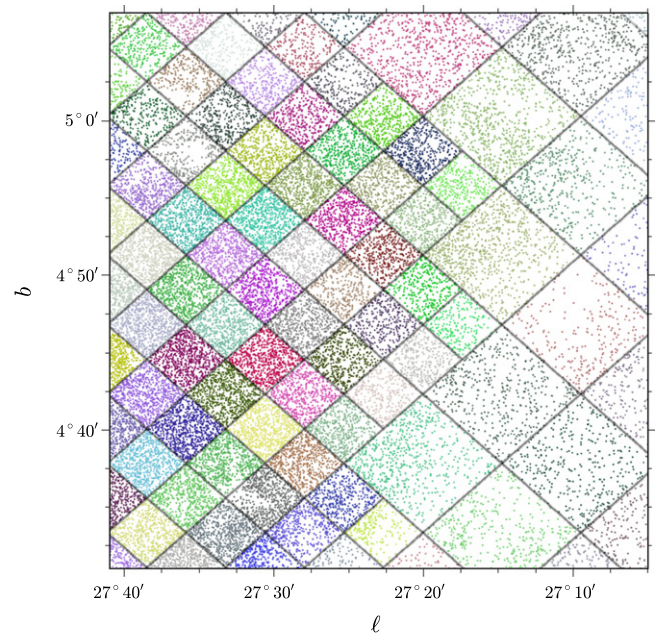

Figure 2. Assigning stars to different pixels, in which the line of sight dust profile will be fit independently. Here, each dot represents a point-source detection, and the dots are colored by the HEALPix pixel they are assigned to. The pixel scale varies in order to keep the number of stars per pixel roughly constant. Note that some pixels show significant extinction-induced variations in stellar density.

influences the number of stars that are observed in any given region of the sky. Sale (2015), for example, discusses how a catalog of stars observed in a survey can be described as a Poisson point process whose rate is determined by the distribution of stars in the Galaxy, the three-dimensional distribution of dust, and the survey selection function. Yet we are using the number of stars observed in each part of the sky to

determine how to pixelize the sky. Our pixelization is therefore set, in part, by an observable (the density of stars across the sky) that is a consequence of the model. Nevertheless, as we are treating each pixel independently, and we would like to keep pixels as small as possible without reducing the number of stars per pixel below a few hundred, this violation is difficult to avoid. We expect the impact of this violation to be small in most regions of the sky. In regions with large sub-pixel variation in dust reddening, however, our map may be biased toward lower reddenings, as we preferentially detect stars in regions of the pixel with lower reddening.

The typical resolution of the map is $6.'8 \times 6.'8$, corresponding to an $\mathtt{nside} = 512$ HEALPix pixelization (Gorski et al. 2005). At this resolution, there can still be significant power in the dust density spectrum below the pixel scale. This can pose problems for our method, especially in the vicinity of dense clouds and filamentary structures, where the sub-pixel angular variation is largest. In order to deal with sub-pixel angular variation in the dust density, we relax our assumption that all stars lie along the same dust column by allowing each star to deviate from the local "average" dust column by a small amount. The reddening of star i is parameterized as

$$E_i = (1 + \delta_i) E(\mu_i; \boldsymbol{\alpha}), \tag{6}$$

where δ_i is the fractional offset of the star from the local dust column, $E(\mu_i; \boldsymbol{\alpha})$.

In effect, our model is therefore that within each HEALPix pixel, the reddening is a white noise process, with a mean that increases piecewise linearly with distance. Each star samples this white noise process at a particular distance and angular position within the pixel. The parameter δ_i is then understood as the fractional residual (from the mean reddening in the pixel at the given distance) of the reddening column at the angular location and distance of star i.

We put a Gaussian prior on δ_i, with zero mean and standard deviation dependent to the scale of the pixel (allowing more variation in larger pixels) and the local dust column (allowing greater fractional variation in regions of greater reddening). In detail,

$$p(\delta_i | \mu_i, \boldsymbol{\alpha}) = \mathcal{N}(\delta_i | 0, \sigma_\delta), \tag{7}$$

with

$$\sigma_\delta = a E(\mu_i; \boldsymbol{\alpha}) + b. \tag{8}$$

Here, a and b are parameters that we set in order to match the variation we see at the given pixel angular scale in the Planck radiance-based two-dimensional dust map. We compute the rms scatter within HEALPix pixels of difference scales, finding that the scatter is well described by setting the coefficients a and b to

$$\log_{10} a = 0.88 \log_{10}\left(\frac{\varphi}{1'}\right) - 2.96, \tag{9}$$

$$\log_{10} b = 0.58 \log_{10}\left(\frac{\varphi}{1'}\right) - 1.88, \tag{10}$$

where φ is the angular pixel scale, defined as the square-root of the pixel solid angle.

We have found through trial and error that it is preferable to impose a minimum scatter of $\sim$10% on the in-pixel dust column. Additionally, if we allow δ_i to approach unity, we

clearly risk the possibility of scattering a star to a negative dust column. We therefore never allow $\sigma_\delta > 0.25$.

Although the complication introduced in this section adds an additional parameter, δ_i, for each star, it can be achieved with minimal additional computational resources. We are introducing a Gaussian scatter in the reddening of each star from the "average" reddening in the pixel, and an appropriate Gaussian smoothing of the individual stellar probability density surfaces, $p(\mu_i, E_i | \boldsymbol{m}_i)$, achieves this effect. Appendix D details how the individual stellar probability surfaces are smoothed.

2.3. Summary of Model

In summary, our model of each sightline contains the following elements:

1. The increase in the "average" reddening in each distance bin: $\boldsymbol{\alpha} = \{\Delta E_k | k = 1, \ldots, n_{\mathrm{bins}}\}$.
2. The distance modulus, μ_i, stellar type, Θ_i and fractional offset, δ_i from the "average" reddening of each star i.

The reddening of star i is therefore determined both by $\boldsymbol{\alpha}$ and δ_i. Together with the distance and type of the star, one can obtain model apparent magnitudes for the star, and thus a likelihood:

$$p(\boldsymbol{m}_i | \mu_i, \boldsymbol{\Theta}_i, \delta_i, \boldsymbol{\alpha}). \tag{11}$$

We also have per-star priors on distance and stellar type, given by a smooth model of the distribution of stars throughout the Galaxy, and a prior on the offset of each star from the local reddening column:

$$p(\mu_i, \boldsymbol{\Theta}_i, \delta_i | \boldsymbol{\alpha}) = p(\mu_i, \boldsymbol{\Theta}_i) p(\delta_i | \mu_i, \boldsymbol{\alpha}). \tag{12}$$

We finally have a log-normal prior on the increase in reddening in each distance bin along each sightline, $p(\boldsymbol{\alpha})$, whose mean is chosen to match a smooth model of the distribution of dust throughout the Galaxy.

The posterior on the line of sight reddening along one sightline is then given by

$$p(\boldsymbol{\alpha} | \{\boldsymbol{m}\}) = p(\boldsymbol{\alpha}) \prod_{i=1}^{n_{\mathrm{stars}}} p(\boldsymbol{m}_i | \mu_i, \boldsymbol{\Theta}_i, \delta_i, \boldsymbol{\alpha}) $$
$$\times p(\mu_i, \boldsymbol{\Theta}_i) p(\delta_i | \mu_i, \boldsymbol{\alpha}). \tag{13}$$

We pre-compute the likelihood and prior terms for the individual stars, and then sample in $\boldsymbol{\alpha}$. The details of how this is done are given in Green et al. (2014), and in Appendix D of this paper.

2.4. Improving Stellar Inferences using the 3D Map

In order to create the 3D dust map, we first probabilistically infer the distance and reddening to each star individually, with no information about the 3D structure of the dust. After creating the 3D dust map, however, we have a very strong constraint on how reddening should increase with distance, and this should impact our inferences about individual stellar distances and reddenings. In order to improve our stellar parameter inferences, we should replace our initial assumption about stellar reddening (a flat prior) with a new one that favors stellar reddenings close to the measured distance-reddening relation along the sightline. This can be done, in practice, by reweighting the probability density functions we initially

calculated for each star. In this section, we therefore define a reweighting of the Markov Chain samples for the individual stars which takes into account the line of sight reddening.

Let us first consider the case in which we fix the line of sight reddening profile. In the formalism used here, that means that we fix the parameters $\boldsymbol{\alpha}$, expressing the line of sight reddening profile as $E(\mu; \boldsymbol{\alpha})$. The reddening of an individual star is then determined by the distance modulus of the star, μ, as well as the fractional deviation, δ, of the stellar reddening from the local dust column. As before, there are also parameters representing the stellar type, $\boldsymbol{\Theta}$. In this formulation, the posterior density of the stellar parameters, given its photometry $\boldsymbol{m}$ and the line of sight reddening profile, is determined by

$$p(\mu, \boldsymbol{\Theta}, \delta | \boldsymbol{m}, \boldsymbol{\alpha}) \propto p(\boldsymbol{m} | \mu, \boldsymbol{\Theta}, \delta, \boldsymbol{\alpha})$$
$$p(\mu, \boldsymbol{\Theta}) p(\delta | \mu, \boldsymbol{\alpha}). \quad (14)$$

We already have Markov chain samples in distance, reddening, and stellar type for each star, for a model which does not take the line of sight reddening into account. We would like to apply weights to these samples, so that they correspond to the model sketched out directly above. As shown in detail in Appendix E, the correct reweighting of the stellar Markov chain samples, assuming a particular line of sight reddening profile $\boldsymbol{\alpha}$, is

$$w_k \propto \frac{\mathcal{N}(\delta_k | 0, \sigma_\delta)}{E(\mu_k; \boldsymbol{\alpha})}, \quad (15)$$

where k indexes the sample, and δ_k is the fractional offset of the sample reddening, E_k, from the line of sight reddening, $E(\mu_k; \boldsymbol{\alpha})$, at distance μ_k. The numerator in the above weight corresponds to a prior on the offset of the stellar reddening from the local dust column, while the denominator comes from the Jacobian transformation from reddening to fractional offset from the local reddening.

Reweighting each of the stellar parameter samples by the above factor, and normalizing the sum of the weights to unity, we obtain an inference for the stellar parameters, conditioned on a particular line of sight reddening profile, $E(\mu; \boldsymbol{\alpha})$. We marginalize over the line of sight reddening profile by repeating this procedure for each sampled value of $\boldsymbol{\alpha}$, summing the weight applied to each stellar sample.

It might be objected that since we infer $\boldsymbol{\alpha}$ using the photometry of all the stars in the given pixel, the samples we drew for $\boldsymbol{\alpha}$ in our initial processing are already dependent on the photometry for the star whose samples we are reweighting. Put more formally, the prior we use here when marginalizing over line of sight reddenings is the inference we obtained earlier, $p(\boldsymbol{\alpha} | \{\boldsymbol{m}\})$, which is conditional on all the photometry in the pixel, $\{\boldsymbol{m}\}$. We are using the photometry of a given star to infer the line of sight reddening, and then again to infer the stellar parameters, conditional on that line of sight reddening.

This would be a problem if the photometry of any given star significantly affected the line of sight reddening inference. However, if photometry from a large number of stars informs the line of sight reddening, then no one star should have a significant impact on the inferred line of sight reddening profile. In a more formally correct formulation of the problem, we would first infer the line of sight reddening using all the stars in the pixel except the star whose parameters we wish to infer, and we would then infer the parameters for that star, conditional on the inferred line of sight reddening profile. As each pixel contains hundreds of stars, however, we expect our
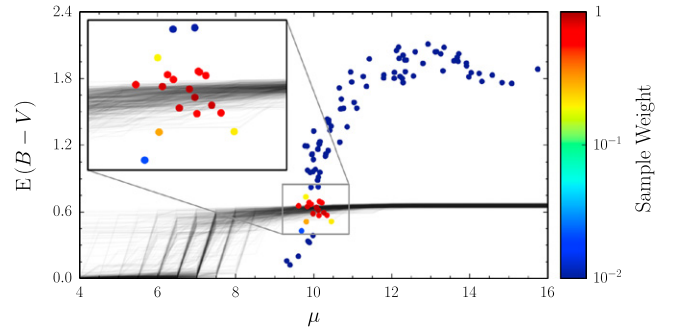

Figure 3. Reweighting Markov Chain samples for an individual star, based on the line of sight reddening profile. Each black line is a reddening profile drawn from the posterior on line of sight reddening. The dots are posterior samples of parameters (including distance and reddening) for one star viewed in isolation, i.e., not conditioned on the line of sight reddening. The samples are reweighted in order to condition them on the line of sight reddening, assigning greater weight to samples that are consistent with the reddening profile. Including information about the line of sight reddening can significantly reduce the uncertainties in stellar distance, reddening and, type.

procedure to approximate this formally correct procedure closely.

Figure 3 shows the result of reweighting the samples for one star. Depending on the line of sight reddening profile and the distribution of the unweighted stellar parameter samples, individual stellar inferences can be tightened dramatically by taking the line of sight reddening into account. Our knowledge of the parameters describing one star is therefore dependent not only on its photometry, but also on the photometry of its neighbors in the sky. Because neighboring stars lie along the same dust column, inferences for nearby stars are coupled through the requirement that the dust column increase with distance.

3. DATA

3.1. Pan-STARRS 1

PS1 is a 1.8 m optical and NIR telescope located on Mount Haleakala, Hawaii (Hodapp et al. 2004; Kaiser et al. 2010). The telescope is equipped with the GigaPixel Camera 1, consisting of an array of 60 CCD detectors, each 4800 pixels on a side (Onaka et al. 2008; Tonry & Onaka 2009). From 2010 May to 2014 April, the majority of the observing time was dedicated to a multi-epoch 3π steradian survey of the sky north of $\delta = -30°$ (K. C. Chambers 2015, in preparation). The 3π survey observes in five passbands g_{P1}, r_{P1}, i_{P1}, z_{P1}, and y_{P1}, similar to the SDSS (York et al. 2000), with the most significant difference being the replacement of the Sloan u band with a NIR band, y_{P1}. The PS1 filter set spans 400–1000 nm (Stubbs et al. 2010). The images are processed by the PS1 Image Processing Pipeline (Magnier 2006), which performs automatic astrometry (Magnier et al. 2008) and photometry (Magnier 2007). The data is photometrically calibrated to better than 1% accuracy (Schlafly et al. 2012; Tonry et al. 2012). The 3π survey reaches typical single-exposure depths of 22 mag (AB) in g_{P1}, 21.5 mag in r_{P1} and i_{P1}, 20.8 mag in z_{P1}, and 20 mag in y_{P1}. The resulting homogeneous optical and NIR coverage of three quarters of the sky makes the Pan-STARRS1 data ideal for studies of the distribution of the Galaxy's dust.

3.2. 2MASS

2MASS is a uniform all-sky survey in three NIR bandpasses, J, H, and K_s (Skrutskie et al. 2006). The survey derives its name from the wavelength range covered by the longest-wavelength band, K_s, which lies in the longest-wavelength atmospheric window not severely affected by background thermal emission (Skrutskie et al. 2006). The survey was conducted from two 1.3 m telescopes, located at Mount Hopkins, Arizona and Cerro Tololo, Chile, in order to provide coverage for both the northern and southern skies, respectively. The focal plane of each telescope was equipped with three 256×256 pixel arrays, with a pixel scale of $2'' \times 2''$. Each field on the sky was covered six times, with dual 51 ms and 1.3 s exposures, achieving a 10σ point-source depth of approximately 15.8, 15.1 and 14.3 mag (Vega) in J, H and K_s, respectively. Calibration of the survey is considered accurate at the 0.02 mag level, with photometric uncertainties for bright sources below 0.03 mag (Skrutskie et al. 2006).

The addition of 2MASS data requires an expansion of the stellar model described in Green et al. (2014) to cover the 2MASS J, H, and K_s passbands, as well as an estimate of the survey selection function for 2MASS. We address these additions to our model in Appendices A and B.

3.3. Source Selection

We match each PS1 source to the nearest source in the 2MASS point-source catalog, rejecting stars at greater than $2''$ separation. We require detection in at least four passbands, two of which must be PS1 passbands. In order to reject extended sources, we require that $m_{\mathrm{psf}} - m_{\mathrm{aperture}} < 0.1$ mag in at least two PS1 passbands. We additionally reject sources flagged as extended sources in 2MASS.

Individual PS1 passbands are rejected if they have a photometric uncertainty greater than 0.2 mag, or if the sources are beyond or close to the saturation limit for PS1 (here considered 14.5, 14.5, 14.5, 14, and 13 mag in $grizy_{\mathrm{P1}}$, respectively). In 2MASS passbands, we make the recommended "high-reliability catalog" selection cuts,[8] in addition to the following requirements:

1. `contamination/confusion flag = 0`,
2. `galaxy contamination flag = 0`.

Our final catalog contains 798,611,689 sources, of which 32% are detected in four passbands, 49% in five passbands, and 19% in six or more passbands.

3.4. Pixelization

We divide the sky into HEALPix pixels (Gorski et al. 2005), adjusting the pixel scale in order to keep the number of stars per pixel roughly constant (see Figure 2). Our procedure is to begin with `nside = 64` pixels, and then subdivide each pixel recursively, as long as the number of stars exceeds some threshold, dependent on the pixel scale. We use thresholds given in Table 1, chosen to allow us to reach a resolution of `nside = 512` with a relatively small number of stars, but to avoid going to higher resolutions unless the stellar density is much higher. We reject pixels with fewer than ten stars. Such

---

Table 1
Pixelization

nside	Max. Stars Pixel^{-1}	Solid Angle At Resolution (deg^2)	# Of Pixels
64	200	77	91
128	250	90	430
256	300	11,980	228,373
512	800	16,071	1,225,471
1024	1200	2957	901,971
2048	⋯	66	80,956
Total	⋯	31,240	2,437,292

pixels comprise a negligible fraction of the sky. In all, we assign just under 800 million stars, covering just over three quarters of the sky, to 2.4 million pixels, with an average of 327 stars per pixel.

4. RESULTS

Applying the method described in Green et al. (2014), with the modifications described above, to PS1 and 2MASS stellar photometry, we produce a three-dimensional map of dust reddening, covering the three-quarters of the sky north of $\delta = -30°$.

4.1. Distance Slices of Map

In Figures 4 and 5, we present the differential reddening in four spherical shells, centered on the Sun. Each panel shows the median dust reddening in a different range of solar-centric distances. Due to perspective, dust at high Galactic latitudes resides nearby, as Galactic dust lies in a thin disk. We recover the wealth of structure seen in the ISM across a wide range of scales, from thin filaments to large cloud complexes. Large, coherent cloud complexes appear at consistent distances.

Figure 6 gives a closer view of the anticentral region ($\ell \sim 180°$). Different features appear clearly in each rendered distance slice. The Perseus, Taurus, and Auriga cloud complexes dominate the anticentral region in the closest distance slice, while the Orion molecular cloud complex ($\ell \sim 210°$, $b \sim -15°$) and the California nebula ($\ell \sim 160°$, $b \sim -8°$) appear very strikingly in the second distance slice. Of particular interest is the ring-like structure that Orion A and B appear to be embedded within. This ring-like structure is only apparent when the background dust is removed. In particular, the northeast portion of the ring is confused with the plane of the Galaxy in projection. Schlafly et al. (2014a) discusses the "Orion ring," including possible formation scenarios for the ring, in greater depth.

Figure 7 shows the Galactic plane from $\ell = 60°$ to $155°$ in more detail. The Cepheus flare, which lies at the center of the frame, at $95° \lesssim \ell \lesssim 110°$, $b \approx 15°$, separates into two components at different distances, visible in the first and third panels. Using a modified version of the method used in this paper along individual sightlines, Schlafly et al. (2014b) places the two components of the Cepheus flare at distances of 360 ± 35 pc and 900 ± 90 pc. The dust associated with Cygnus X ($\ell \approx 80°$, $b \approx 0°$) appears in the third panel, along with a wealth of fine structure along the Galactic plane.
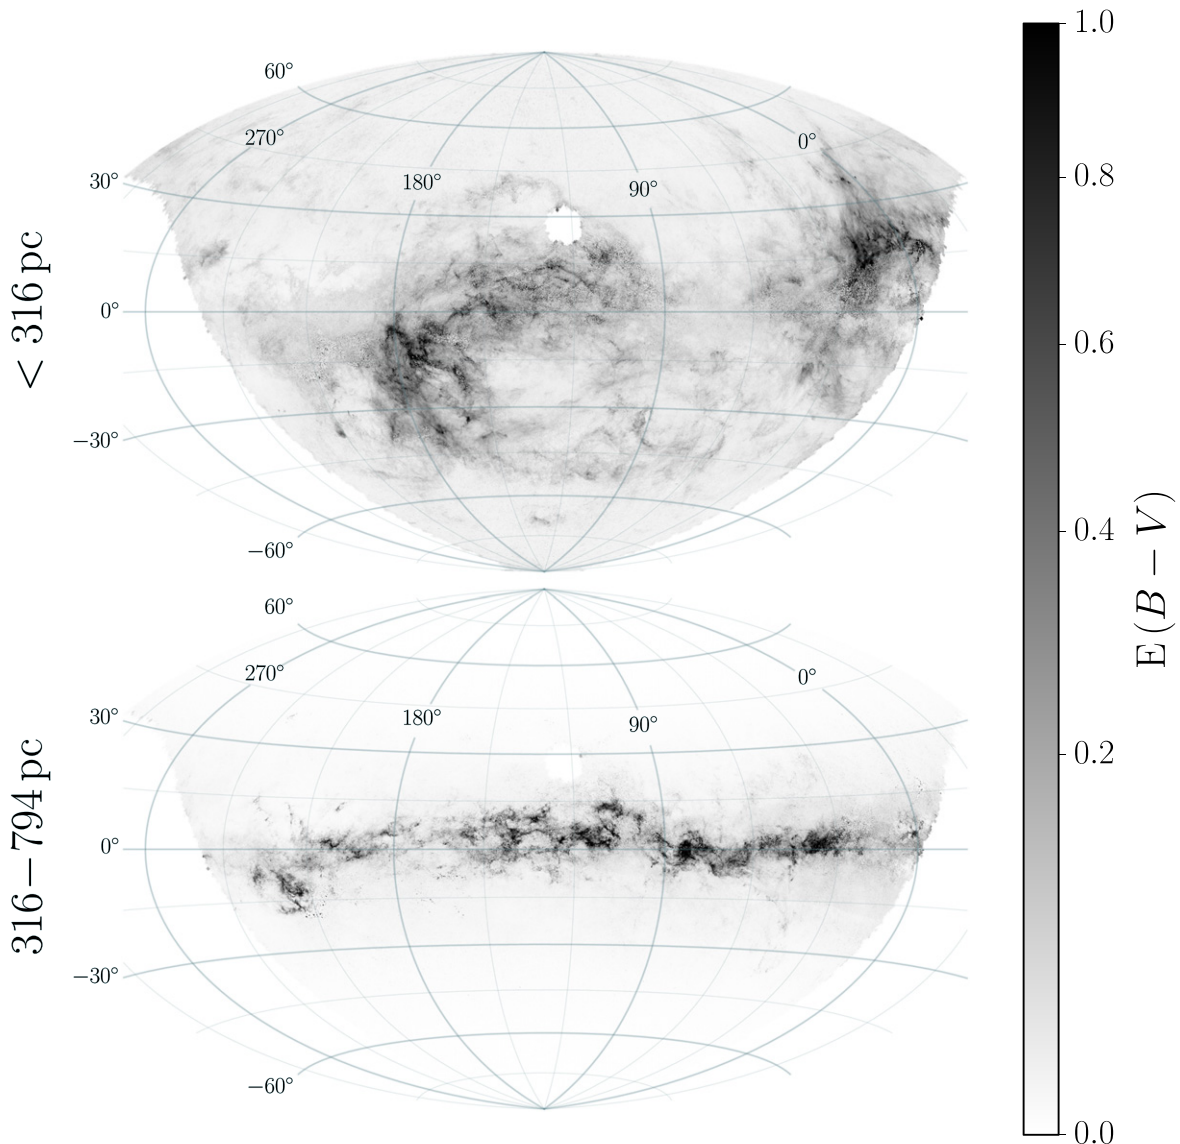
Figure 4. Median differential reddening in two solar-centric distance ranges. The distance breaks are chosen to coincide with distance moduli $\mu = 7.5$ and 9.5, which line up with edges of distance bins in our map. We adopt a square-root stretch, in order to capture both low- and high-reddening features. The hole at $\ell \approx 120°$, $b \approx 30°$ corresponds to declinations above $\sim 84°$, which had not yet been fully processed at the time we created our 3D map.

We note that each pixel is fit independently, and our only prior assumption about the spatial structure of the dust is that if forms a diffuse disk, as shown in Figure 1. The detailed spatial structure in the interstellar medium that our analysis reveals indicates that the PS1 and 2MASS photometry dominates over our priors out to a distance of several kiloparsecs and reddening of $E(B-V) \approx 1.5$ mag. With the assumption of spatial correlations between neighboring pixels, we expect that one would be able to significantly reduce the uncertainty in the map, and achieve better distance resolution (see, e.g., Sale & Magorrian 2014).

4.2. Maximum and Minimum Reliable Distances in Map

Our 3D dust map is based on measurements of stellar distances and reddenings. Beyond the most distant stars, we have no sensitivity to dust, and in front of the nearest stars, we have no information about the distance to the dust. Therefore, we estimate the minimum and maximum distance to which our map is reliable by locating the nearest and farthest stars in each

pixel. Outside of this distance range, our line of sight reddening inferences are dominated by our priors. Using the improved stellar parameter inferences (described in Section 2.4), we define the minimum reliable distance in each line of sight as the distance out to which there are N_{closer} observed main-sequence stars, and the maximum reliable distance as the distance beyond which there are N_{farther} observed main-sequence stars.

For this calculation, we count each Markov Chain sample in stellar distance as a fraction of a star. We exclude stars that fail to converge, for which the model is a very poor match to the data (as determined by the Bayesian evidence for the star; see Green et al. 2014), or which are inconsistent with the inferred line of sight reddening profile. We consider a star inconsistent with the line of sight reddening inference if none of the 100 stored Markov Chain samples of the stellar distance and reddening is within a fractional distance σ_δ (the modeled intra-pixel scatter in the dust column; see Section 2.4) of the line of sight reddening profile. Such objects are likely not well fit by any of our stellar templates, or alternatively signal that there is
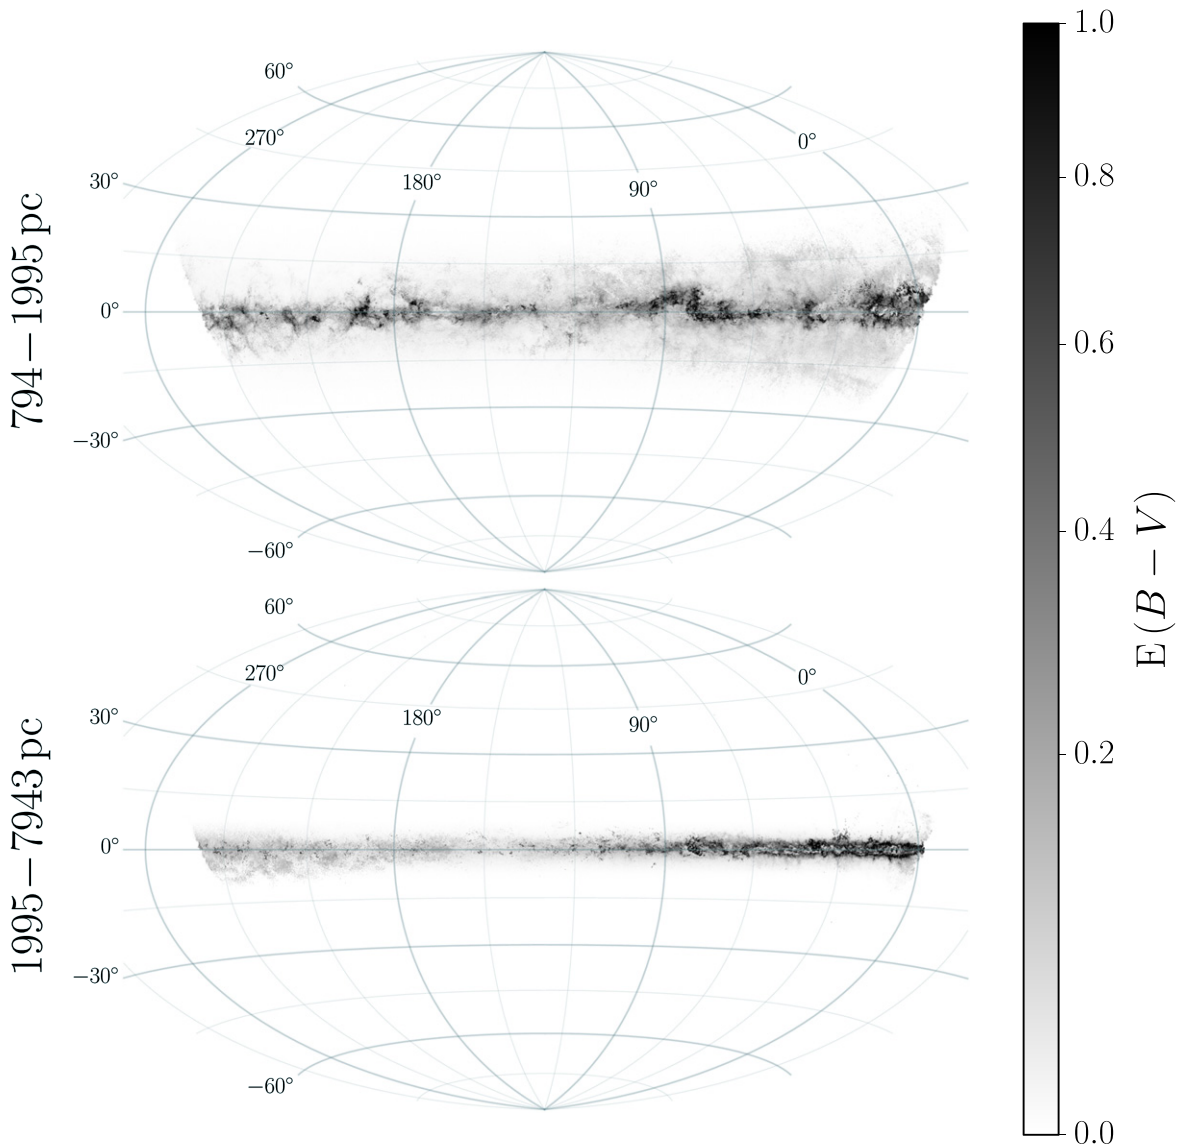
Figure 5. Same as in Figure 4, but with breaks at distance moduli $\mu = 11.5$ and 14.5.

more variation in the dust column at fine angular scales than we allow.

In determining the minimum and maximum reliable distances in the 3D dust map, we use only main-sequence stars. This is because we consider our inferences for giants to be less reliable than our inferences for dwarfs. The colors and luminosities of giants depend more strongly on metallicity and age, the latter of which we do not model.

In the left two panels of Figure 8, we show the results obtained by requiring $N_{\rm closer} = 2$ and $N_{\rm farther} = 10$. The results are qualitatively similar for other choices of these parameters, as long as they are small compared to the typical number of stars in a pixel. The closest reliable distance is set almost entirely by the angular pixel scale. The nearby density of stars is, to a very rough estimation, uniform, meaning that the distance to the closest star in a pixel is a function primarily of the solid angle covered by the pixel. Accordingly, the top left panel of Figure 8 is essentially a map of pixel solid angle, with boundaries in distance following pixel scale boundaries.

The farthest reliable distance of the 3D dust map is strongly influenced not only by the pixel scale, but also by the

distribution of stars throughout the Galaxy, the 3D distribution of dust and the survey depth. The survey depth plays a larger role in the far limit because the most distant observed stars lie preferentially near the limiting survey magnitude. The closest stars, by contrast, are distributed more evenly in apparent magnitude. Accordingly, the high-latitude patches with particularly shallow maximum reliable distances (e.g., the shallow patch near $\ell = 205°$, $b = -55°$) correspond to areas which are missing one or more PS1 bands, or which have fewer observation epochs in the PS1 3π survey, and therefore worse coverage.

As a check on these estimated nearest and farthest reliable distance maps, we construct equivalent maps using simulated stellar catalogs. For each pixel in our map, we use our Galactic priors and the estimated limiting magnitudes for $grizy_{\rm P1}$ and 2MASS JHK_s to generate an equal number of stars as actually observed. We then determine the distance in front of which there are $N_{\rm closer}$ and beyond which there are $N_{\rm farther}$ main-sequence stars. In order to remove the complicating factor of the line of sight dust inference, we restrict this test to $|b| > 15°$ and to pixels for which $E(B-V)_{\tau_{353}} < 0.05$ mag, and assume
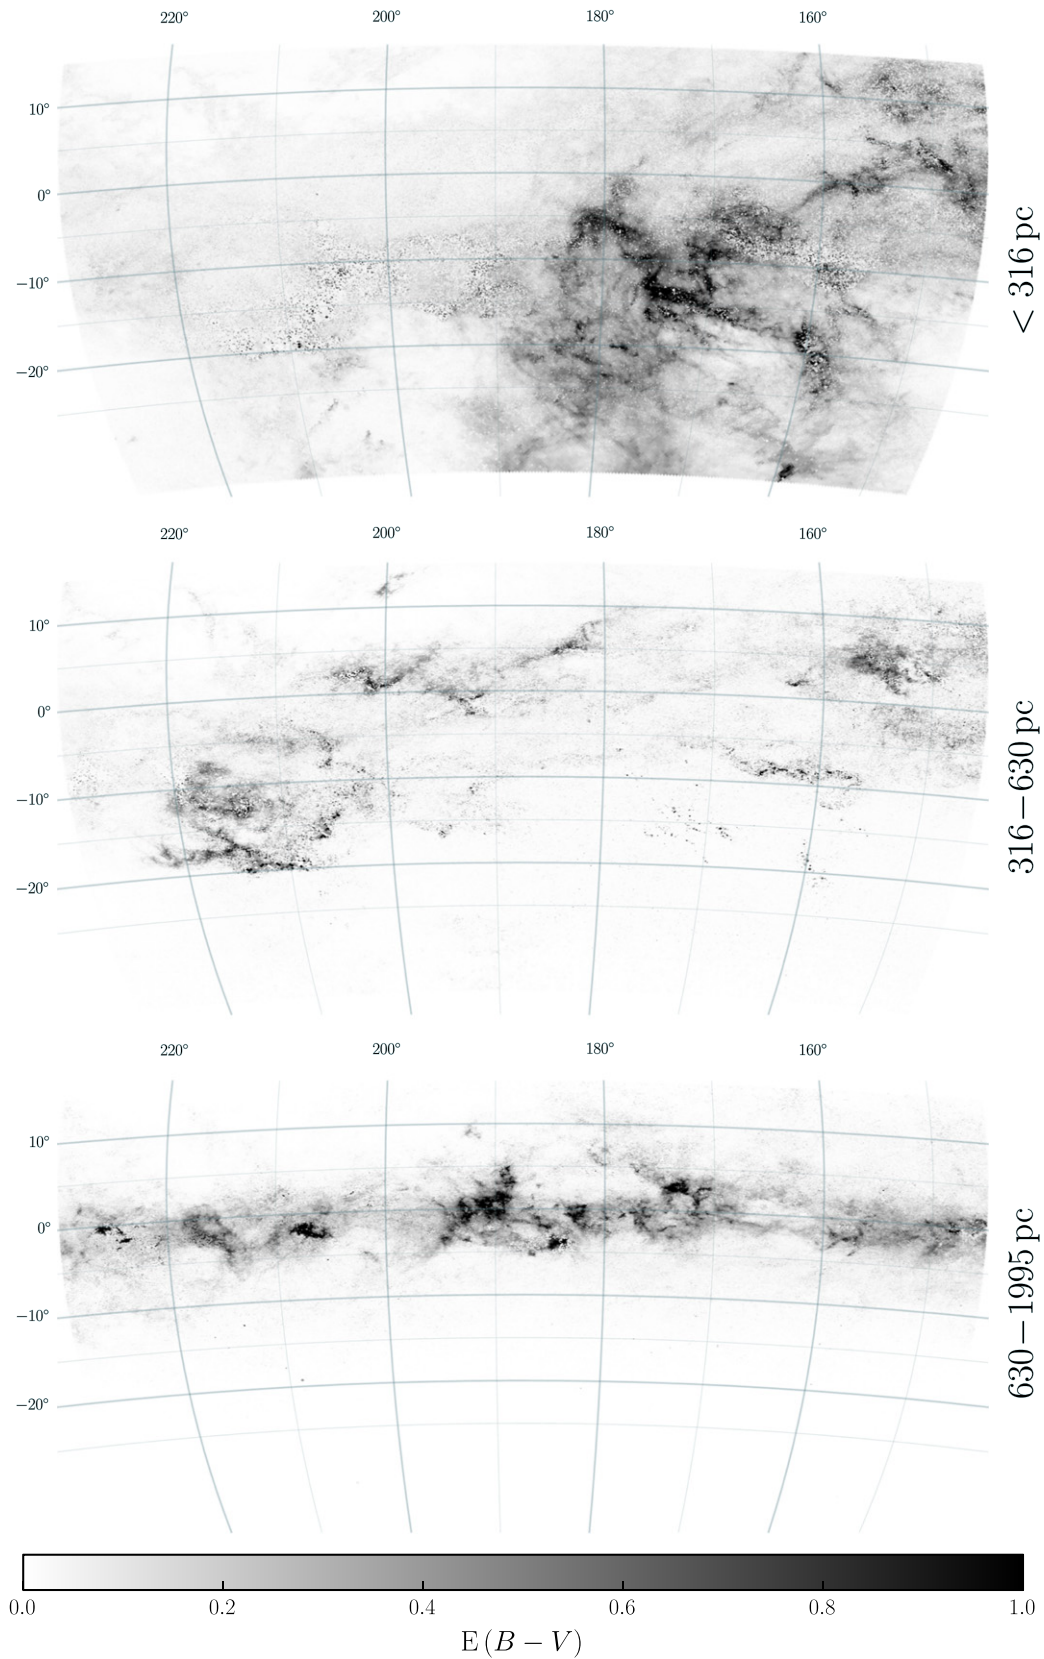
Figure 6. Closer view of the dust in the anticentral region in three distance bins. As in Figures 4 and 5, we plot the median differential reddening. The Taurus–Perseus–Auriga complex is visible in the right half of the nearest distance bin. In the second distance bin, the Orion complex is visible on the left, while the California cloud is visible on the right. Note the ring-like shape of the Orion complex, which is only revealed by 3D mapping when confusion from background dust is removed. See Schlafly et al. (2014a) for a discussion of this feature. Monoceros R2 appears beyond Orion, in the third distance bin, flanked by the plane of the Galaxy at a yet greater distance.
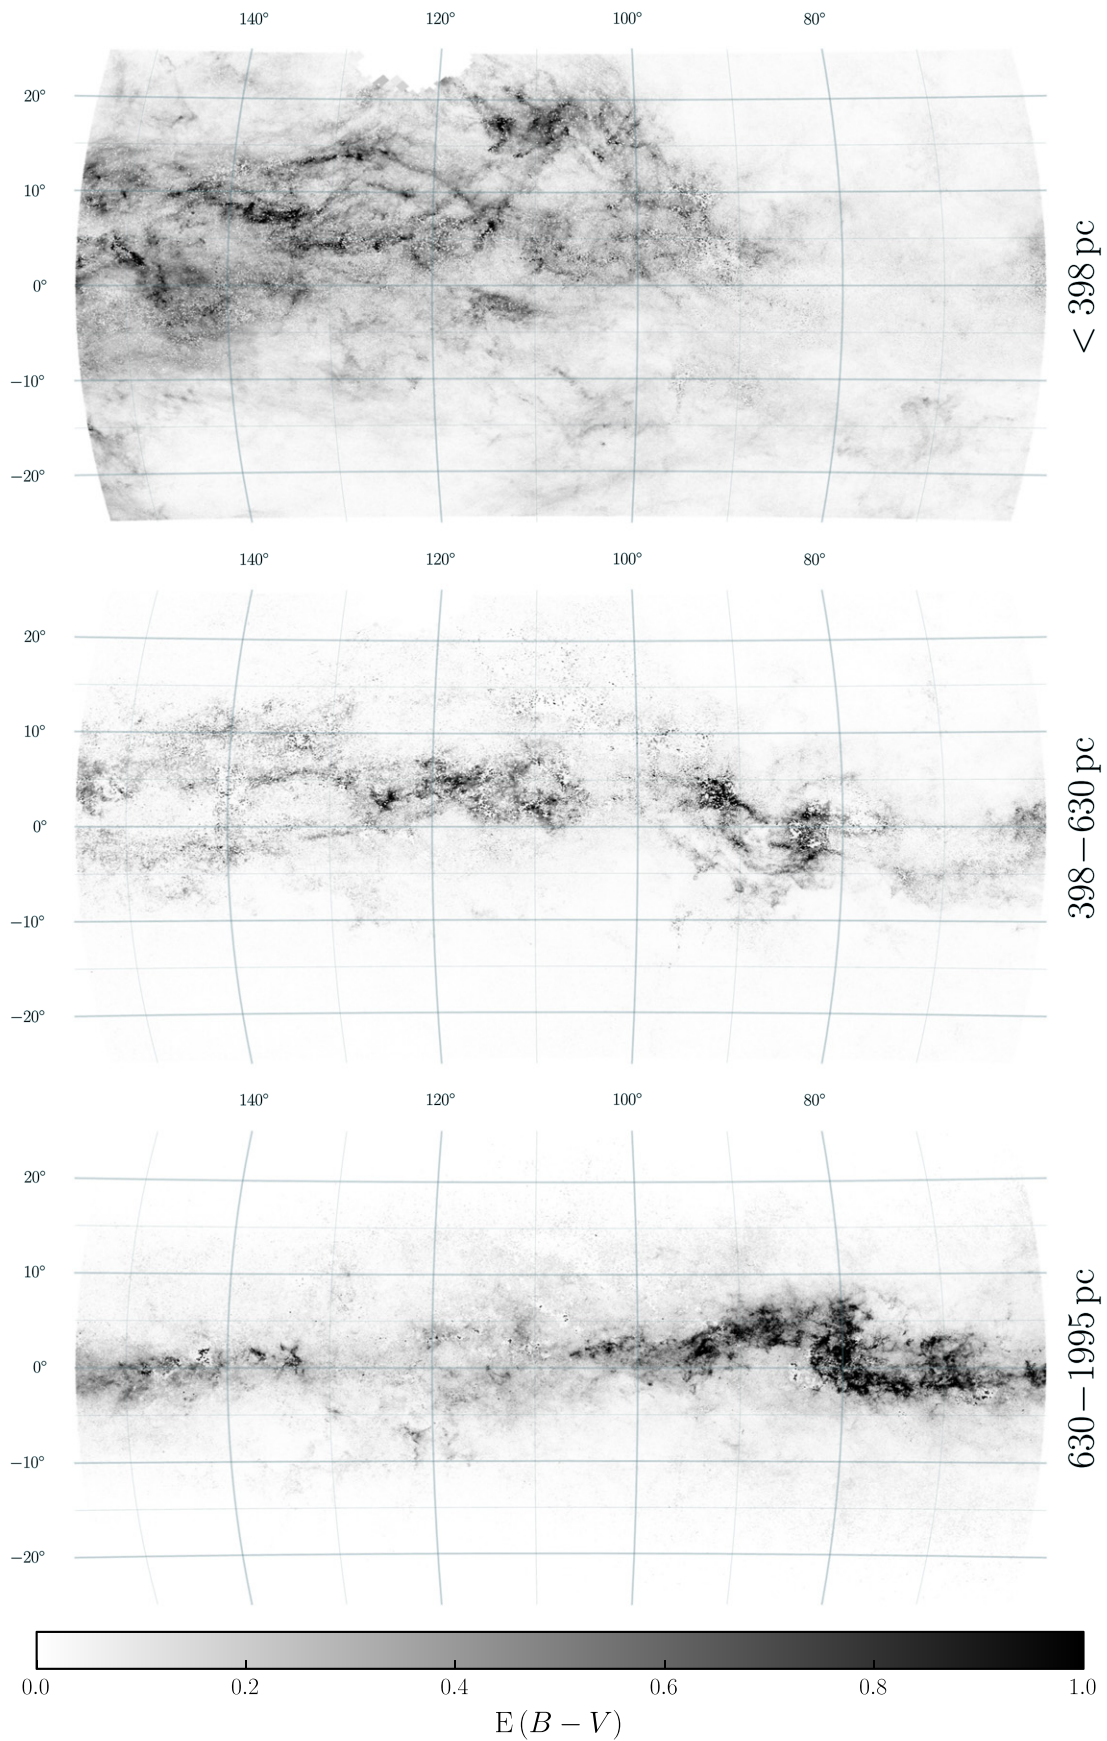
Figure 7. Closer view of the dust in the direction of Cepheus and Polaris flares and the eastern portion of the Great Rift, including Cygnus X. The Cepheus flare $(95° \lesssim \ell \lesssim 110°, b \approx 15°)$ separates into two components in distance, visible in the first and third panels. The dust associated with the Cygnus X region $(\ell \approx 80°, b \approx 0°)$ appears in the third panel.
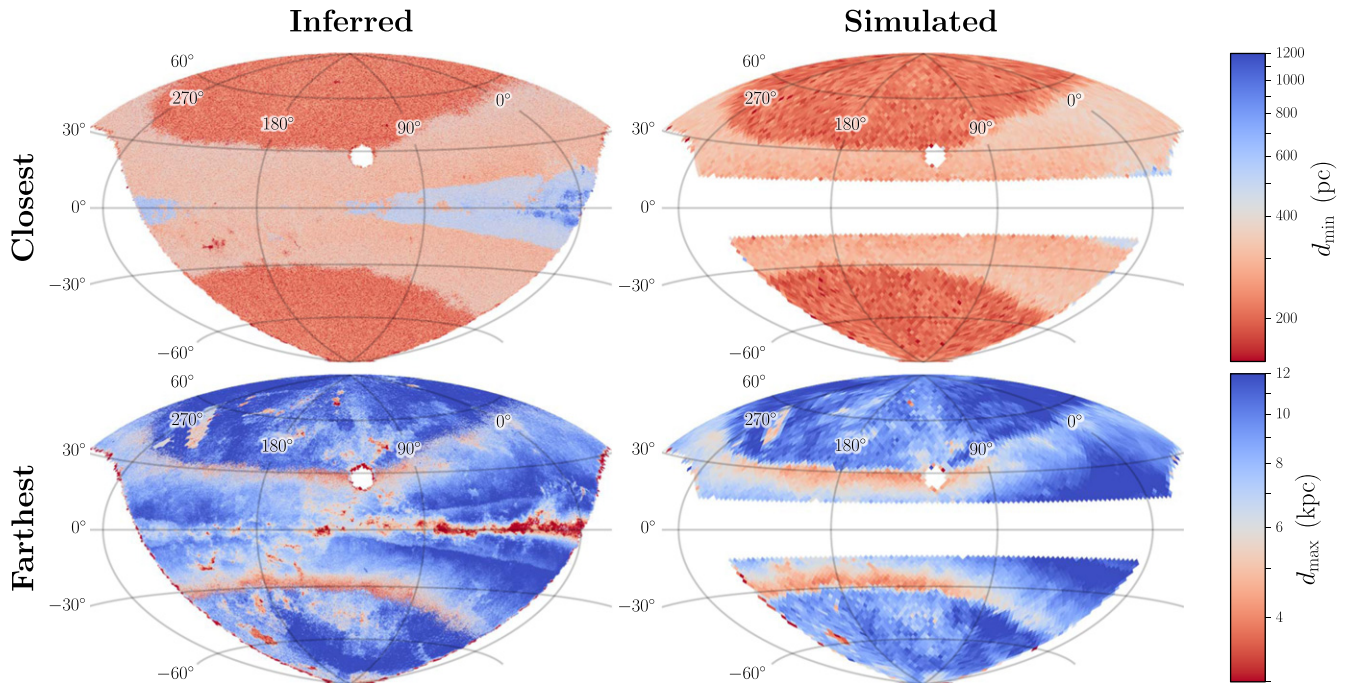
Figure 8. Minimum and maximum reliable distances in the 3D dust map. The left two panels are based on the inferred distribution of stars along each line of sight, while the right two panels are based on simulated stellar catalogs generated from our Galactic priors. We consider the map reliable if there are at least two stars closer than the given distance, and at least ten stars beyond the distance. Pixel size, survey depth and completeness, stellar density throughout the Galaxy, and the presence of dust all affect the distribution of observed stars along the line of sight, and therefore the distance range over which the map is reliable. Sharp transitions in reliable distance occur at pixel size boundaries, where the number of stars per pixel changes discontinuously, as well as in patches of the sky with missing PS1 passbands.

for our simulated catalogs that there is zero dust extinction. In these comparisons, we find that a halo number density close to the value given in Jurić et al. (2008) is required to reproduce inferred map depths of Figure 8. The results for a halo strength of $n_h = 0.004$, binned down to `nside = 64`, are shown in the right two panels of Figure 8.

These results indicate that while we dialed down the halo strength in our priors, the data, as reflected in our photometric stellar inferences, nonetheless prefers a stronger halo. In future work, we will investigate the implications of our stellar and line of sight dust inferences for a global Galactic structure model.

4.3. Stellar Inferences

As described in Section 2.4, we reweight the naively inferred parameters for each star in order to take the line of sight reddening profile into account. In order to demonstrate the improvement in stellar inferences after this re-weighting procedure, we compare our reddening inferences to a set of independently determined stellar reddening standards, as in Green et al. (2014). For this, we use the set of stellar reddening standards developed by Schlafly & Finkbeiner (2011). The Sloan Extension for Galactic Understanding and Exploration (SEGUE; Yanny et al. 2009), part of SDSS-II, used moderate-resolution spectra to classify 240,000 stars. Empirically adjusted SDSS colors based on model atmospheres can then be compared with observed SDSS photometry to obtain reliable reddening estimates. We select the same sample of 200,000 SEGUE target stars as Schlafly & Finkbeiner (2011), which excludes white dwarfs and M dwarfs (the former because they are not contained in our model, and the latter because of their unreliable spectral classification).

In Figure 9, we compare reddening samples drawn from our PS1+2MASS-based stellar inferences to samples drawn from

the SEGUE-based estimates. which have Gaussian uncertainties. The agreement between the two reddening estimates improves significantly after reweighting our PS1+2MASS-based stellar inferences. The improvement is most pronounced at low reddenings. The improvement is negligible for $E(B-V) \gtrsim 1$ mag, due to the fact that we allow the individual stars to deviate from the local line of sight reddening profile by about 10%. Unlike the SEGUE-based reddenings, the PS1+2MASS-based reddening inferences are constrained to be non-negative, leading to a negative slope in the residuals near zero reddening.

5. COMPARISON WITH PREVIOUS DUST MAPS

Schlafly et al. (2014c) compares an earlier version of our 3D dust map with the two-dimensional, FIR emission-based SFD (Schlegel et al. 1998) and Planck dust maps (Planck Collaboration et al. 2014). Here, we repeat this comparison using our new 3D dust map, and additionally compare our dust map with previous 3D dust maps, which are also based on stellar photometry.

5.1. 2D Dust Maps

It is possible to obtain a 2D dust map from our 3D map by projecting out distance, i.e., by taking the cumulative reddening out to some large distance. Such a map is of particular interest for extragalactic astronomy, where Milky Way dust is essentially a foreground screen to be removed. Several all-sky 2D maps of dust reddening already exist, among them Burstein & Heiles (1982), based on H I emission and galaxy number counts, SFD and two recent reddening maps derived from Planck data (Planck Collaboration et al. 2014), which model dust optical depth and temperature from FIR emission,
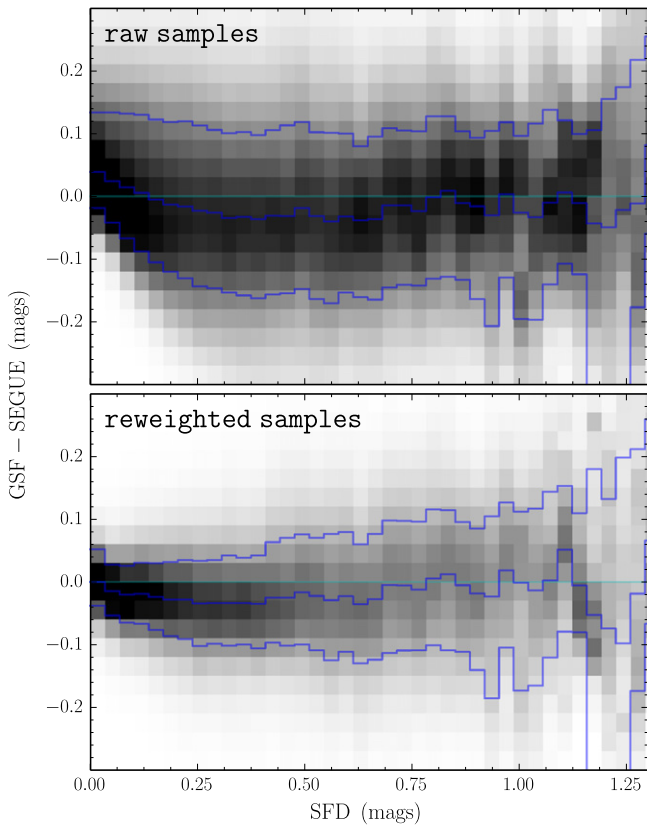
Figure 9. Histograms of the residuals of our PS1+2MASS-based stellar reddenings (referred to as "GSF" here), vs. reddening estimates obtained by comparing SEGUE spectral classifications with SDSS photometry. The histograms are spread out along the *x*-axis by the local SFD reddening. The blue lines trace the 16th, 50th, and 84th percentiles of the residuals in each bin of SFD reddening. In the upper panel, we use unweighted samples drawn from the individual stellar parameter Markov Chains. In the lower panel, we use reweighted samples, as described in the text.

and then calibrate a dust optical depth to reddening relation. NIR stellar colors have also been used to determine dust reddening more directly, as in the NICE/NICER/NICEST family of algorithms (Lada et al. 1994; Lombardi & Alves 2001; Lombardi 2009, respectively), and Rowles & Froebrich (2009).

Schlafly et al. (2014c) compares a 2D projection of an earlier version of our 3D dust map with the widely used SFD reddening map, as well as the newer Planck reddening maps. We repeat a number of the same tests for our new 3D dust map.

We begin, however, with a comparison between the map presented in this paper and the map presented in Schlafly et al. (2014c). The most important differences between the 3D maps used here and in Schlafly et al. (2014c) are the addition of NIR 2MASS photometry in our newer map, and that we sample here from the full posterior distribution on line of sight reddening, rather than finding the maximum-likelihood line of sight reddening. Because Schlafly et al. (2014c) requires that each star be detected in g_{P1}, and because we incorporate NIR 2MASS photometry alongside PS1 photometry, our map reaches to deeper dust extinctions. This is apparent in the right panel of Figure 10, where our new map predicts more reddening both in the inner Galactic plane and in dense dust clouds off the plane, such as Orion, Taurus and Perseus.

At reddenings below $E(B-V) \lesssim 0.08$ mag, our new map predicts significantly less reddening than Schlafly et al.

(2014c), as can be seen in the left panel of Figure 10. Beyond $E(B-V) \approx 0.1$ mag, our reddening scales agree very closely, with our new map predicting ~4% more reddening. The scatter between the two reddening measures comes to ~0.04 mag, with a maximum median offset of 0.04 mag at $E(B-V) \approx 0.08$ mag, decreasing gradually to an offset of less than 0.01 mag between the two measures at $E(B-V) = 1$ mag. The behavior of the residuals suggests that at very low reddenings, our new 3D dust map prefers essentially zero reddenings too strongly. This may be due to the stronger priors on dust reddening used in the present work, or due to slight differences in our new combined PS1-2MASS stellar locus, versus the PS1 stellar locus used in Schlafly et al. (2014c).

Next, we compare our inferred cumulative reddening out to 5 kpc with the SFD map, the Planck $\tau_{353\,\mathrm{GHz}}$-based reddening (hereafter, denoted simply as τ_{353}, with units of magnitudes of $E(B-V)$), and the Planck radiance-based map (hereafter denoted $\mathcal{R}$, likewise with units of magnitudes of $E(B-V)$). All three of these maps are derived by modeling dust emission, and should therefore have different types of systematic errors than our stellar reddening-based dust map.

The upper panel of Figure 11 shows the SFD reddening across the footprint of our map, while the lower panel shows the residuals after subtracting off our 5 kpc map. Off the Galactic plane, the residuals are small, while larger residuals are found in the plane of the Galaxy, where there is significant dust past 5 kpc, and where PS1 does not necessarily detect stars beyond all of the dust. Near the Galactic center, one noticeable feature is a blue halo, where we infer more dust than SFD. The residuals are correlated with dust reddening in this area, indicating a scale offset between the two maps, rather than a constant offset.

The "blue halo" is again apparent if we compare our map to the Planck τ_{353} map. In Figure 12, we compare our inferred reddening at 5 kpc with SFD, τ_{353}, and $\mathcal{R}$. The left panels map the residuals across the PS1 survey footprint on a square-root stretch, emphasizing small-amplitude differences. In the left panels, the Planck-based maps have been scaled by a constant factor to match our inferred PS1 reddening for $|b| > 20°$.

As the "blue halo" occurs in the direction of the nearby Aquila Rift, it is possible that the cloud has anomalous dust properties. For example, the grain size distribution or composition may vary, so that $R_V = 3.1$ is not a good assumption in the Aquila Rift. Another possibile explanation for the "blue halo" is that our stellar inferences may be systematically biased toward greater reddenings in this direction due to limitations in our Galactic model. Our priors do not, for example, include a radial metallicity gradient in the disk components of the Galaxy, which could lead to biased metallicity estimates for stars toward the Galactic center. We leave this question for future investigation.

In the right panels of Figure 12, we plot the residuals of our inferred reddening to 5 kpc with SFD and the Planck emission-based maps as a function of reddening. Here, we do not scale the Planck maps by any factor to bring them into alignment with our map. To conduct this comparison, we compare the emission-based maps to multiple random realizations drawn from the posterior on reddening to 5 kpc from our 3D dust map. We restrict our comparison to high-Galactic-latitude regions ($|b| > 30°$), and cut out the ecliptic plane ($|\beta| < 20°$), where imperfectly subtracted Zodiacal light might contaminate the
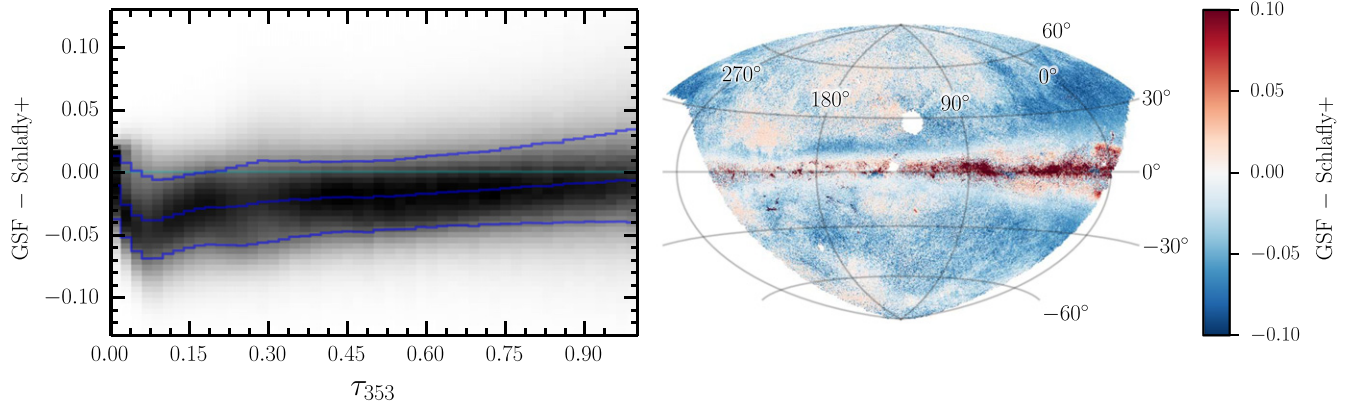
Figure 10. Comparison of our new 3D dust map ("GSF"), integrated to 4.5 kpc, with the 2D map presented in Schlafly et al. (2014c). The left panel shows the difference between our map and Schlafly et al. (2014c) as a function of a third reddening measure with uncorrelated errors, the Planck $\tau_{353\,GHz}$-based dust reddening map. The right panel maps the median residuals between our map and Schlafly et al. (2014c) across the sky. All units are in magnitudes of $E(B-V)$.
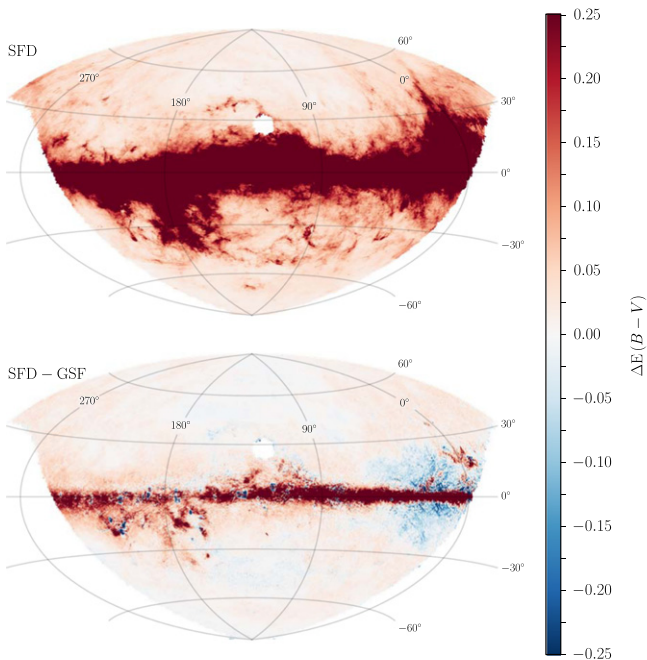

Figure 11. Comparison of the 2D SFD reddening map with our PS1-based map ("GSF"), integrated out to 5 kpc. The top panel shows the SFD reddening over the footprint of the PS1 survey, clipped to 0.25 mag, while the bottom panel shows the residuals after subtracting off the median cumulative reddening out to 5 kpc predicted by our 3D dust map. Both panels use the same color scale.

emission-based reddening maps. When comparing our PS1-based reddening with SFD, we place the Planck $\tau_{353\,GHz}$-based reddening on the x-axis, as its errors are uncorrelated with both maps on the y-axis. When displaying the residuals between our PS1-based reddening and the Planck reddening maps, we place SFD along the x-axis for the same reason.

For reddenings above $E(B-V) \gtrsim 0.05$ mag, we see broadly similar residuals as found in Schlafly et al. (2014c), with different behaviors above and below $E(B-V) \approx 0.15$ mag. Above $E(B-V) \approx 0.15$ mag, our map predicts about 10% less reddening than SFD, but is in good agreement with τ_{353}. As in Schlafly et al. (2014c), we find an overall difference in scale between our PS1-based map and the Planck $\mathcal{R}$-based map. For $E(B-V) \lesssim 0.05$ mag, we see the same residual between our map and the emission-based maps as found between our map

and Schlafly et al. (2014c). For these small reddenings, our map favors essentially zero reddening too heavily.

The exact behavior of the residuals depends on which regions of the sky are masked in the analysis, indicating that there are spatially dependent systematic differences in the residuals between our reddening map and emission-based maps. However, the essential features remain the same, with different slopes in the residuals below and above $E(B-V) \approx 0.15$ mag, and our map favoring lower reddening below $E(B-V) \approx 0.05$ mag.

5.2. Marshall et al. (2006)

Marshall et al. (2006) developed a method to determine the reddening-distance relation along individual lines of sight by comparing 2MASS $J - K_s$ stellar colors to those of simulated catalogs based on the Besançon model of the Galaxy (Robin et al. 2003). Marshall et al. (2006) then applied this method to a regular grid of sightlines separated, by 15′ covering the region $|\ell| < 100°$, $|b| < 10°$. The result is a 3D map of reddening in the inner Galaxy, extending to a maximum extinction of $\sim1.4-3.75$ mag in the 2MASS K_s band (equivalent to $\sim4.5-12$ mag in $E(B-V)$), and a maximum distance of ~7 kpc. Because Marshall et al. (2006) use only giants in their analysis, the dust map they produce has little information in the nearest kiloparsec. We will refer to this map as the "Marshall map."

Our dust map overlaps with the Marshall map in the approximate region $0° < \ell < 100°$, $|b| < 10°$. Figure 13 shows the cumulative reddening at increasing distances in both the Marshall map and our 3D dust map. When converting from extinction to reddening, we assume $A_{K_s} = 0.320 \times E(B-V)$, as calculated by Yuan et al. (2013) for a 7000 K source spectrum at $E(B-V) = 0.4$ mag, using the Cardelli et al. (1989) reddening law and assuming $R_V = 3.1$. We mask our map beyond our predicted maximum reliable distance, as determined in Section 4.2. In the large-distance limit, our maps show good qualitative agreement outside of the masked areas. Figure 14 shows the differential reddening in the two maps in bins of increasing distance.

In the nearest two to three kiloparsecs, our map shows clearly differentiated structures at discrete distances that are spread over several distance bins in the Marshall dust map. For example, the Cygnus rift (located at $\ell \sim 80°$, $b \sim 0°$) appears clearly in our map in the distance bin spanning $0.5-1$ kpc,
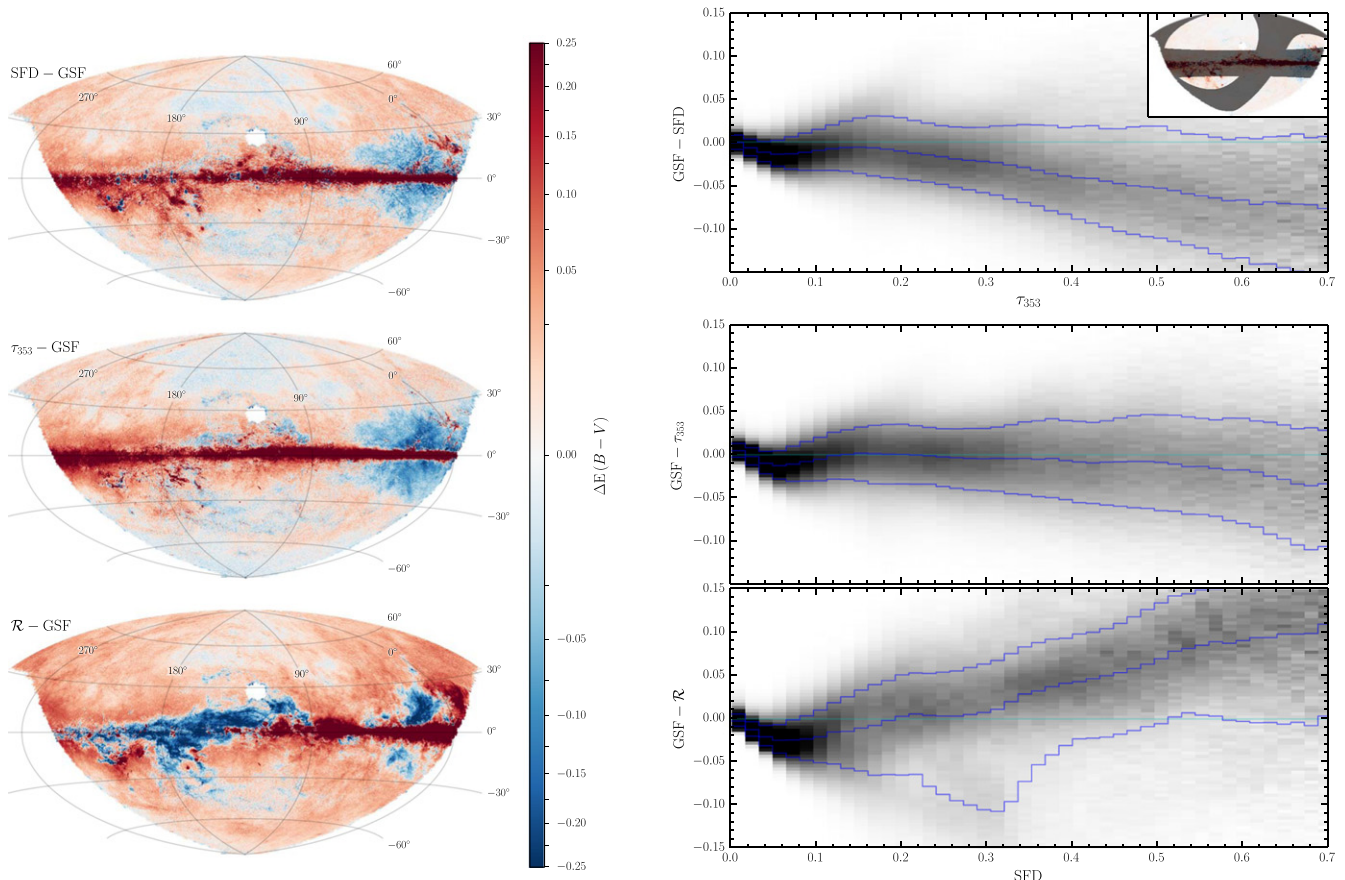
Figure 12. Comparison of our PS1-based reddening to 5 kpc (denoted by "GSF") with SFD, the Planck $\tau_{353\,GHz}$-based reddening, and the Planck radiance-based map (denoted by $\mathcal{R}$). The left panels map the residuals (between each map and the median integrated reddening in our map) on a square-root stretch. The panels on the right show the histogram of the residuals (between each emission-based map and posterior samples from our map) as a function of $E(B-V)$. In the upper right panel, we put the $\tau_{353\,GHz}$-based reddening estimate, in magnitudes, along the x-axis, since its errors should be largely uncorrelated with the errors in the PS1 and SFD reddening estimates. In the bottom two panels on the right, we use the SFD reddening estimate, in units of magnitudes, as our proxy for reddening, likewise because its errors should be uncorrelated with those of the quantities along the y-axis. An inset in the top-right panel shows the regions that are masked in this analysis. The detailed behavior of the residuals, particularly at large reddenings, depends on which regions are masked, indicating that there are systematic differences in the residuals between our reddening map and emission-based map in different regions of the sky.

while in the Marshall map, it is spread over all the distance bins closer than $\sim$3 kpc, and is therefore not cleanly separated from superimposed dust structures at greater distances. The greater distance resolution of our map at nearby distances is most likely due to the fact that we use both main-sequence stars and giants, while Marshall et al. (2006) relies solely on giants, which are saturated nearby and form a larger fraction of the observable stellar population at greater distances.

In addition to better distance resolution in the first few kiloparsecs, the greater source density of PS1 relative to 2MASS allows us to achieve better angular resolution than the Marshall map. Figure 15 demonstrates the difference in angular resolution between the two maps. In most of the region of overlap between our reddening map and the Marshall map, we achieve an angular resolution of 3$'$.4, as opposed to the constant 15$'$ resolution of the Marshall map. This allows us to resolve detailed filamentary structure not seen in the latter map. As dust reddening can vary significantly on small angular scales, this increased angular resolution will be important in practice for correctly de-reddening extinguished sources in regions with complex dust structure.

Deep in the plane of the Galaxy, where high extinction reduces source counts, our finer angular resolution limits the distance to which our map can trace dust, compared to the

Marshall map. Although the PS1 3π survey is deeper than 2MASS, the 2MASS passbands are less affected by dust extinction, and the advantage of PS1 decreases in regions of high extinction. In such regions, the greater depth of PS1 does not fully compensate for our smaller pixels, limiting the depth to which we trace dust deep in the Galactic plane. The fact that our map derives its most accurate reddening information from main-sequence stars also limits the maximum extinction to which it is reliable.

5.3. Lallement et al. (2013)

Combining distance and reddening estimates for $\sim$23,000 stars with the assumption of spatial correlation in dust density, Lallement et al. (2013) infer reddening in 3D out to a distance of $\sim$800 pc from the Sun. We find close morphological agreement between our 3D dust map and that of Lallement et al. (2013), with some differences which are worth taking note of.

In Figure 16, we show the distribution of dust and stars in a slice 25 pc above the Galactic plane, level with the Sun. We show the median dust density. In order to generate the stellar locations, we draw a sample at random from the improved distance posterior of each star (see Section 2.4), and select
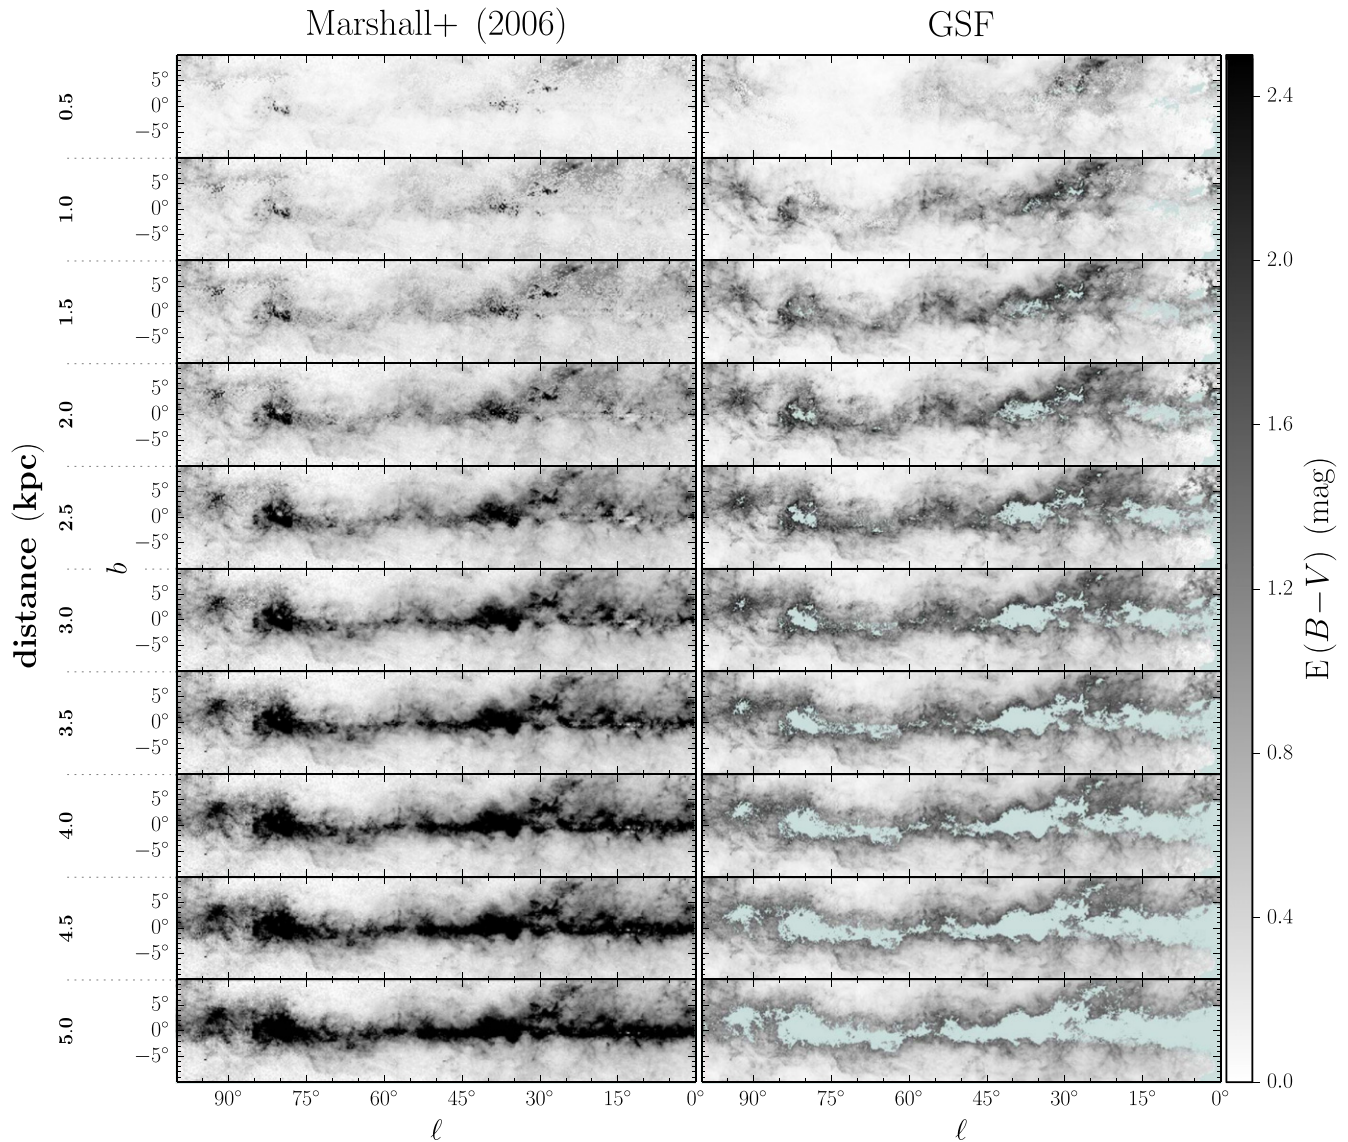
Figure 13. Comparison of the median cumulative reddening out to increasing distances in the Marshall map (left panel) and our map (right panel). Regions beyond the maximum reliable distance in our map are masked out in blue in the right panel. At large distances, the two maps agree qualitatively, with the masked regions in our map corresponding to the most heavily obscured regions in the Marshall et al. (2006) map. The Marshall map has greater depth, but lower angular resolution, and lower distance resolution in the nearest two to three kiloparsecs.

only those stars that lie within 5 pc of the chosen plane. For display purposes, we only display one out of every thousand stars.

Just like Lallement et al. (2013), we find cavities in the dust density in the directions of $\ell = 70°$ and $225°$. The overall morphology of the dust structure in the right planel of Figure 16 matches that of Figure 1 in Lallement et al. (2013).

The most obvious difference between our maps and those of Lallement et al. (2013) is the different voxel shapes employed in our work and theirs. Lallement et al. (2013) use small cubic voxels, and assume a spatial correlation function that favors similar dust densities in nearby voxels. This allows them to densely sample the reddening distribution, with voxels that are not directly constrained by stellar reddening measurements being constrained by neighboring voxels. In contrast, we infer the dust distribution in each line of sight separately, without assuming spatial correlations in the dust density, as laid out in Green et al. (2014). While our map has excellent angular

resolution, it has distance bins with a width of about 25%, giving our voxels their pencil-beam shape.

Although we see roughly the same structures, such as cavities in the reddening distribution centered on $\ell = 70°$ and $225°$, the distances we derive for a number of clouds is greater than the distances Lallement et al. (2013) find. In particular, while Lallement et al. (2013) place the Cygnus rift between 500 and 600 pc, we place it at a distance of 800–1000 pc.

6. ACCESSING THE MAP

Our 3D dust map can be accessed at http://argonaut.skymaps.info. The website provides an interface for querying individual lines of sight, as well as the ability to download the entire map and software to read it. We also provide an API through the website, which allows users to query the map remotely with a few lines of code, without the need to download the entire data cube. The data is also accessible
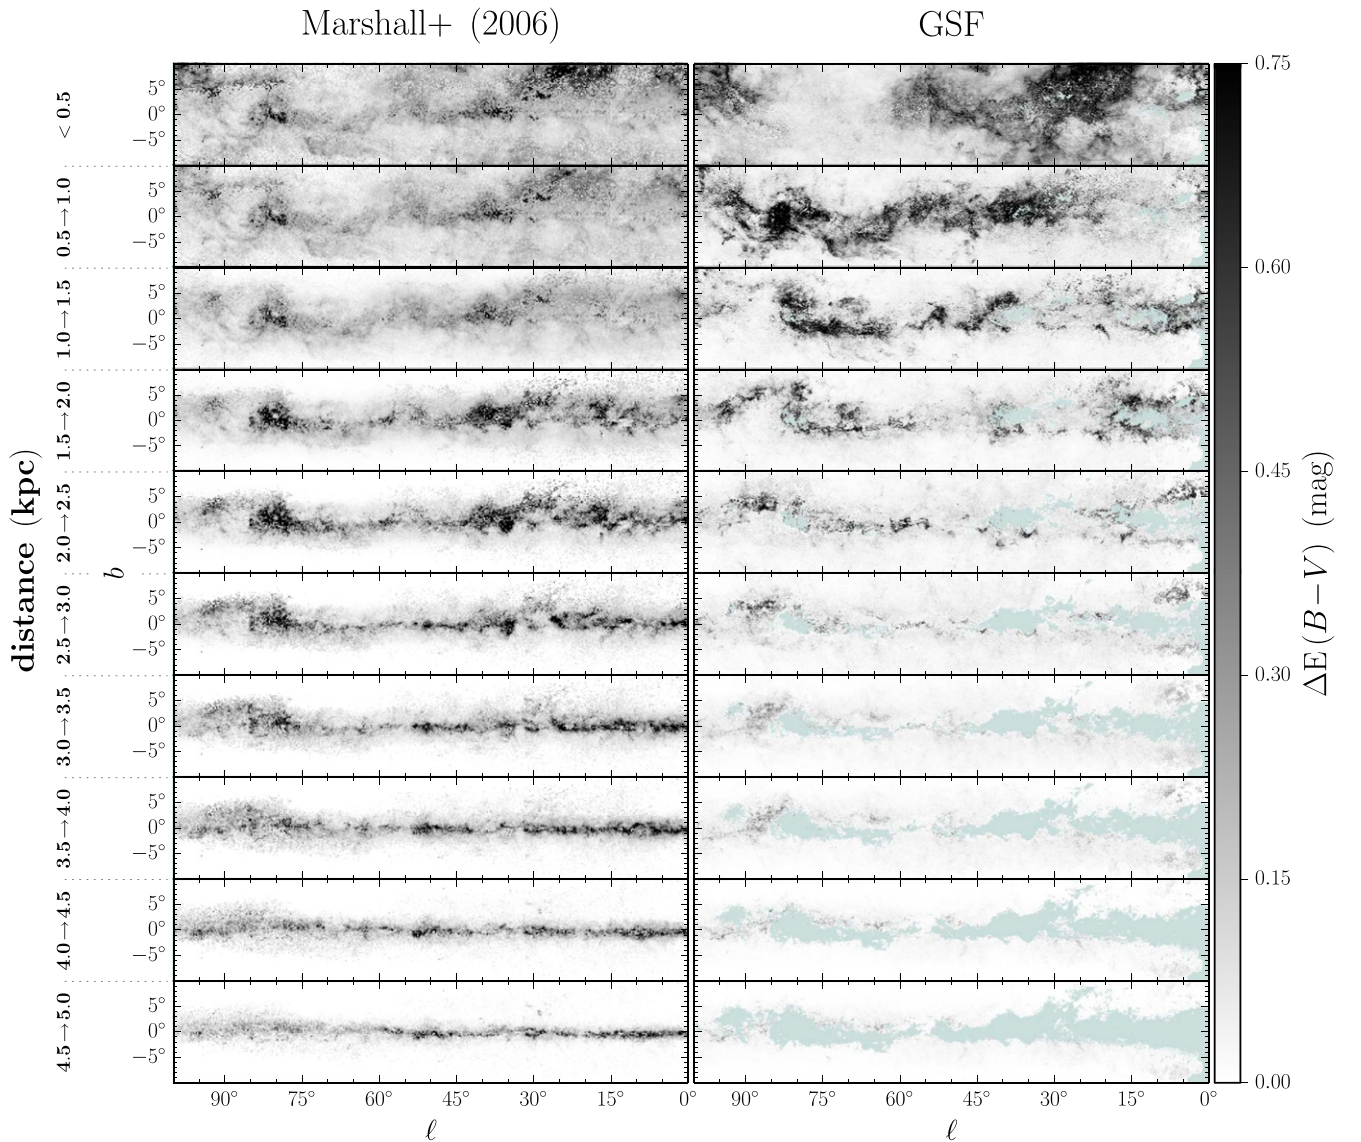
Figure 14. Comparison of the median differential reddening in bins of increasing distance in the Marshall map (left panel) and our map (right panel). In the nearest two to three kiloparsecs, we achieve much better distance resolution, as evidenced by the differentiated structures visible in successive distance bins.

at http://dx.doi.org/10.7910/DVN/40C44C, through the Harvard Dataverse.

6.1. Data Cube

The basic data product our map contains is samples of the differential reddening in 31 distance bins, in each of 2.4 million pixels. The data structure is thus

$$\left(2.4 \times 10^6 \text{ pixels}\right) \times (500 \text{ samples}) \times (31 \text{ distance bins}).$$

This allows us to determine the probability density of the cumulative reddening to any distance (within a few kiloparsecs) in the 3π steradians covered by the PS1 survey.

We encourage users to use the Markov chain samples of reddening versus distance provided by our interface in their analyses, as the samples contain the full statistical information generated by our method. These samples can be queried using our web API, downloaded as an ASCII table for individual lines of sight, or accessed directly in the complete data cube provided for download.

An additional, and larger, data set that we produce while generating the 3D dust map is a library of photometrically determined stellar parameters. As described in Green et al. (2014), we infer the probability distribution of the distance modulus, reddening, metallicity, and absolute r_{P1}-magnitude for each star. We thus have a second data cube with shape

$$\left(8 \times 10^8 \text{ stars}\right) \times (100 \text{ samples}) \times (4 \text{ parameters}).$$

In addition, we store quality assurance information for each star, including whether the Markov Chain converged during the fitting procedure, and the Bayesian evidence for the stellar model, which is similar to the χ^2 statistic in maximum-likelihood fitting. Point sources with poor evidence are likely either of stellar types not contained in our model, such as very young stars or binary systems, or are not stars (e.g., white dwarfs, quasars, unresolved galaxies).
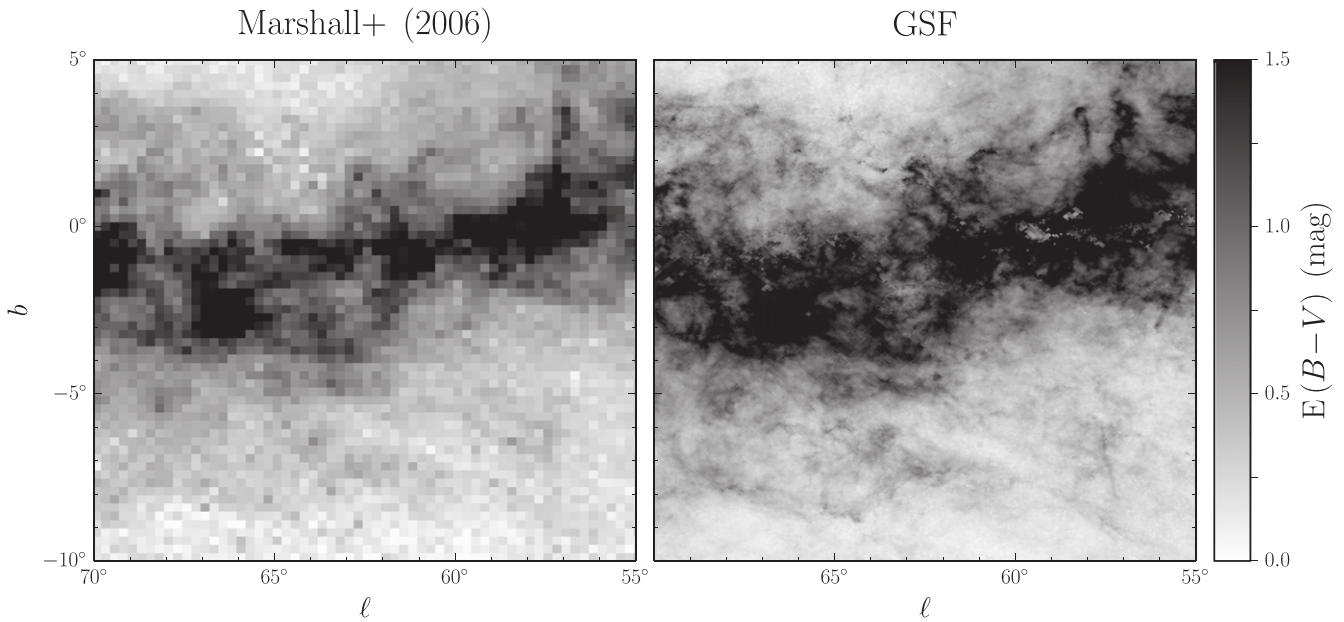
Figure 15. Zoom-in of median cumulative reddening to 3 kpc, showing the difference in angular resolution between the Marshall 3D reddening map (left panel) and our map (right panel), as well as the lower noise of the latter. The Marshall dust map has an angular resolution of 15′, while our dust map has a typical angular resolution of 3′.4 in the region of the Galactic plane displayed above. The reliability mask has not been applied our map in this figure.
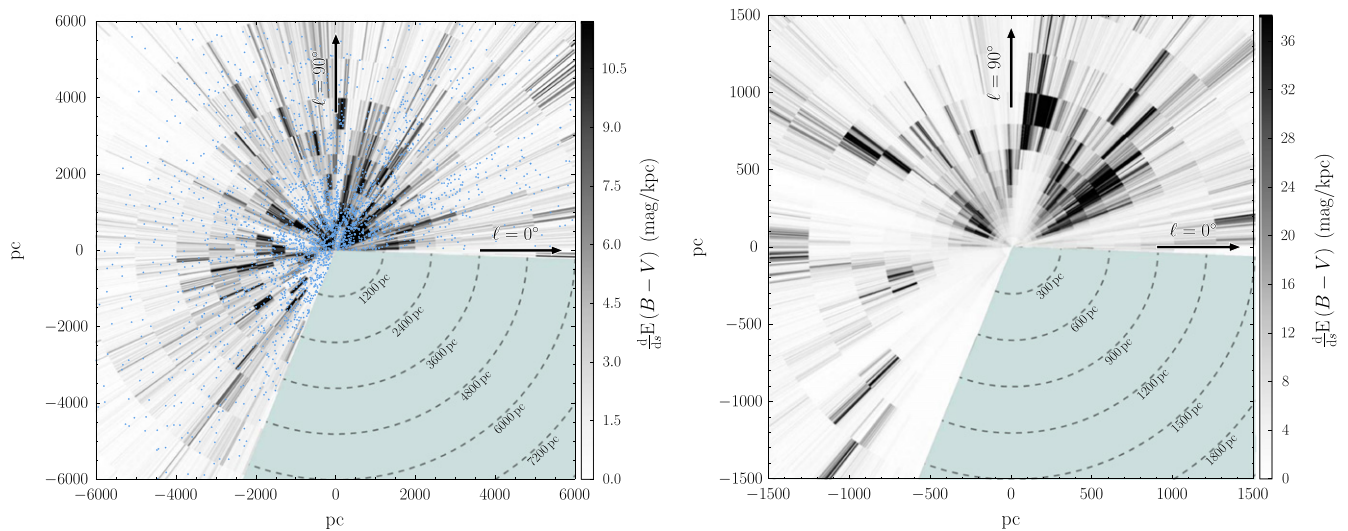

Figure 16. Dust density in a 10 pc-thick slice lying 25 pc above the Galactic plane (i.e., level with the Sun), with positions of stars within this slice overplotted in the left panel. In detail, we show the median (over multiple realizations of the 3D dust map) of the reddening column density along each sightline passing through the slice (i.e., into the page). The Sun is at the origin of each panel, with the right panel being a zoomed-in version of the left (without the overplotted stars). Only every 1000th star has been plotted. The dust map is reliable out to distances at which the stellar density becomes too small to trace the dust column. Compare with Figure C2 of Lallement et al. (2013), in which the stellar density is concentrated within ~200 pc of the Sun.

7. CONCLUSIONS

We have presented a three-dimensional map of dust reddening covering three quarters of the sky, based on photometric inferences for ~800 million stars with high-quality multiband photometry in PS1. The map provides a window into the structure of the interstellar medium, revealing detailed structure from the smallest scales in our map, 3′.4, all the way to large cloud complexes spanning many degrees. We provide interfaces to query and download the map at http://argonaut. skymaps.info.

Projecting our map down to two dimensions, we find good agreement with emission-based dust maps at high Galactic latitudes, where we expect our method to trace the entire dust

column. Comparison with the 3D map of nearby dust presented in Lallement et al. (2013) shows the same morphological features. In comparison with the 3D dust map of the inner Galactic plane presented in Marshall et al. (2006), our map has greater angular resolution and superior distance resolution within ~3 kpc. However, due to our reliance primarily on main-sequence stars, as opposed to giants, our map does not penetrate to as great a depth as Marshall et al. (2006).

Our dust mapping technique can be extended to take advantage of photometric surveys beyond PS1 and 2MASS. The Dark Energy Survey (The Dark Energy Survey Collaboration 2005) is surveying 5000 deg² of the southern sky, largely complementary to the PS1 footprint, in a similar filter set. The

LSST (Ivezic et al. 2008) will provide deep *ugrizy* photometry for the sky south of $\delta \approx 34°.5$. In the nearer future, the ESA *Gaia* mission (Lindegren et al. 1994) will provide multiband photometry, geometric parallax distance measurements, and proper motions for one billion stars. Parallax distances, where available, will vastly improve stellar distance estimates, breaking the dwarf-giant degeneracy in particular. Kinematic information from *Gaia* will provide additional information about stellar distances and types.

A reliable map the 3D distribution of dust is important for studies of stellar populations within the Galaxy, as well as streams and the global morphology of the Milk Way. In future work, we will leverage the stellar inferences produced as a by-product of our map to study the morphology of the Galaxy. We expect that the three-dimensional dust map presented here will find many different uses not yet envisioned by the authors.

APPENDIX A
PS1 AND 2MASS STELLAR TEMPLATES

Green et al. (2014) describes how stellar templates are compiled for the PS1 passbands. Briefly, metallicity-independent main-sequence stellar colors were obtained by fitting a stellar locus in color–color space, and metallicity-dependent absolute magnitudes were obtained from the metallicity-dependent photometric parallax relation given in Ivezić et al. (2008). For the giant branch, linear fits to globular cluster color–magnitude diagrams from Ivezić et al. (2008) were used.

In order to include 2MASS photometry in our data set, we require joint PS1-2MASS stellar templates. We compile these templates using a nearly identical procedure as in Green et al. (2014). We begin by selecting ∼1 million stars with $E(B-V)_{SFD} < 0.1$ mag, detections in all PS1 and 2MASS passbands, and photometric errors less than 0.5 mag in every passband. The resulting sample has a median reddening of 0.016 mag in $E(B-V)$. After dereddening the photometry, we fit a stellar locus in seven-dimensional color space, using the algorithm laid out in Newberg & Yanny (1997). The resulting stellar locus is plotted in Figure 17.

As before, we apply a metallicity-dependent photometric parallax relation to obtain a set of stellar templates, indexed by an absolute r_{P1}-magnitude and [Fe/H]. In order to obtain templates for the giant branch which incorporate 2MASS, we again use the templates from Ivezić et al. (2008), replacing stellar colors with our new joint PS1-2MASS stellar locus colors. We use $r_{P1} - i_{P1}$ to match each giant template to a color template in our PS1-2MASS stellar locus.

In the 2MASS passbands, we adopt the reddening coefficients $R_J = 0.786$, $R_H = 0.508$, and $R_{K_s} = 0.320$, calculated by Yuan et al. (2013) for a 7000 K source spectrum at $E(B-V) = 0.4$ mag, using the Cardelli et al. (1989) reddening law and assuming $R_V = 3.1$.

APPENDIX B
2MASS SELECTION FUNCTION

In order to avoid Malmquist bias in photometric inferences, it is necessary to characterize the survey selection function. In Green et al. (2014), we characterized this as the detection probability in each passband as a function of apparent magnitude and position on the sky. Denote the detection or non-detection of a point source on a particular location on the sky in passband i as a binary variable, S_i. Denote the intrinsic (or modeled) apparent magnitude of the star in passband i as $m_{mod,i}$. The detection probability is written as

$$p\left(S_i \middle| m_{mod,i}\right). \qquad (16)$$

For point-source detections in PS1, we modeled the detection probability based on the sky background, read noise, and point-spread function. As the 2MASS point-source catalogs do not contain this information, we model the 2MASS selection function empirically, from the histogram of detections as a function of apparent magnitude across the sky.

There are two ways to approach this problem. The first approach is to construct a full forward model, drawing stellar types, locations, and reddenings from our Galactic and stellar priors, generating model photometry for the simulated stars, and applying a trial selection function and photometric errors to obtain a sample of simulated "observed" apparent magnitudes. One would then vary the selection function until the histogram of detections versus observed apparent magnitude matched observed histogram in a given region of the sky. This method suffers from its reliance on our relatively crude priors on Galactic reddening, and its sensitivity to any errors in our Galactic and stellar priors.

We therefore opt for a simpler approach. We make the assumption that near the limiting magnitude, the true sky density of objects is a smooth function of apparent magnitude. In small patches of the sky, we locate the turnoff in the histogram of detections, $m_{to,i}$, as a function of apparent magnitude. In the range $-3.5 < m_i - m_{to,i} < -0.1$, we model the logarithm of the number of detections in each apparent magnitude bin as a third-order polynomial in magnitude. This is our smooth estimate of the intrinsic sky density of sources as a function of apparent magnitude. For each 2MASS passband,
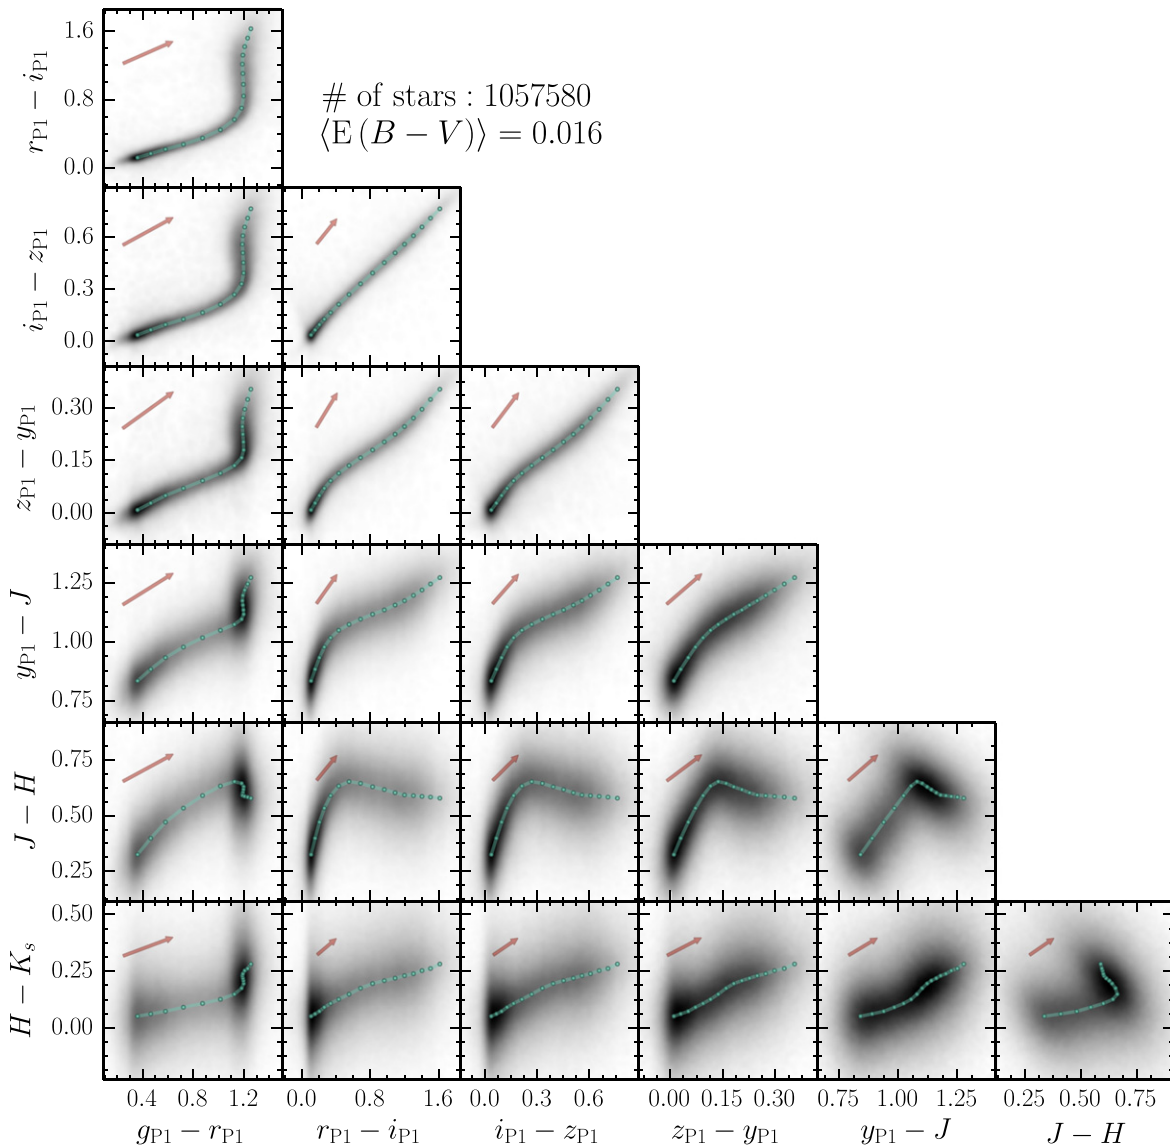
Figure 17. Stellar locus fit in joined PS1–2MASS color space. For this fit, ~1 million stars in the vicinity of the north Galactic Pole, with SFD reddening less than 0.1 mag in $E(B-V)$, were used. The fitted stellar locus anchor points are plotted over the density of stars in each color–color projection. A reddening vector with magnitude 0.5 mag in $E(B-V)$ is overplotted for each combination of colors.

we define the limiting magnitude as the bin in which the observed sky density per unit magnitude falls below 50% of the estimated true sky density per unit magnitude.

We use this procedure to construct maps of the 2MASS J, H, and K_s limiting magnitudes at two HEALPix resolutions, `nside = 32` and 64. In each region of the sky, we adopt the `nside = 64` map, unless the given pixel contains fewer than 1000 detections in the given passband, in which case we switch to the lower resolution, `nside = 32` map of limiting magnitude. Figure 18 shows the resulting, multi-resolution map of limiting magnitude in each passband.

APPENDIX C
MARKOV CHAIN "SWAP" PROPOSALS

Because of the way in which we parameterize the line of sight distance versus reddening profile, there can be strong anticorrelations between the different parameters. The differential reddening in one distance bin is often strongly anticorrelated with the differential reddening in neighboring distance bins, because transferring some differential reddening from one bin to a neighboring bin keeps the cumulative reddening constant in following distance bins. Similarly, if the data is well fit by a large jump in reddening in one bin, it is often also well fit by a jump in reddening in a neighboring bin. These strong anticorrelations between the differential reddening in neighboring distance bins can slow down Markov chain convergence when fitting the line of sight reddening profile.

We therefore introduce a new type of MCMC proposal step, which we term the "swap" proposal. In order to generate a proposal state for the Markov chain, we swap the differential reddening in two distance bins, as shown in Figure 19.

The swap proposal is symmetric, allowing the usual Metropolis–Hastings acceptance probability to be used. Call the current state X and the proposal state Y, and denote the two distance bins that were swapped to generate Y as i and j. The
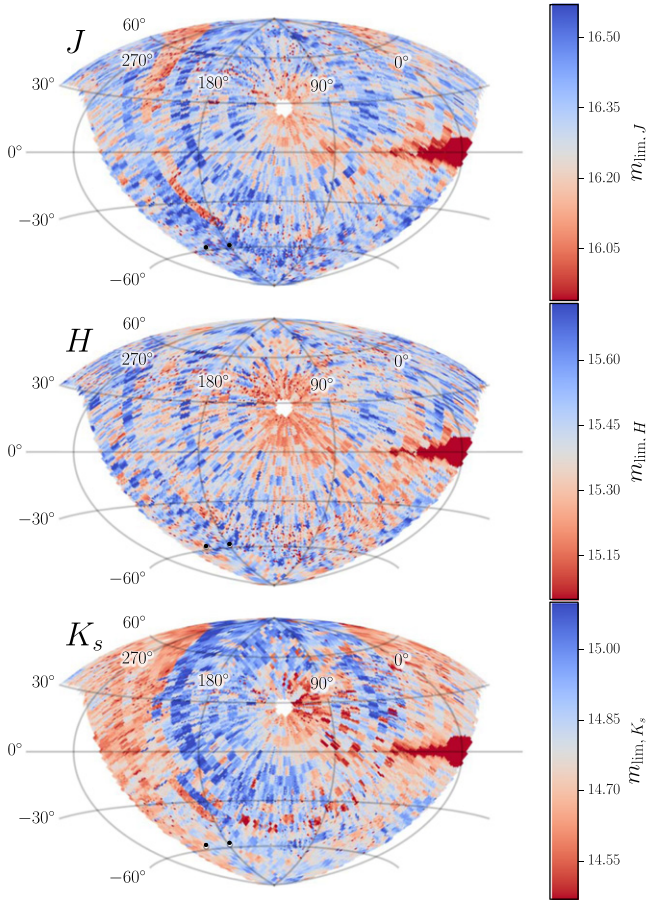
Figure 18. Limiting magnitude of the 2MASS survey in all three passbands, across the PS1 survey footprint. The limiting magnitude is defined as the apparent magnitude at which 50% of point sources are detected. Effective survey depth is relatively uniform over most of the sky, but falls off precipitously toward the Galactic Center due to source crowding.
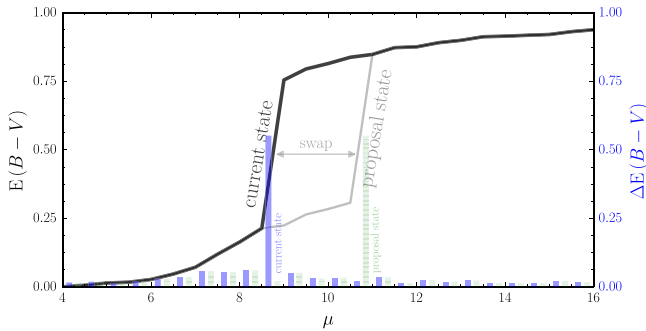

Figure 19. Generation of an MCMC proposal state by the "swap" proposal method. The proposal state is identical to the current state, up to a swap in differential reddening in two bins. The black curve shows the cumulative reddening of the current state, while the light gray curve shows the cumulative reddening of the proposal state. The blue bars show the current differential reddening in each distance bin, while the striped green bars show the proposed differential reddening in each distance bin. Note that the two are identical, up to a swap between two distance bins.

probability of proposing Y, $Q(X \rightarrow Y)$, is simply the probability that bins i and j are selected to be swapped. As long as the choice of i and j has nothing to do with the current state of the chain, the reverse step, from Y to X, is equally likely. That is, $Q(Y \rightarrow X)$ is also simply equal to the probability that distance bins i and j are selected to be

swapped. The acceptance probability, $A(X \rightarrow Y)$, is therefore given by

$$A(X \rightarrow Y) = \min\left(1, \ \frac{p(Y)}{p(X)} \frac{Q(Y \rightarrow X)}{Q(X \rightarrow Y)}\right)$$

$$= \min\left(1, \ \frac{p(Y)}{p(X)}\right), \tag{17}$$

where $p(X)$ and $p(Y)$ are the probability densities of X and Y, respectively. This is simply the usual Metropolis–Hastings acceptance probability.

The addition of swap proposals, alongside Metropolis–Hastings proposals and affine stretch proposals, allows the Markov chain to mix more quickly. Many distance-reddening curves for the sightline plotted in Figure 3, for example, differ primarily in which distance bin the large jump in reddening occurs. In such a sightline, the addition of swap proposals allows the Markov chain to transition quickly between probable states, greatly improving convergence times.

APPENDIX D
SCATTER IN THE LINE OF SIGHT REDDENING PROFILE

As sketched out in 2.2, we allow each star to deviate slightly from the line of sight distance–reddening relation in the pixel. Allowing stars at the same distance to have somewhat different reddenings renders our model more robust to sub-pixel angular variations in dust density. What follows below is a detailed derivation of how marginalization over δ_i, the deviation of star i from the average reddening profile in the pixel, translates into a Gaussian smoothing of the stellar probability density function in distance–reddening space, $p(\mu_i, E_i | \boldsymbol{m}_i)$. Beginning with the full posterior density on line of sight reddening given our stellar photometry, and including a deviation, δ_i, for each star i,

$$p(\boldsymbol{\alpha} \,|\, \{\boldsymbol{m}\}) \propto p(\boldsymbol{\alpha}) \prod_i \int d\mu_i \, d\boldsymbol{\Theta}_i \, d\delta_i \, p\left(\boldsymbol{m}_i \big| \mu_i, \boldsymbol{\Theta}_i, \boldsymbol{\alpha}, \delta_i\right)$$
$$\times \, p\left(\mu_i, \boldsymbol{\Theta}_i, \delta_i \big| \boldsymbol{\alpha}\right). \tag{18}$$

Here, $\boldsymbol{\Theta}_i$ is the type of star i. Taking just the integrand, and ignoring the subscript i for the moment,

$$I \equiv p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, \boldsymbol{\alpha}, \delta) p(\mu, \boldsymbol{\Theta}, \delta|\boldsymbol{\alpha}) \tag{19}$$

$$= p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, \boldsymbol{\alpha}, \delta) p(\mu, \boldsymbol{\Theta}|\boldsymbol{\alpha}) p(\delta|\mu, \boldsymbol{\Theta}, \boldsymbol{\alpha}). \tag{20}$$

As in Green et al. (2014), we assume that

$$p(\mu, \boldsymbol{\Theta}|\boldsymbol{\alpha}) = p(\mu, \boldsymbol{\Theta}), \tag{21}$$

with the only complication being the survey completeness limit, which is dealt with explicitly in that paper, but which has no impact on the modification being discussed here. We also assume that

$$p(\delta|\mu, \boldsymbol{\Theta}, \boldsymbol{\alpha}) = p(\delta|\mu, \boldsymbol{\alpha}), \tag{22}$$

since δ is a property of the sub-pixel variation in dust density, and should be unrelated to stellar type. Then,

$$I = p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, \boldsymbol{\alpha}, \delta) p(\mu, \boldsymbol{\Theta}) p(\delta|\mu, \boldsymbol{\alpha}). \tag{23}$$

The likelihood term (the first term on the right-hand side (rhs)), can be rewritten as

$$p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, E(\mu; \boldsymbol{\alpha}, \delta)), \qquad (24)$$

since the modeled apparent magnitude is simply determined by the stellar distance, type, and reddening, and the individual stellar reddening is determined by the line of sight reddening profile, the stellar distance, and the fractional deviation of the stellar reddening from the local reddening.

By Bayes' rule, the product

$$p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, E(\mu; \boldsymbol{\alpha}, \delta))p(\mu, \boldsymbol{\Theta}) \qquad (25)$$

is proportional to the posterior density on distance, reddening and stellar type for an individual star, in the presence of a flat prior on reddening. Thus,

$$I \propto p(\mu, \boldsymbol{\Theta}, E|\boldsymbol{m})p(\delta|\mu, \boldsymbol{\alpha}), \qquad (26)$$

evaluated at $E = (1 + \delta)E(\mu; \boldsymbol{\alpha})$. After marginalizing over stellar type, Θ, we are left with

$$\int d\Theta\, I \propto p(\mu, E) = (1 + \delta)E(\mu; \boldsymbol{\alpha}|\boldsymbol{m})p(\delta|\mu, \boldsymbol{\alpha}). \qquad (27)$$

Recall that the prior on δ is a Gaussian centered on zero, with width determined by $E(\mu; \boldsymbol{\alpha})$. Then,

$$\int d\Theta\, I \propto p(\mu, E = (1 + \delta)E(\mu; \boldsymbol{\alpha})|\boldsymbol{m})p(\delta|E(\mu; \boldsymbol{\alpha})). \qquad (28)$$

Marginalizing over δ now leaves us with a smoothed version of the individual stellar posterior density in distance and reddening:

$$\int d\Theta\, d\delta\, I = \tilde{p}(\mu, E(\mu; \boldsymbol{\alpha})|\boldsymbol{m}), \qquad (29)$$

where we have defined the "smoothed" individual stellar posterior probability density

$$\tilde{p}(\mu, E) = \int d\delta\, p(\mu, (1 + \delta)E|\boldsymbol{m})p(\delta|E). \qquad (30)$$

In Equation (3), we therefore replace the individual stellar posterior densities with "smoothed" posterior densities:

$$p(\boldsymbol{\alpha}\,|\{\boldsymbol{m}\}) \propto p(\boldsymbol{\alpha})\prod_i \int d\mu_i\, \tilde{p}\big(\mu_i, E(\mu_i; \boldsymbol{\alpha})\,\big|\, \boldsymbol{m}_i\big). \qquad (31)$$

Our method relies on calculating the individual stellar posterior probability densities for all stars along a line of sight first, before sampling from the line of sight reddening distribution, as described in detail in Green et al. (2014). The above derivation shows that this new addition to our method, allowing stars to deviate from the line of sight reddening profile, requires an intermediate step, in which the individual stellar posterior densities are smoothed in reddening, according to Equation (30).

APPENDIX E
STELLAR SAMPLE REWEIGHTING

Here, we derive in detail how to reweight the Markov Chain samples resulting from the naive stellar inferences, in order to condition the stellar parameters on the line of sight reddening profile. This discussion complements Section 2.4.

Given a fixed line of sight reddening profile parameterized by, $E(\mu; \boldsymbol{\alpha})$, each star is described by a distance μ, stellar type

$\boldsymbol{\Theta}$, and fractional offset δ from the reddening profile. The posterior of these parameters, conditioned on the star's photometry and the line of sight reddening profile, is given by

$$p(\mu, \boldsymbol{\Theta}, \delta|\boldsymbol{m}, \boldsymbol{\alpha}) \propto p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, \delta, \boldsymbol{\alpha})p(\mu, \boldsymbol{\Theta}, \delta|\boldsymbol{\alpha}). \quad (32)$$

The second term on the rhs can be broken down into two parts,

$$p(\mu, \boldsymbol{\Theta}, \delta|\boldsymbol{\alpha}) \propto p(\mu, \boldsymbol{\Theta})p(\delta|\mu, \boldsymbol{\alpha}). \qquad (33)$$

Thus,

$$\begin{aligned} p(\mu, \boldsymbol{\Theta}, \delta|\boldsymbol{m}, \boldsymbol{\alpha}) &\propto p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, \delta, \boldsymbol{\alpha})p(\mu, \boldsymbol{\Theta}) \\ &\times p(\delta|\mu, \boldsymbol{\alpha}). \end{aligned} \qquad (34)$$

The two stellar parameterizations, $\{\mu, \boldsymbol{\Theta}, \delta, \boldsymbol{\alpha}\}$ and $\{\mu, \boldsymbol{\Theta}, E\}$, are equivalent, as the stellar distance modulus, μ, the line of sight reddening parameters, $\boldsymbol{\alpha}$, and the star's fractional offset, δ, from the line of sight reddening are sufficient to determine the stellar reddening, E. We would like to reweight the samples we store in the space $\{\mu, \boldsymbol{\Theta}, E\}$, so that they correspond to the posterior density given above, in Equation (34). As described in Green et al. (2014), the samples we store in our initial processing are drawn from the posterior density

$$p(\mu, \boldsymbol{\Theta}, E|\boldsymbol{m}) \propto p(\boldsymbol{m}|\mu, \boldsymbol{\Theta}, E)p(\mu, \boldsymbol{\Theta})p(E), \qquad (35)$$

with a flat prior on E, so that $p(E) = \text{const}$. Transforming to the parameterization $\{\mu, \boldsymbol{\Theta}, \delta, \boldsymbol{\alpha}\}$, we have to transform the flat prior $p(E)$ to an equivalent prior, $p(\delta|\mu, \boldsymbol{\alpha})$. This prior is given by

$$p(\delta|\mu, \boldsymbol{\alpha}) = p(E)\frac{\partial E}{\partial \delta}\bigg|_{\mu, \boldsymbol{\alpha}} \propto \frac{\partial E}{\partial \delta}\bigg|_{\mu, \boldsymbol{\alpha}}, \qquad (36)$$

where we have taken out the constant factor $p(E)$ on the rhs and replaced the equality with a proportionality. The stellar reddening is related to μ, $\boldsymbol{\alpha}$ and δ by Equation (6). Taking the derivative w.r.t. δ,

$$\frac{\partial E}{\partial \delta}\bigg|_{\mu, \boldsymbol{\alpha}} = \frac{\partial}{\partial \delta}\big[(1 + \delta_i)E(\mu_i; \boldsymbol{\alpha})\big] = E(\mu_i; \boldsymbol{\alpha}). \qquad (37)$$

The flat prior in E thus implies a prior on δ equal to

$$p(\delta|\mu, \boldsymbol{\alpha}) \propto E(\mu_i; \boldsymbol{\alpha}). \qquad (38)$$

In order to sample from Equation (34), we can weight the samples from our initial processing by the ratio of the prior in the new parameterization, where reddening is conditioned on the line of sight reddening profile, to the prior used in the initial sampling, where a flat prior on reddening is used. From the above calculation, this ratio is given by

$$\frac{p_{\text{new}}(\delta|\mu, \boldsymbol{\alpha})}{p_{\text{old}}(\delta|\mu, \boldsymbol{\alpha})} \propto \frac{p_{\text{new}}(\delta|\mu, \boldsymbol{\alpha})}{E(\mu_i; \boldsymbol{\alpha})}. \qquad (39)$$

The functional form of $p_{\text{new}}(\delta|\mu, \boldsymbol{\alpha})$ is given by Equation (7).

REFERENCES

Alves, J. a., Lombardi, M., & Lada, C. J. 2014, A&A, 565, A18
Andrieu, C., & Roberts, G. O. 2009, AnSta, 37, 697
Beaumont, M. a. 2003, Genetics, 164, 1139
Berry, M., Ivezić, v., Sesar, B., et al. 2011, ApJ, 757, 166
Burstein, D., & Heiles, C. 1982, AJ, 87, 1165
Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, ApJ, 345, 245
Chen, B.-Q., Liu, X.-W., Yuan, H.-B., et al. 2014, MNRAS, 443, 1192
Draine, B. 2003, ARA&A, 41, 241
Drimmel, R., & Spergel, D. 2001, ApJ, 20, 181

Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2014, ApJ, 783, 114
Hanson, R. J., & Bailer-Jones, C. a. L. 2014, MNRAS, 438, 2938
Hodapp, K. W., Kaiser, N., Aussel, H., et al. 2004, AN, 325, 636
Hogg, D., Myers, A., & Bovy, J. 2010, ApJ, 725, 2166
Ivezić, v., Sesar, B., Jurić, M., et al. 2008, ApJ, 684, 287
Ivezic, Z., Tyson, J. a., Acosta, E., et al. 2008, arXiv:0805.2366
Jurić, M., Ivezić, v., Brooks, A., et al. 2008, ApJ, 673, 864
Kaiser, N., Burgett, W., Chambers, K., et al. 2010, Proc. SPIE, 7733, 77330E
Lada, C. J., Lada, E. A., Clemens, D. P., & Bally, J. 1994, ApJ, 429, 694
Lallement, R., Vergely, J.-L., Valette, B., et al. 2013, 1
Lindegren, L., Perryman, M. A., Bastian, U., et al. 1994, Proc. SPIE, 2200, 599
Lombardi, M. 2009, A&A, 493, 735
Lombardi, M., & Alves, J. 2001, A&A, 377, 1023
Lombardi, M., Alves, J., & Lada, C. J. 2006, A&A, 454, 781
Lombardi, M., Alves, J., & Lada, C. J. 2011, A&A, 535, A16
Lombardi, M., Lada, C. J., & Alves, J. 2008, A&A, 489, 143
Lombardi, M., Lada, C. J., & Alves, J. 2010, A&A, 512, A67
Magnier, E. 2006, in The Advanced Maui Optical and Space Surveillance
    Technologies Conf., ed. D. B. Ryan (Maui, HI: Maui Economic
    Development Board), 50
Magnier, E. 2007, in ASP Conf. Ser. 364, The Future of Photometric,
    Spectrophotometric and Polarimetric Standardization, ed. C. Sterken (San
    Francisco, CA: ASP), 153
Magnier, E. A., Liu, M., Monet, D. G., & Chambers, K. C. 2008, in IAU Symp.
    248, A Giant Step: From Milli- to Micro-arcsecond Astrometry, ed.
    W. J. Jin, I. Platais & M. A. C. Perryman (Paris: IAU), 553
Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., & Picaud, S. 2006,
    A&A, 453, 635

Newberg, H., & Yanny, B. 1997, ApJS, 113, 89
Onaka, P., Tonry, J. L., Isani, S., et al. 2008, Proc. SPIE, 7014, 70140
Peek, J. E. G., & Graves, G. J. 2010, ApJ, 719, 415
Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2014, A&A, 571, A11
Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, A&A, 409, 523
Rowles, J., & Froebrich, D. 2009, MNRAS, 395, 1640
Sale, S. E. 2012, MNRAS, 427, 2119
Sale, S. E. 2015, arXiv:1505.07626
Sale, S. E., Drew, J. E., Barentsen, G., et al. 2014, MNRAS, 443, 2907
Sale, S. E., & Magorrian, J. 2014, MNRAS, 445, 256
Schlafly, E. F., & Finkbeiner, D. P. 2011, ApJ, 737, 103
Schlafly, E. F., Finkbeiner, D. P., Jurić, M., et al. 2012, ApJ, 756, 158
Schlafly, E. F., Finkbeiner, D. P., Schlegel, D. J., et al. 2010, ApJ,
    725, 1175
Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2014a, ApJ, 799, 29
Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2014b, ApJ, 786, 29
Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2014c, ApJ, 789, 15
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163
Stubbs, C. W., Doherty, P., Cramer, C., et al. 2010, ApJS, 191, 376
The Dark Energy Survey Collaboration 2005, IJMPA, 20, 3121
Tonry, J., & Onaka, P. 2009, in Advanced Maui Optical and Space
    Surveillance Technologies Conf., ed. D. B. Ryan (Maui, HI: Maui
    Economic Development Board), 40
Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, ApJ, 750, 99
Yanny, B., Rockosi, C., Newberg, H. J., et al. 2009, AJ, 137, 4377
York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, AJ, 120, 1579
Yuan, H.-B. B., Liu, X.-W. W., & Xiang, M.-S. S. 2013, MNRAS,
    430, 2188