

MEASURING DISTANCES AND REDDENINGS FOR A BILLION STARS: TOWARD A 3D DUST MAP FROM PAN-STARRS 1

GREGORY MAURICE GREEN¹, EDWARD F. SCHLAFLY², DOUGLAS P. FINKBEINER¹, MARIO JURIĆ³, HANS-WALTER RIX²,
WILL BURGETT⁴, KENNETH C. CHAMBERS⁴, PETER W. DRAPER⁵, HEATHER FLEWELLING⁴, ROLF PETER KUDRITZKI⁴,
EUGENE MAGNIER⁴, NICOLAS MARTIN⁶, NIGEL METCALFE⁵, JOHN TONRY⁴, RICHARD WAINSCOT⁴, AND CHRISTOPHER WATERS⁴

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA; ggreen@cfa.harvard.edu

² Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

³ LSST Corporation, 933 North Cherry Avenue, Tucson, AZ 85721, USA

⁴ Institute for Astronomy, University of Hawaii at Manoa, Honolulu, HI 96822, USA

⁵ Department of Physics, Durham University, South Road, Durham DH1 3LE, UK

⁶ Observatoire Astronomique de Strasbourg, 11 rue de l'Université, F-67000 Strasbourg, France

Received 2013 September 12; accepted 2014 January 5; published 2014 February 21

ABSTRACT

We present a method to infer reddenings and distances to stars based only on their broad-band photometry, and show how this method can be used to produce a three-dimensional (3D) dust map of the Galaxy. Our method samples from the full probability density function of distance, reddening, and stellar type for individual stars, as well as the full uncertainty in reddening as a function of distance in the 3D dust map. We incorporate prior knowledge of the distribution of stars in the Galaxy and the detection limits of the survey. For stars in the Pan-STARRS 1 (PS1) 3π survey, we demonstrate that our reddening estimates are unbiased and accurate to ~ 0.13 mag in $E(B - V)$ for the typical star. Based on comparisons with mock catalogs, we expect distances for main-sequence stars to be constrained to within $\sim 20\%$ – 60% , although this range can vary, depending on the reddening of the star, the precise stellar type, and its position on the sky. A later paper will present a 3D map of dust over the three quarters of the sky surveyed by PS1. Both the individual stellar inferences and the 3D dust map will enable a wealth of Galactic science in the plane. The method we present is not limited to the passbands of the PS1 survey but may be extended to incorporate photometry from other surveys, such as the Two Micron All Sky Survey, the Sloan Digital Sky Survey (where available), and in the future, LSST and *Gaia*.

Key words: dust, extinction – Galaxy: structure – ISM: structure – methods: statistical – stars: distances

Online-only material: color figures

1. INTRODUCTION

A long-standing goal of astronomy has been to understand the structure and formation of galaxies. Studies of external galaxies have begun to paint a detailed, global picture of the forces at work in shaping galaxies but lack the resolution and sensitivity to probe individual stars. In the Milky Way, meanwhile, measurements of the positions and types of millions of stars have been assembled, though these stars probe only a fraction of our Galaxy's volume. Moreover, the positions of these stars are especially uncertain in the Galactic disk, where the bulk of the stars reside, owing to the presence of dust, which obscures and reddens the light from these stars.

Accordingly, wide-field surveys like the Sloan Digital Sky Survey (SDSS; York et al. 2000), which observed millions of stars, have focused on the structure of the stars at high Galactic latitudes, mostly outside the Galactic disk (e.g., Jurić et al. 2008; Ivezić et al. 2008a). These studies have revealed an abundance of substructure in the Galactic halo (Belokurov et al. 2006) and have led to new constraints on the structure of the Galaxy's halo (e.g., Law & Majewski 2010) and the identification of challenges to the standard picture of Galaxy formation (e.g., the Missing Satellites Problem; Simon & Geha 2007). Still, the bulk of the Galaxy's stars reside and formed in the disk, and it is unclear how much the Galaxy's halo can inform the processes at work there. Bovy et al. (2012a, 2012b, 2012c) derive smooth models for the Galactic disk, using subpopulations observed in the Sloan Extension for Galactic Understanding and Exploration (SEGUE; Yanny et al. 2009). Better photometric distance

estimates for large, magnitude-limited samples of heavily dust-obscured stars will aid investigation into the smooth structure of the disk, as well as possible disk substructure.

This paper presents a technique to simultaneously infer the distances and reddenings to stars embedded in dust, to enable study of the properties and structure of our Galaxy's disk from optical surveys of resolved stars. We exploit prior knowledge of the types and distribution of stars in the Galaxy in a Bayesian framework to deliver the full probability density function of distance and reddening and stellar type for each star. We derive a principled Bayesian technique to infer reddening as a function of distance, using stars as tracers of the dust column.

We are not the first in this area. Marshall et al. (2006) produced a three-dimensional (3D) extinction map of the Galactic plane by comparing the Two Micron All Sky Survey (2MASS) $J - K_s$ stellar colors to those of simulated catalogs based on the Besançon model of the Galaxy (Robin et al. 2003). Our work is most similar to that of Sale (2012), who presents related techniques and applies them to simulated IPHAS data (Drew et al. 2005). The method presented here for obtaining stellar reddenings and distances is also related to that of Berry et al. (2011), who highlight the large amount of 3D structure in the Galaxy's dust using data from SDSS. The work of Bailer-Jones (2011), likewise, presents a similar technique using broad-band photometry and *Hipparcos* parallaxes (Perryman et al. 1997). Vergely et al. (2001) and Lallement et al. (2003) map out the 3D distribution of clouds in the Local Bubble by measuring absorption lines imprinted by the interstellar medium on the spectra of stars with *Hipparcos* parallaxes. Lallement et al. (2013)

uses $\sim 23,000$ stellar parallaxes and reddening estimates from a number of sources to infer the 3D distribution of dust opacity out to 800–1000 pc in the plane of the Galaxy, and ~ 300 pc out of the plane. Our work differentiates itself from these primarily in that it is adapted to studying the approximately one billion stars with high-quality Pan-STARRS 1 (PS1) photometry, which cover three quarters of the sky and two thirds of the Galactic plane, representing an unprecedented resource for studies of the Galaxy’s disk.

The paper is organized as follows. Section 2 describes the PS1 survey. Section 3 develops the Bayesian formalism required to produce a 3D reddening map. Sections 4.1 and 4.2 describe the stellar and Galactic models we employ. Section 5 describes the practical implementation of our model to produce a 3D map. Then, in Section 6, we conduct tests of our method with mock photometry, and in Section 7, we validate our model with real photometry.

2. PAN-STARRS 1 SURVEY

In this work, we derive the distances and reddenings to stars observed by PS1. PS1 is a 1.8 m optical and near-infrared telescope located on Mount Haleakala, Hawaii (Kaiser et al. 2010; Hodapp et al. 2004). The telescope is equipped with the GigaPixel Camera 1, consisting of an array of 60 CCD detectors, each 4800 pixels on a side (Tonry & Onaka 2009; Onaka et al. 2008). The majority of the observing time is dedicated to a multi-epoch 3π steradian survey of the sky north of $\delta = -30^\circ$ (K. C. Chambers et al., in preparation). The 3π survey observes in five passbands, g_{P1} , r_{P1} , i_{P1} , z_{P1} , and y_{P1} , which together span 400–1000 nm (Stubbs et al. 2010). The images are processed by the PS1 Image Processing Pipeline (IPP; 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 resulting homogeneous optical and near-infrared coverage of three quarters of the sky makes the PS1 data ideal for studies of the distribution of the Galaxy’s dust.

3. LINE-OF-SIGHT REDDENING PROFILE

Here, we describe the basic assumptions that we make in order to produce a 3D dust map. We show that the problem can be decomposed into two steps. In the first step, we determine the probability density function of distance and reddening for each star (see Section 4). In the second step, we use information from stars on the same small patch of sky to infer reddening as a function of distance in the given direction.

We make the assumption that stars which are close to one another in 3D space are behind the same column of dust. By grouping together stars which are close-by on the sky, we are able to use the stars as tracers of the total dust column at different distances in a particular direction on the sky. So long as the dust column does not vary significantly over small angular scales, then this is a valid assumption.

Let us denote the reddening profile along a particular line of sight by

$$E(\mu; \alpha), \quad (1)$$

where E , the color excess in some pair of passbands (e.g., $E(B - V)$), is a function of distance modulus μ , and α denotes any fitting parameters defining the reddening profile. These parameters could be, for example, the dust density in

each distance bin, and could in principle include the value of R_V , which parameterizes the wavelength dependence of the extinction law (Cardelli et al. 1989; Fitzpatrick 1999). The extinction A in any set of passbands is then assumed to be a function of the reddening E :

$$A = A(E, R_V). \quad (2)$$

Along one line of sight, denote the photometry of star i by \mathbf{m}_i and the set of all observed stellar magnitudes as $\{\mathbf{m}\}$. We wish to determine how the model parameters α for the line-of-sight reddening profile depend on the stellar photometry $\{\mathbf{m}\}$. That is, we wish to determine

$$p(\alpha | \{\mathbf{m}\}). \quad (3)$$

Using Bayes’ rule,

$$p(\alpha | \{\mathbf{m}\}) = \frac{p(\{\mathbf{m}\} | \alpha) p(\alpha)}{p(\{\mathbf{m}\})}. \quad (4)$$

The likelihood $p(\{\mathbf{m}\} | \alpha)$ is the probability density of obtaining the set of observed magnitudes $\{\mathbf{m}\}$, given the reddening profile defined by α . The likelihood of the entire set of stellar observations is just the product of the individual likelihoods, i.e.,

$$p(\{\mathbf{m}\} | \alpha) = \prod_i p(\mathbf{m}_i | \alpha). \quad (5)$$

This is just a statement that the photometry of one star does not influence the photometry of any other star. Plugging this into Equation (4), and dropping the normalizing factor $p(\{\mathbf{m}\})$,

$$p(\alpha | \{\mathbf{m}\}) \propto p(\alpha) \prod_i p(\mathbf{m}_i | \alpha). \quad (6)$$

We now introduce nuisance parameters describing the distance to and intrinsic type of each star, and then marginalize over these parameters to obtain the likelihood. We denote the distance modulus to star i by μ_i , and the parameters describing the stellar type by Θ_i . For an individual star,

$$p(\mathbf{m} | \alpha) = \int d\mu d\Theta p(\mathbf{m}, \mu, \Theta | \alpha) \quad (7)$$

$$= \int d\mu d\Theta p(\mathbf{m} | \mu, \Theta, \alpha) p(\mu, \Theta | \alpha) \quad (8)$$

$$= \int d\mu d\Theta p(\mathbf{m} | \mu, \Theta, E(\mu; \alpha)) p(\mu, \Theta). \quad (9)$$

In the last step, we have assumed that the joint prior on the distance and intrinsic stellar type are independent of the reddening profile. Up to a normalizing constant, the above integrand is equivalent to the posterior density

$$p(\mu, E, \Theta | \mathbf{m}) \quad (10)$$

for an individual star, where the prior on E is flat. Intuitively, this is because the prior on reddening is on the fitting parameters α , rather than the reddening of an individual star. After marginalizing over the stellar type Θ ,

$$p(\mathbf{m} | \alpha) \propto \int d\mu p(\mu, E(\mu; \alpha) | \mathbf{m}). \quad (11)$$

Plugging the above into Equation (6), we find that the full posterior density for α is given by

$$p(\alpha | \{m\}) \propto p(\alpha) \prod_i \int d\mu_i p(\mu_i, E(\mu_i; \alpha) | m_i), \quad (12)$$

where we have defined the function

$$p(\mu_i, E_i | m_i) \equiv \frac{1}{Z_i} \int d\Theta_i p(m_i | \mu_i, \Theta_i, E_i) p(\mu_i, \Theta_i), \quad (13)$$

which is equivalent to the posterior probability density of finding a single star at distance μ_i and reddening E_i , where the prior on reddening is flat. Here, Z_i is a normalizing constant. Effectively, it is the Bayesian evidence for star i , a measure of how likely the star is to be drawn from the model. A higher evidence indicates that the data is more consistent with the model, while a low value for Z_i indicates that the model does not well describe the object i . Our strategy for sampling from a posterior of the form given by Equation (12) is to pre-compute Equation (13) for each star by Markov chain Monte Carlo (MCMC) sampling and then to sample from $p(\alpha | \{m\})$. This approach has two advantages. First, it factorizes the full problem into a series of smaller problems of lower dimension, potentially speeding up the computation. The second advantage of this two-step approach is that it allows outlier rejection on the basis of the Bayesian evidence for each star before proceeding to the second step. Point sources which do not fit the chosen stellar model (e.g., blue stragglers, white dwarfs, and galaxies mistakenly classified as stars) can be filtered out by imposing a cut on the evidence Z_i (see Section 5.2). More principled approaches to reducing the influence of outliers exist (e.g., Hogg et al. 2010), though in the context of our problem, they are significantly more computationally expensive to implement.

A more general limitation to the approach taken here is that it does not allow the simultaneous fitting of parameters describing the stars and the spatial variation in dust properties. One could imagine simultaneously constraining stellar types and distances, as well as the dust density and R_V parameter throughout space. By fitting the dust properties in many voxels simultaneously, one would be able to place priors on the density power spectrum of the dust and to infer R_V as a function of position in the Galaxy. By fitting dust properties throughout the entire volume of the Galaxy simultaneously, one could even attempt to constrain global parameters, such as the dust scale height and scale length. This would represent a hierarchical approach to the problem of creating a 3D dust map (see, for example, Kruschke 2011 for a discussion of hierarchical Bayesian models). At the highest level in the hierarchy, one has global parameters, which describe the overall dust distribution and density power spectrum. One level below in the hierarchy, one would have parameters describing the dust properties in each voxel in the Galaxy. At the lowest level, one could have the type and distance for each star. Such a hierarchical approach is appealing because it takes into account spatial correlations in dust properties and because it directly fits the global structure of the Galaxy's dust component. However, this hierarchical approach potentially requires much greater computational power than the approach we take in this paper, as it does not allow one to process each star individually and treat each line of sight separately, greatly increasing the dimensionality of parameter space. This paper therefore confines itself to fitting each star independently and then combining the information from each star along any given line of sight to determine the reddening profile as a function of distance.

4. INDIVIDUAL STARS

Now that we have factorized the problem of determining the line-of-sight reddening profile into one of determining $p(\mu_i, E_i, \Theta_i | m_i)$ for each star, we need to determine the individual stellar likelihoods and priors. In the following, we will drop the subscript i , as it is assumed that we are dealing with one star.

Using Bayes' Rule,

$$p(\mu, E, \Theta | m) \propto p(m | \mu, E, \Theta) p(\mu, E, \Theta). \quad (14)$$

The likelihood, $p(m | \mu, E, \Theta)$, is the probability density of a star having apparent magnitudes m , given a distance, reddening, and stellar type. The likelihood is thus dependent on our model of intrinsic stellar colors, which we discuss below (in Section 4.1). The priors, $p(\mu, E, \Theta)$, are dependent on our model of the distribution of stars of different types throughout the Galaxy. We discuss our Galactic model in Section 4.2.

4.1. Stellar Model

In our model, each star is described by two intrinsic parameters, its absolute magnitude M_r in the PS1 r band and its metallicity $[\text{Fe}/\text{H}]$. In terms of our previous notation, $\Theta = (M_r, [\text{Fe}/\text{H}])$. Given a set of stellar templates, $\mathbf{M}(M, [\text{Fe}/\text{H}])$, which map intrinsic stellar type to a set of absolute magnitudes, one obtains theoretical apparent magnitudes

$$m_{\text{mod}} = \mathbf{M}(M_r, [\text{Fe}/\text{H}]) + A(E, R_V) + \mu. \quad (15)$$

The extinction vector $A(E, R_V)$ used in this work follows a Fitzpatrick (1999) reddening law with $R_V = 3.1$, adapted to the PS1 *grizy_{p1}* filter set by Schlafly & Finkbeiner (2011). The likelihood of observing apparent magnitudes m with Gaussian uncertainties σ is then given by

$$p(m | \mu, M_r, [\text{Fe}/\text{H}], E) = \mathcal{N}(m | m_{\text{mod}}, \sigma). \quad (16)$$

In this paper, we use the notation $\mathcal{N}(x | \mu, \sigma)$ to denote the probability density of a multivariate normal with mean μ and standard deviations σ , evaluated at x .

We adopt a set of empirical stellar templates based on stellar observations in PS1, with photometric parallaxes and metallicities derived from the work of Ivezić et al. (2008a). That work determines the absolute magnitude of a star as a function of its intrinsic color and metallicity using observations of globular clusters in SDSS. These globular clusters are uniformly old, and as a result, our stellar templates are appropriate only for old populations and do not include age as a parameter. This means that young blue stars are not included and that the morphology of the subgiant and giant branches is only approximate. Accordingly, our giant branch distances are less reliable than our main sequence distances. Nevertheless, we include the giant branches in our models because any given star may indeed be a distant giant rather than a nearby dwarf.

In detail, we fit a spline to the colors of stars (in four-color space) near the North Galactic Pole to derive the shape of the stellar locus. All main-sequence stars are required to have intrinsic colors lying along this one-dimensional curve. We then associate each position along the main sequence with an absolute magnitude and a metallicity using the relations of Ivezić et al. (2008a), which give absolute magnitude as a function of color and metallicity. Models for the giants are obtained from linear fits of absolute magnitude to color and metallicity by




Figure 1. Model stellar colors as a function of absolute r magnitude and metallicity in Pan-STARRS 1 passbands. The stellar templates are based on PS1 color-color relations, and color is related to absolute magnitude and metallicity by SDSS observations of globular clusters (Ivezić et al. 2008a). Our empirical templates therefore assume an old stellar population. While the main sequence below the turnoff is nearly invariant with age, the giant branch and the location of the turnoff do, in reality, vary considerably with age. For this reason, we expect our inferences for main-sequence stars to be more accurate than those for giants. The narrowness of the kink at $M_r \simeq 2.4$ is an artifact of our models (see Section 4.1).

(A color version of this figure is available in the online journal.)

Ž. Ivezić (2011, private communication), based on observations of globular clusters. These giant branch fits are joined to the main sequence via a cubic interpolating polynomial for $4 > M_r > 2.35$. We are able to use relations derived from SDSS because of the close similarity between the PS1 and SDSS filter sets; we transform from the PS1 to SDSS colors using the color transformations of D. P. Finkbeiner et al. (in preparation), which have residuals of less than about 1% across the full range of stellar types considered in this work.

The resulting stellar templates are shown in Figure 1, which gives the templates' colors as a function of their absolute magnitude and metallicity. Our empirical approach produces a close match to the observed colors of stars, and comparison with globular and open clusters indicates that the absolute magnitudes are accurate along the main sequence (see Section 7.2). An alternative approach would have been to adopt template colors from a library of synthetic spectra, such as the Padova & Trieste Stellar Evolution Code (PARSEC; Bressan et al. 2012), giving us access to age as an additional stellar parameter. However, the synthetic libraries have difficulty reproducing the colors of M dwarfs in detail. We choose therefore to adopt a set of empirical models that match the colors of most stars well, though in future work a hybrid approach may be best suited to the problem.

4.2. Galactic Model

We now present the priors that we place on the intrinsic and extrinsic parameters describing each star. We factorize the priors

as follows:

$$p(\mu, M_r, [\text{Fe}/\text{H}]) = p(\mu)p([\text{Fe}/\text{H}]|\mu)p(M_r). \quad (17)$$

We describe the distance prior in Section 4.2.1, the metallicity prior in Section 4.2.2, and the prior on absolute magnitude in Section 4.2.3.

4.2.1. Distance

For a given line of sight, the prior probability of finding a star in a small range, $d\mu$, in distance modulus is proportional to the number of stars per unit distance modulus per unit solid angle in the direction of the pixel:

$$p(\mu) \propto \frac{dN}{d\mu d\Omega} = \frac{dN}{dr d\Omega} \frac{dr}{d\mu} = n(\mu)r^2 \frac{dr}{d\mu} \propto 10^{3\mu/5} n(\mu). \quad (18)$$

Here, r denotes physical distance from the Sun, and $n(\mu)$ is the stellar number density at distance modulus μ along the chosen line of sight. The distance prior is thus controlled by the line-of-sight number density of stars, as well as a volume factor, $10^{3\mu/5}$, which grows with distance. This latter term takes into account that the volume represented by a beam of constant, small width in distance modulus grows with distance. For a typical line of sight, the prior is driven upward by the volume factor at small distances, while farther out, it is suppressed by the decline in density in the outer reaches of the Galaxy. For constant density, $n(\mu)$, Equation (18) simply reduces to the Euclidean counts equation.

The distance prior thus requires us to calculate the stellar number density at arbitrary locations in the Galaxy. We employ the three-component Galactic model developed in Jurić et al. (2008), which comprises a thin disk, a thick disk, and an oblate halo. In cylindrical coordinates centered on the Galactic Center, and with the Galactic plane defining $Z = 0$, each disk component has number density of the form

$$n_i(R, Z) = n_{0,i} e^{-(R/R_{0,i} + |Z|/Z_{0,i})}, \quad (19)$$

where $R_{0,i}$ and $Z_{0,i}$ are the scale radius and height, respectively, of each disk component, and $n_{0,i}$ is the stellar number density of each component at the Galactic Center. In the solar neighborhood,

$$n_i(R_\odot, Z_\odot) = n_{0,i} e^{-(R_\odot/R_{0,i} + |Z_\odot|/Z_{0,i})}. \quad (20)$$

We can thus write the number density of each disk component in terms of locally defined quantities, which are more readily measurable than quantities defined at the Galactic Center:

$$n_i(R, Z) = n_i(R_\odot, Z_\odot) e^{-[(R-R_\odot)/R_{0,i} + (|Z|-|Z_\odot|)/Z_{0,i}]}, \quad (21)$$

where R_\odot and Z_\odot are the Galactocentric solar radius and height in cylindrical coordinates, respectively. From this point onward, we will denote $n_{\text{thin}}(R_\odot, Z_\odot)$ as n_\odot , and write

$$n_{\text{thick}}(R_\odot, Z_\odot) \equiv f_{\text{thick}} n_\odot. \quad (22)$$

The halo is assumed to have stellar number density

$$n_{\text{halo}}(R, Z) = n_\odot f_h \left(\frac{R_{\text{eff}}}{R_\odot} \right)^{-\eta}, \quad (23)$$

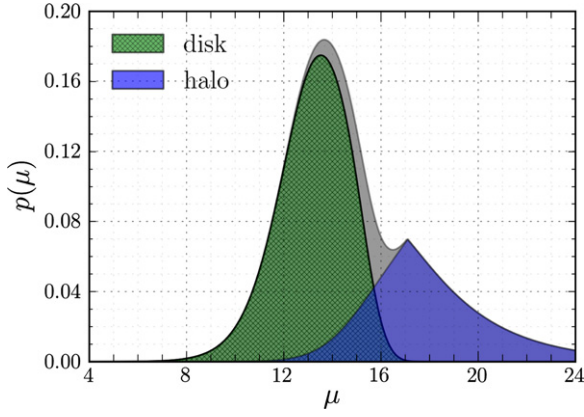


Figure 2. Distance prior for $(\ell, b) = (90^\circ, 10^\circ)$. The contributions of the disk and halo are shown individually in green and purple, respectively, while the total prior is given by the gray contour. The break in the contribution from the halo is due to the use of a broken power law for the number density of stars in this component.

(A color version of this figure is available in the online journal.)

Table 1
Stellar Number Density Parameters

Thin Disk		Thick Disk		Halo	
R_{thin}	2150 pc	R_{thick}	3261 pc	R_{br}^a	27.8 kpc
Z_{thin}	245 pc	Z_{thick}	743 pc	q_h^a	0.70
		f_{thick}	0.13	f_h	0.003
				η_{inner}^a	2.62
				η_{outer}^a	3.80

Note. ^a Values from Sesar et al. (2011). All other adopted values are from Jurić et al. (2008).

with

$$R_{\text{eff}} \equiv \sqrt{R^2 + (Z/q_h)^2 + R_\epsilon^2}. \quad (24)$$

Here, q_h controls the oblateness of the halo, and η controls the steepness of the power law. Following Sesar et al. (2011), the power law breaks at $R_{\text{eff}} = R_{\text{br}}$, becoming steeper. We therefore define η_{inner} and η_{outer} , corresponding to the halo steepness inside and outside of the break. We introduce the distance scale R_ϵ over which the inner region of the halo is smoothed, in order to remove the singularity at the Galactic Center. We choose R_ϵ to be 500 pc; at this scale, it has a negligible effect on the halo density in the regions that Jurić et al. (2008) studied but it prevents the halo from dominating over the disk at the Galactic Center and is the same scale adopted by Robin et al. (2003).

We employ the parameters given in Table 10 of Jurić et al. (2008) and Sesar et al. (2011). These are listed in Table 1. The adopted value of f_h , on the low end of the possible range inferred in Jurić et al. (2008), is chosen to better match observed PS1 color–magnitude diagrams at high Galactic latitudes. Throughout, we use $R_\odot = 8$ kpc and $Z_\odot = 25$ pc. The shape of the distance prior for a line of sight centered on $\ell = 90^\circ$ and $b = 10^\circ$ is shown in Figure 2.

4.2.2. Metallicity

We adopt the model of Galactic metallicity developed in Ivezić et al. (2008a) and Bond et al. (2010), assigning separate metallicity distributions to the disk and halo. The metallicity distribution of the disk varies with height above the Galactic plane and is thus dependent on line-of-sight distance. The

Table 2
Metallicity Parameters

	Disk		Halo
a_D	−0.89	a_H	−1.46
σ_D	0.20	σ_H	0.30
c	0.63		
Δa	0.14		
Δ_μ	0.55		
H_μ	0.5 kpc		

metallicity prior takes the form

$$p([\text{Fe}/\text{H}] | \mu) = p([\text{Fe}/\text{H}] | \mu, \text{disk})p(\text{disk} | \mu) + p([\text{Fe}/\text{H}] | \mu, \text{halo})p(\text{halo} | \mu). \quad (25)$$

The membership probabilities are simply

$$p(\text{disk} | \mu) = \frac{n_{\text{thin}}(\mu) + n_{\text{thick}}(\mu)}{n_{\text{thin}}(\mu) + n_{\text{thick}}(\mu) + n_{\text{halo}}(\mu)}, \quad (26)$$

$$p(\text{halo} | \mu) = 1 - p(\text{disk} | \mu), \quad (27)$$

which can be calculated on the basis of the preceding discussion (Section 4.2.1).

For the disk, stellar metallicity is distributed as a sum of Gaussians. The mean of each Gaussian varies with height above the Galactic plane, so that the prior is best written in terms of cylindrical coordinates:

$$p([\text{Fe}/\text{H}] | Z, \text{disk}) = c \mathcal{N}([\text{Fe}/\text{H}] | a(Z), \sigma_D) + (1 - c) \mathcal{N}([\text{Fe}/\text{H}] | a(Z) + \Delta a, \sigma_D), \quad (28)$$

where

$$a(Z) = a_D + \Delta_\mu e^{-|Z|/H_\mu}. \quad (29)$$

The parameters a_D , controlling the central disk metallicity, and Δ_μ and H_μ , describing the vertical metallicity gradient in the disk, are defined in Table 2. We assume the halo metallicity to be spatially invariant, and distributed as

$$p([\text{Fe}/\text{H}] | \text{halo}) = \mathcal{N}([\text{Fe}/\text{H}] | a_H, \sigma_H). \quad (30)$$

The parameters, a_H and σ_H , describing the halo metallicity, are given in Table 2. The metallicity is plotted in Figure 3 as a function of height above the Galactic midplane in the solar neighborhood.

4.2.3. Absolute Magnitude

We use the r -band absolute magnitude to parameterize the luminosity of each star. The joint prior on luminosity and distance is

$$p(\mu, M_r) \propto \frac{dN(\mu, M_r)}{d\mu dM_r}. \quad (31)$$

The luminosity function is assumed to be same in the halo and both disk components, and furthermore independent of position. The priors on distance and luminosity are then separable, so that

$$p(\mu, M_r) = p(\mu)p(M_r), \quad (32)$$

with $p(M_r) = \text{LF}(M_r) \propto dN/dM_r$.




Figure 3. Metallicity prior, $p([\text{Fe}/\text{H}] | Z)$, in the solar neighborhood ($R = 8$ kpc). High above the plane of the Galaxy, where the halo dominates, the metallicity distribution has a constant mean and variance. In the plane, where the disk dominates, the mean decreases with scale height. Adapted from Figure 9 of Ivezić et al. (2008a).

(A color version of this figure is available in the online journal.)

We adapt the PS1 luminosity functions provided by the PARSEC (Bressan et al. 2012), assuming a Chabrier (2001) log-normal initial mass function. We average over luminosity functions for populations with ages of $\tau = 7 \pm 2$ Gyr and metallicities of $[\text{Fe}/\text{H}] = -0.5 \pm 0.5$ dex. We denote the luminosity function for a population of age τ and metallicity $[\text{Fe}/\text{H}]$ as $\text{LF}(M_r | \tau, [\text{Fe}/\text{H}])$. The luminosity function we adopt is then

$$\text{LF}(M_r) \propto \int d\tau \int d[\text{Fe}/\text{H}] \text{LF}(M_r | \tau, [\text{Fe}/\text{H}]) \times \exp \left[-\frac{(\tau - \tau_0)^2}{2\sigma_\tau^2} - \frac{([\text{Fe}/\text{H}] - [\text{Fe}/\text{H}]_0)^2}{2\sigma_{[\text{Fe}/\text{H}]}^2} \right], \quad (33)$$

with $\tau_0 = 7$ Gyr, $\sigma_\tau = 2$ Gyr, $[\text{Fe}/\text{H}]_0 = -0.5$ dex, and $\sigma_{[\text{Fe}/\text{H}]} = 0.5$ dex. In principle, it is possible to make the luminosity function depend on metallicity, by not averaging over $[\text{Fe}/\text{H}]$ in Equation (33). For simplicity, we assume here that the luminosity function is universal.

4.2.4. Reddening

As indicated in Section 3, the manner in which we have factorized the line-of-sight reddening problem requires us to place a flat prior on the color excess, E , for each star. The priors on the reddening profile are imposed on the parameters which control the line-of-sight reddening, rather than on individual stellar reddenings. For example, if one divides each line of sight into N distance bins and assigns a different dust density ρ_i to each bin, then the reddening prior would take the form $p(\rho_1, \rho_2, \dots, \rho_N)$.

4.2.5. Survey Selection Function

The distance prior developed above only asks how many stars are in a thin shell at each distance. However, for a magnitude-limited survey, we would like instead to know the number of observable stars at a given distance. We should assign zero prior probability to the possibility of a star being observed which our instrument cannot detect. The fact that a star has been observed

by a given instrument therefore tells us something about its stellar type, distance, and extinction. Using the notation from Sale (2012), we define the vector S for each star, where S_i is true if a star has been observed in passband i , and false if the star is not detected in that passband. The PS1 data set used in this paper is not based on forced photometry, so there is a separate probability of a source being detected in each passband. If forced photometry were conducted, there would be one single probability $p(S)$, equal to the probability of detecting the source in at least one of the passbands. Including this information in the single-star posterior, Equation (13), we get

$$p(\mu, E, \Theta | m_{\text{obs}}, S) \propto p(m_{\text{obs}} | \mu, E, \Theta, S) \times p(\mu, E, \Theta | S). \quad (34)$$

But the prior is now just

$$p(\mu, E, \Theta | S) \propto p(S | \mu, E, \Theta) p(\mu, E, \Theta), \quad (35)$$

so in full,

$$p(\mu, E, \Theta | m_{\text{obs}}, S) \propto p(m_{\text{obs}} | \mu, E, \Theta, S) p(\mu, E, \Theta) \times p(S | \mu, E, \Theta). \quad (36)$$

The first term is simply the likelihood we found earlier, since the knowledge that the star has been detected has no effect on the apparent magnitudes the model predicts, assuming the stellar type, distance, and reddening are known. That is to say,

$$m_{\text{mod}} = M_{\text{mod}}(\Theta) + A(E) + \mu, \quad (37)$$

and

$$p(m_{\text{obs}} | \mu, E, \Theta, S) = \mathcal{N}(m_{\text{obs}} | m_{\text{mod}}, \sigma). \quad (38)$$

The only element of the calculation which changes when we take into account Malmquist bias is therefore the prior, which picks up an extra factor of

$$p(S | \mu, E, \Theta) = p(S | m_{\text{mod}}) = \prod_i p(S_i | m_{\text{mod}, i}). \quad (39)$$

This is the survey selection function.

If forced photometry were used instead, then we would have a single detection parameter S , denoting that the source was detected in at least one passband, and the survey selection function would be

$$p(S = \text{true} | \mu, E, \Theta) = p(S | m_{\text{mod}}) = 1 - \prod_i p(S_i = \text{false} | m_{\text{mod}, i}) \quad (40)$$

in place of the expression in Equation (39).

We therefore require an estimate of the completeness of the survey in each band, as a function of apparent magnitude. We determine completeness by comparison with point-source detections in the 275 deg² SDSS Stripe 82 survey (York et al. 2000; Annis et al. 2011). The co-added Stripe 82 images go more than a magnitude deeper than the individual PS1 3π images, allowing us to use Stripe 82 detections as a complete catalog of point sources past the detection limits of the PS1 3π survey. For each Stripe 82 source, we determine whether there is a PS1 detection within $1''$. The completeness fraction of the PS1 3π




Figure 4. Completeness of the PS1 3π survey, as a function of magnitudes past a locally estimated 5σ magnitude limit. The completeness is estimated by comparison with SDSS Stripe 82 (York et al. 2000; Annis et al. 2011). The shaded curve shows the median completeness, with 1σ range of completeness in each bin, based on estimates in $27'$ pixels. The solid black line shows our fit to the completeness curve. The dashed black line shows the effect of adding a small floor to our fit, which takes into account an assumed small rate of false coincidences between PS1 and Stripe 82 detections.

(A color version of this figure is available in the online journal.)

survey is the percentage of Stripe 82 detections with a PS1 match.

We determine the completeness fraction of the PS1 3π survey as a function of

$$\Delta m \equiv m - m_{\text{lim}}, \quad (41)$$

where m is the PS1 magnitude, and m_{lim} is an estimate of the local PS1 5σ magnitude limit, based on the point spread function of nearby PS1 detections and local sky and read noise. We divide the SDSS Stripe 82 footprint into HEALPix $n_{\text{side}} = 128$ pixels (with $\sim 27'$ scale). In each pixel, we select all Stripe 82 detections classified as stars and transform their $ugriz$ magnitudes to $grizy_{\text{P1}}$, using color transformations derived by D. P. Finkbeiner (in preparation), based on standard-star catalogs. In each pixel, we determine the PS1 limiting $grizy_{\text{P1}}$ magnitudes from the median limiting magnitudes estimated for individual PS1 detections. For each Stripe 82 detection in the pixel, we obtain Δm in each band by subtracting the local limiting magnitude from the transformed detection magnitude. In each passband, we bin Stripe 82 detections by Δm , obtaining an empirical estimate of the completeness in each bin from the number of PS1 matches.

We find that the completeness fraction is reasonably well fit by

$$p(S_i = \text{true} | m_{\text{mod}, i}) = \left[1 + \exp \left(\frac{m_{\text{mod}, i} - m_{\text{lim}, i} - \Delta m_1}{\Delta m_2} \right) \right]^{-1}, \quad (42)$$

where $m_{\text{lim}, i}$ is a limiting magnitude calculated for each point-source detection in the PS1 3π survey, equal to the magnitude of a source that would be detected at 5σ in one exposure, given the sky and read noise. $\Delta m_1 = 0.16$ mag and $\Delta m_2 = 0.2$ mag are fitting parameters. The positive value of Δm_1 indicates

that the PS1 pipeline goes somewhat deeper than our naive estimate m_{lim} . The same fitting parameter values reproduce the completeness curve in all five PS1 passbands reasonably well, reflecting the consistency of the PS1 optics and pipeline across the entire filter set. The empirically measured completeness fraction and our fit are plotted for each passband in Figure 4.

5. SAMPLING METHOD

5.1. Individual Stars

We use MCMC sampling to explore the parameter space for individual stars. The sampling must be performed with great care, owing to two features of the distributions $p(\mu, E)$. First, the distributions are invariably highly elongated, and second, they are often multimodal. The elongation stems from the close alignment between the reddening vector and the stellar locus in the PS1 bands, as shown in Figure 5. The multimodality has two causes. First, the reddening vector in general intersects the gri stellar locus in two locations, creating a degeneracy between blue and red main-sequence stars. Second, the PS1 bands do not distinguish dwarfs from giants, leading to the possibility that a star can be either a faraway red giant or a nearby red dwarf.

These degeneracies are easier to visualize if we consider only three passbands, as shown in Figure 5. In this reduced space, stellar photometry is fully described by a single overall observed magnitude and two colors. If we observe a star at a given location in color-color space, we can then move backward along the reddening vector until we intersect the stellar locus. One can then compare the observed apparent magnitude with the absolute magnitude of the stellar locus at the point of intersection. One thus obtains both a reddening and distance for the star. If there are multiple intersections, then the observed star could be of different intrinsic types and thus have different reddening and distance combinations. The relative probability of each mode is,




Figure 5. Sketch of how photometric parallax works, for illustrative purposes, adapted from Berry et al. (2011). A star is observed at location 1 in color-color space. Its dereddened colors may lie along any point on the gray line, parallel to the reddening vector. The intersections of this line with the model stellar locus, labeled 2 and 3, represent the most likely intrinsic stellar types. The posterior density for the star will thus have two modes—one at larger distance and lesser reddening (2) and one at smaller distance and greater reddening (3). For simplicity, we assume Solar metallicity in this example. This is how one would make a distance and reddening determination by eye. Our more rigorous Bayesian method takes into account photometric uncertainties, as well as priors on stellar type and Galactic structure.

(A color version of this figure is available in the online journal.)

in practice, given by the space of stellar types lying close to the gray line, as well as the priors applied to the problem.

We perform an MCMC sampling of these surfaces using a custom C++ implementation of the affine-invariant sampler introduced in Goodman & Weare (2009) and recently given in a python implementation by Foreman-Mackey et al. (2012). We employ both short-range “stretch” and long-range “replacement” moves (Goodman & Weare 2009). The long-range moves allow mixing between widely separated modes in parameter space but are more computationally expensive than the short-range “stretch” steps. The replacement moves are closely related to the Normal Kernel Coupler of Warnes (2001). We have found that, with the addition of long-range “replacement” steps, the affine-invariant sampler is capable of handling the multimodality of the problem and that it is well suited to the strong degeneracies in parameter space. For each star, we sample from each stellar posterior density in four independent runs and check convergence with the Gelman–Rubin diagnostic (Gelman & Rubin 1992). The Gelman–Rubin diagnostic essentially verifies that the variance between the means of separate chains is small compared with the variance within the chains. If independent MCMC runs produce significantly different estimates of the parameter means, one or more of the runs must not have converged. Each run employs 20 samplers, with a mix of 80% stretch steps and 20% replacement steps, and 2000 steps per sampler. The first 1000 steps from each sampler are discarded as burn-in. Thus, excluding the burn-in phase, a total of 80,000 samples are drawn for each star across the four chains, with the number of independent samples being lower. On four cores of a 2.67 GHz Intel Xeon X5650 with 12 MB of L3 cache, our run time per star is typically 0.15 s per run per core, or 0.8 CPU seconds for four independent runs.

5.2. Bayesian Evidence and Outlier Rejection

When an observed object does not match our stellar model, the inferences we draw on its distance, reddening, and stellar type are unreliable. One means of quantifying the reliability of our inferences for an individual star is to compute the evidence

$$Z \equiv \int d\mu dE d\Theta p(\mathbf{m}, \mathbf{S} | \mu, E, \Theta) p(\mu, E, \Theta), \quad (43)$$

which is the probability density of drawing the observed magnitudes \mathbf{m} from the stellar model and observing the star in the PS1 survey. A low evidence indicates that the observed point source is a member of a stellar population not included in our model (e.g., a young blue giant or an unresolved binary system with colors that do not match any stellar template), is not a star (e.g., a white dwarf or galaxy), that the errors in the photometry have been underestimated, or that the reddening vector is inaccurate. Here, our approach is similar to Berry et al. (2011), which identifies objects which do not fit the stellar model by a threshold χ^2 statistic. There is no direct analogue for the χ^2 statistic in a Bayesian framework but evidence may serve a similar purpose in model comparison.

Note that, as we do not include priors on the extinction to individual stars but rather on the line-of-sight reddening profile, we do not strictly calculate the evidence. Instead, we calculate the evidence of a model with a prior on E with wide support (i.e., a prior which allows E to take on a wide range of values), such that the prior is nearly constant across all relevant reddenings:

$$p(E) \approx \begin{cases} a & 0 \leq E \lesssim E_0 \\ af(E) & E \gtrsim E_0 \end{cases}, \quad (44)$$

where $f(E)$ is some integrable function whose precise behavior is unimportant, a is a normalizing constant, and E_0 is some large extinction. We thus effectively calculate the evidence Z for a model of this form, up to a constant factor a , which is the same for every star. This allows for outlier rejection, based on comparisons between the evidence for different stars.

We employ a modified harmonic mean estimate, which is obtained directly from the Markov chain produced in sampling the posterior density and thus requires little additional computation (Gelfand & Dey 1994; Robert & Wraith 2009). This method is presented in more detail in Appendix A.

5.3. Line-of-Sight Fitting

We choose the HEALPix pixelization scheme (Gorski et al. 2005) as our method of dividing the sky into individual lines of sight. Once we have determined $p(\mu, E | \mathbf{m})$ for each star in a given HEALPix pixel and rejected stars which fall below the evidence cut, we can apply Equation (12) to determine the posterior probability of the parameters α describing the reddening profile. We parameterize the reddening profile as a piecewise-linear function in distance modulus, with $\alpha_i = \Delta E^{(i)}$ describing the rise in r -band reddening in distance segment i . We split up each line of sight into 30 distance segments of equal width in μ , with the closest distance being at $\mu = 4$, corresponding to 63 pc, and the furthest distance being at $\mu = 19$, corresponding to 63 kpc. It must be cautioned that, in general, our method does not tightly constrain reddening at this latter distance, where PS1 observes very few stars. In addition to requiring that reddening increase monotonically with distance, we apply a wide log-normal prior on the differential reddening in each distance bin, $\{\Delta E^{(i)}\}$.

We use the affine-invariant sampler to draw a representative sample of possible reddening profiles. We sample from the posterior density given by Equation (6). We can thus produce a 3D reddening map which includes the uncertainty in reddening as a function of distance.

6. TESTS WITH MOCK PHOTOMETRY

The first and most straightforward test of our method is to generate mock photometry for stars of varying stellar type, distance, and extinction, and to see how well we can recover those parameters. This is less of a test of the particular stellar model used than a demonstration that photometry alone is capable of sufficiently constraining stellar parameters. We find that our method is capable of accurately recovering both single-star parameters and the line-of-sight reddening profile.

6.1. Generating Mock Catalogs

In order to generate a mock photometric catalog for a particular region on the sky, we begin by drawing intrinsic stellar types (metallicities and absolute r_{p1} magnitudes) and distances from our priors. We assign a reddening to each star, either according to an assumed distance–reddening relationship, or from a reddening distribution we define, depending on the purpose of the mock catalog. For each star in the catalog, we generate model magnitudes, as described in Section 4.1.

We determine which passbands each star is detected in, according to our probabilistic PS1 completeness model, Equation (42). In the remainder of the paper, we reject simulated stars which do not have five-band detections.

We then apply magnitude-dependent Gaussian photometric errors to each simulated star to obtain observed magnitudes. The error we apply to each star in each passband is a function of the model apparent magnitude:

$$\sigma^2(m) = \sigma_{\text{floor}}^2 + \sigma_0^2 \exp \left[\frac{2(m - m_{\text{lim}})}{\Delta m_3} \right]. \quad (45)$$

As before, m_{lim} is the 5σ limiting magnitude in the given passband. We set the error floor to $\sigma_{\text{floor}} = 0.02$. For PS1 passbands, we find that $\sigma_0 = 0.16$ and $\Delta m_3 = 0.8$ give a reasonable fit to the photometric uncertainties.

Our final catalog thus contains noisy observed magnitudes of each star, along with the photometric uncertainty in each passband. As a final step, we plug the observed magnitudes back into Equation (45) to obtain a new estimate of the photometric uncertainties for each star. The final catalog that we pass to our pipeline thus reports inexact photometric uncertainties, much as a realistic catalog would. The usefulness of these mock catalogs is that they allow us to generate a large amount of photometry for stars whose “true” distances and reddenings are known.

6.2. Single-star Tests

We illustrate the typical appearance of single-star posterior distributions in distance and reddening in Figure 6. For each of the four simulated stars in Figure 6, the “true” distance and reddening are indicated by a dot, while the background heat map shows the probability density inferred by our pipeline.

In order to determine how far off our estimates are on average, we define the centered probability density

$$\tilde{p}(\Delta\mu, \Delta E) \equiv \frac{1}{N} \sum_{i=1}^N p(\mu_i^* + \Delta\mu, E_i^* + \Delta E), \quad (46)$$




Figure 6. Distance and reddening estimates for four simulated stars. The joint posterior in distance and reddening is shown as a heat map. As this is mock photometry, we know the “true” distances and reddenings for the stars, which are shown as green dots. The true stellar parameters lie in regions of inferred high probability, as expected. The shape of the probability density functions traces that of the stellar locus. The probability density at closer distances corresponds to the main sequence, with increasing reddening compensating for the bluer intrinsic colors as one travels up the stellar locus. The peak in reddening corresponds to the main-sequence turnoff. Distances beyond the turnoff correspond to the giant branch.

(A color version of this figure is available in the online journal.)

where μ_i^* and E_i^* are the true distance modulus and reddening, respectively, for star i . For a simulated line-of-sight, this function gives the average probability density of our inference being offset from the true stellar parameters by $(\Delta\mu, \Delta E)$. We plot \tilde{p} for a typical line of sight in Figure 7. The typical spread of $\Delta\mu$ and ΔE varies across different lines of sight, but the centered probability density generally peaks at the origin, as should be expected. In the bottom two panels of Figure 7, we show the effect of applying flat priors to the stellar parameters, in place of the priors developed in Section 4. The effect is to widen and bias the inferred probability density functions. Due to the near-alignment of the reddening vector with the PS1 stellar locus for much of the main sequence, this bias remains even for stars with low observational uncertainties. The stellar priors are thus important in correctly inferring stellar parameters from PS1 photometry. In the right two panels of Figure 7, we only use inferences for stars with low signal-to-noise detections. The low signal-to-noise population is generated using an inflated error model that applies three times the normal observational uncertainty to the mock photometry. As expected, the inferred parameters for such stars are less constrained but they are nonetheless unbiased.

In Table 3, we present typical uncertainties in the inferred distance and reddening of individual stars. To do this, we generate mock catalogs along two different lines of sight. For our high-Galactic-latitude target, we choose the North Galactic Pole, where the stellar population is dominated by the halo. Here, we apply reddenings of $E(B - V) \lesssim 0.1$ to the simulated stars. For the low-Galactic-latitude target, we choose $\ell = 45^\circ$, $b = 0^\circ$, and draw reddening uniformly from the range $0 \leq E(B - V) \leq 2$. We run the two mock catalogs through our pipeline and compare the inferred distances and reddenings, drawn from the posterior probability density $p(\mu, A)$, to the true values. For this test, we allow our inferred reddenings to go slightly negative ($E(B - V) > -0.25$) to avoid introducing




Figure 7. Centered and stacked probability densities on a linear scale for 5000 stars along a simulated line of sight pointed at $\ell = 90^\circ$, $b = 20^\circ$. In the bottom panels, we fit the stars using flat priors, so that only the likelihood function comes into play. In the right panels, we show inferences for low signal-to-noise detections, generated using three times the normal observational uncertainties. Removing the priors biases the inferred distances and, to a lesser extent, reddenings. Inferred distances and reddenings for stars with high signal-to-noise detections have smaller uncertainties. In each panel, each stellar probability density function has first been centered on the true distance and reddening, before the probability densities for the stars have been summed, as described in the text. The “X” shape of the stacked probability densities in the top-left panel reflects the existence of separate giant and dwarf modes. The feature stretching from the bottom left to the top right corresponds to the main sequence, while the perpendicular feature corresponds to the giant mode. The histograms bordering each panel show the distribution of $\Delta\mu$ and $\Delta E(B - V)$, with the 15.87% to 84.13% region shaded.

(A color version of this figure is available in the online journal.)

Table 3
Uncertainty in Inferred Distances and Reddenings

	Low Latitude		High Latitude	
	$\frac{\Delta d}{d}$ ^a	$\Delta E(B - V)$	$\frac{\Delta d}{d}$	$\Delta E(B - V)$
$-1 < M_r \leq 4$	$-20^{+41\%}_{-30\%}$	$-0.03^{+0.07}_{-0.12}$	$-37^{+48\%}_{-32\%}$	$-0.03^{+0.09}_{-0.11}$
$4 < M_r \leq 6$	$6^{+55\%}_{-33\%}$	$0^{+0.12}_{-0.20}$	$12^{+79\%}_{-37\%}$	$0^{+0.10}_{-0.12}$
$6 < M_r \leq 8$	$23^{+97\%}_{-34\%}$	$0.17^{+0.28}_{-0.36}$	$8^{+92\%}_{-21\%}$	$0.04^{+0.22}_{-0.20}$
$8 < M_r \leq 10$	$33^{+261\%}_{-39\%}$	$0.29^{+0.62}_{-0.42}$	$4^{+31\%}_{-14\%}$	$0.02^{+0.33}_{-0.09}$
$10 < M_r \leq 12$	$5^{+14\%}_{-15\%}$	$0.02^{+0.13}_{-0.13}$	$1^{+15\%}_{-12\%}$	$0^{+0.10}_{-0.09}$
Dwarfs ^b	$9^{+63\%}_{-33\%}$	$0.01^{+0.16}_{-0.20}$	$4^{+50\%}_{-18\%}$	$0.01^{+0.17}_{-0.12}$
All Stars	$-3^{+55\%}_{-35\%}$	$-0.01^{+0.12}_{-0.16}$	$2^{+45\%}_{-23\%}$	$0.01^{+0.16}_{-0.12}$

Notes.

^a $(\Delta d/d)$ is given in percent. See Equation (48).

^b Dwarfs are defined here as all stars in the range $4 < M_r \leq 12$.

a bias into the inferred values. We give the median, and 15.87th and 84.13th percentiles of $\Delta d/d$ and $\Delta E(B - V)$, equivalent to the one-standard-deviation range for a Gaussian distribution.

Uncertainties in distance modulus can be transformed to uncertainties in distance by making use of the relation

$$d = (10 \text{ pc}) 10^{\mu/5}. \quad (47)$$

Let $\mu_{\text{inferred}} = \mu_{\text{true}} + \Delta\mu$. Then,

$$\frac{\Delta d}{d} \equiv \frac{d_{\text{inferred}} - d_{\text{true}}}{d_{\text{inferred}}} = 10^{\Delta\mu/5} - 1. \quad (48)$$

Similarly, we define $\Delta E(B - V)$ as $E(B - V)_{\text{inferred}} - E(B - V)_{\text{true}}$.

Our distance and reddening estimates are unbiased. However, if one selects a subsample of stars of a certain known type, a bias in distance is introduced. Thus, distance estimates for mock dwarf stars are biased low, as the model assigns some probability to the possibility of them being giants. Inversely, distances to giants are biased high. A star drawn at random, however, has an unbiased distance estimate. Distance and reddening constraints depend upon the quality of the photometry and direction on the sky, and therefore vary significantly on a star-per-star basis. It is, in general, more informative to look at the detailed shape of the posterior distribution for a given star in distance and reddening space (see Figure 6).

Next, we test that the true stellar parameters are drawn from the probability density functions we calculate. For each simulated star, we derive the posterior density, $p(\mu, E)$, based on the simulated photometry. Since we know the true distance modulus μ^* and reddening A^* of the star, a natural question is whether μ^* and E^* are drawn from $p(\mu, E)$. This cannot be answered for a single star but we can test this hypothesis for a large number of stars. We assign a percentile to a given star as follows:

$$P(p < p^*) \equiv \int_{p(\mu, E) < p^*} d\mu dE p(\mu, E), \quad (49)$$

where $p^* \equiv p(\mu^*, E^*)$. This represents the probability that the true stellar parameters would be found at a point in (μ, E) -space of lower posterior density. If (μ^*, E^*) lies at the point of maximum posterior density, then $P(p < p^*) = 1$. Conversely, if (μ^*, E^*) lies in a region of vanishing probability density, then $P(p < p^*) \approx 0$. This percentile is therefore similar to a cumulative distribution function and is uniformly




Figure 8. Distribution of percentiles for simulated photometry of 1000 stars, as defined in Equation (49). The percentiles are expected to be drawn from the standard uniform distribution, resulting in each bin being of equal height. In 50% of trials, we would expect all the bins to fall within the dark green band, and in 95% of trials, all the bins should lie within the light green band. The percentiles are consistent with being drawn from a uniform distribution, indicating that we are sampling from the model correctly.

(A color version of this figure is available in the online journal.)

distributed between 0 and 1. A straightforward test of whether $P(p < p^*) \sim U(0, 1)$ is to generate mock photometry for a large number of stars, to calculate the percentile for each simulated star, and to then bin the results. Each bin is expected

to contain the same number of stars, with the precise number of stars in the bins determined by a multinomial distribution. We can therefore derive approximate confidence intervals for the number of stars that should fall into any given bin. Figure 8 shows this test for a set of 1000 simulated stars along a line of sight with Galactic coordinates $\ell = 90^\circ$, $b = 10^\circ$. The results are consistent with our expectations, indicating that our method recovers correct posterior densities for the simulated photometry.

6.3. Mock Line of Sight

Finally, we demonstrate that we are able to recover line-of-sight reddening profiles for simulated photometry. We first invent an arbitrary relationship, $E(B - V)(\mu)$, between distance and reddening. We add in low-level scatter to the distance–reddening relationship, as the reddening relation across one HEALPix pixel may vary. We then generate mock photometry for 150 stars along the line of sight. Following the procedure outlined in Section 3, we use the simulated photometry to determine a posterior density in distance and reddening space for each star, and then combine the information from all of the stars to find the range of allowable reddening profiles. Our final product is thus a set of reddening profiles, drawn from the probability density over reddening profiles (Equation (12)). We parameterize the reddening as a piecewise-linear function in distance and apply a weak log-normal prior to the differential reddening in each distance segment, as described in Section 5.3. The results for one simulated line of sight, shown in Figure 9, demonstrate that we are able to correctly infer the reddening




Figure 9. Recovery of the line-of-sight reddening profile from simulated photometry for 150 stars. The large panel shows the inferred posterior densities of the stars, stacked on top of one another. The contrast is stretched at each distance for purposes of visualization. This gives a picture of the information which is fed into the second stage of our analysis, in which we recover the reddening as a function of distance from the individual stellar probability densities (see Equation (12)). The stacked image, however, plays no direct role in this inference. The curves show possible reddening profiles, conditioned on the mock photometry. The green curve traces the most probable reddening profile. The remaining curves are colored according to the logarithm of their probability density, with blue denoting high probability and red denoting low probability. We recover a reddening profile similar to the original, which had a single cloud of depth $E(B - V) = 0.5$ at distance modulus $\mu = 8.5$ (shown as a dashed black line on the plot). The slight, gradual increase in inferred reddening beyond the cloud is due to the constraint that differential reddening in each bin be non-negative and to the log-normal prior on differential reddening. A priori, having no reddening away from the cloud is unlikely, and this is reflected in our inference. The upper four panels show individual stellar posterior density functions over the same domain, with the same reddening profiles overplotted. Each star is consistent with the range of possible reddening profiles.

(A color version of this figure is available in the online journal.)




Figure 10. Comparison of PS1 stellar colors in the vicinity of the North Galactic Pole with our model colors. Each object is colored according to the evidence Z we compute. Objects represented by red dots have a low probability of being drawn from our stellar model and are rejected for the line-of-sight reddening determination. The solid black line traces our model stellar colors. Our main-sequence model colors do not depend on metallicity, while the model colors for the giant branch have a slight metallicity dependence.

(A color version of this figure is available in the online journal.)

profile for mock data. The method produces the best results at distances where there are many stars. Nearby and at large distances, where there are comparatively few stars to constrain the reddening profile, the uncertainties in reddening can become very large. However, at intermediate distances ($\mu \sim 10$ to ~ 15 for typical lines of sight, corresponding to 1–10 kpc), the fit produces uncertainties on the order of $\Delta E(B-V) \sim 0.05$ mag, consistent with the intrinsic scatter in reddening which we introduce into the simulated photometry.

7. COMPARISON WITH DATA

7.1. Colors

We compare our model colors to PS1 stellar photometry from low-extinction regions at high Galactic latitudes. It is important

to choose low-extinction regions, so that assumptions about the reddening law and what percentage of the total dust column is in front of each star play only a minor role. This allows us to obtain a comparison between the intrinsic colors in our model and of real stars. We de-redden the stellar colors assuming that they are behind the full dust column predicted by Schlegel et al. (1998, SFD).

The results for the North Galactic Pole are shown in Figure 10. We compute an evidence for each star, as described in Section 5.2. As expected, objects which lie far from the stellar locus in color-color space tend to have lower evidence. This helps us to reject objects that are either not stars, are not included in our stellar model (e.g., young and blue horizontal branch (BHB) stars), or that have particularly bad photometry. In the window shown here, 15% of the detected objects would fail an evidence cut of $\ln Z > \ln Z_{\max} - 20$. These objects tend to have problematic photometry for which the PS1 pipeline may have produced inaccurate results, though some are variables, quasars, and unrecognized galaxies, which our technique is not designed to handle. Our line-of-sight reddening inferences are not strongly dependent on the choice of the evidence threshold.

7.2. Distances

Correct distance determination requires not only correct model colors, but correct absolute magnitudes. We therefore compare our stellar models to globular and open clusters. In Figure 11, we compare our model magnitudes to photometry from four globular and open clusters.

We find that the model absolute magnitudes match the main sequence. The model magnitudes are unreliable past the main sequence turnoff, particularly for younger clusters. Our model magnitudes trace the giant branch of intermediate-age clusters—typical of the age of most stars in the Galaxy—somewhat better. Nevertheless, we expect that most of the information in our dust maps will come from the main sequence, where distance and reddening estimates are better constrained. Massive young blue stars and blue horizontal branch stars, which are not included in the stellar templates, are found to have low evidence, allowing them to be identified and excluded from the line-of-sight dust inference. Inclusion of age-dependent stellar models, and therefore more reliable distance and reddening determinations for the most massive stars, could potentially increase the distance to which our dust maps are reliable and is an important direction for future work.




Figure 11. PS1 color-magnitude diagrams of three globular and one open cluster. For each cluster, the model isochrone with the catalog metallicity of the cluster is overplotted. The stellar photometry has been de-reddened and shifted by the catalog distance modulus to produce absolute magnitudes. The reddening vector is plotted in the top left corner of each panel in red for reference. Each star is colored by its evidence, with red stars unlikely to be drawn from our stellar model. In particular, stars which are blueward of the main-sequence turnoff, which are bluer than any star in our template library, have low evidence.

(A color version of this figure is available in the online journal.)




Figure 12. Histogram of the difference between the mode of the Bayesian reddening inference and the SEGUE-derived reddening, as a function of SFD reddening (see Section 7.3). The blue envelopes mark the 15.87 and 84.13 percentiles of the residuals, equivalent to one standard deviation for a normal distribution, while the central blue curve marks the median of the residuals. The left panel compares the mode of the Bayesian posterior density functions with the mean of the SEGUE-derived reddening. The right panel compares random samples drawn from the Bayesian posteriors with random samples drawn from the SEGUE-derived posteriors, which are Gaussian.

(A color version of this figure is available in the online journal.)

7.3. Reddenings

In order to test the accuracy of our reddening inferences for individual stars, we compare our photometric reddening to independently measured reddening for a sample of stars. The SEGUE survey (Yanny et al. 2009), part of SDSS-II, provides a convenient set of stars for which one can independently determine reddening. The SEGUE survey obtained moderate-resolution spectroscopy for 240,000 stars with SDSS photometry. Whereas most of the SDSS footprint is at low reddening, some of the SEGUE targets are at moderate reddening, up to ~ 1 mag in $E(B - V)$. The SEGUE Stellar Parameter Pipeline (SSPP) fits an atmospheric model to each star to derive the temperature, metallicity, and gravity of the star, as well as other parameters (Lee et al. 2008a, 2008b; Allende Prieto et al. 2008). These stellar parameters were used by Schlafly & Finkbeiner (2011) to predict the intrinsic colors of stars and to study the effect of reddening by attributing the differences between the observed and intrinsic colors to dust. We use the reddening estimates of Schlafly & Finkbeiner (2011) in the four independent SDSS colors, in concert with their recommended $R_V = 3.1$ reddening vector (Fitzpatrick 1999), to estimate $E(B - V)$ to each star. The details of deriving the reddening based on SEGUE-determined intrinsic colors and SDSS photometric colors are described in Appendix B.

We use the same sample of SEGUE targets as Schlafly & Finkbeiner (2011). This sample excludes objects targeted as white dwarfs and also removes M dwarfs, for which the stellar parameters are less reliable. We also require that each SEGUE target have a PS1 counterpart. Distances and reddenings to each star are then inferred as described in Section 4.1, and compared to the SEGUE-determined reddenings. For this comparison, we allow our photometric reddening estimates to be negative, for consistency with the SEGUE-derived reddenings. We save 100 reddening samples from the Markov chain for each star, as well as the maximum-posterior density reddening. In the left panel of Figure 12, we compare the maximum-posterior density Bayesian reddening estimate with the SEGUE-derived reddening for 200,000 stars. We bin the stars by the reddening expected from the dust maps of Schlegel et al. (1998, SFD) and plot a histogram of the difference in the two reddening measures in each bin. We place the SFD reddening on the x -axis because it is a good proxy for reddening and is independent of

both of the two reddening estimates we wish to compare, while placing either the SEGUE-derived reddening or the Bayesian reddening along the x axis can introduce spurious trends in the resulting comparison. We find that the mode of our reddening estimates is unbiased over a range of $E(B - V) = 0$ to 1, above which there are too few stars in the SEGUE sample to extend the comparison. The scatter in the difference between the two reddening estimates is approximately 0.12 mag in $E(B - V)$, with the overall estimate being unbiased to within 0.03 mag.

Our Bayesian reddening and distance estimates assume that stars are drawn at random from the observable stars on each line of sight. The SEGUE survey, however, does not target stars at random but instead targets only subsets of stars of particular interest. Moreover, the subset of SEGUE-observed stars for which we have reliable reddening estimates does not extend to the M dwarfs, meaning that our sample of SEGUE-derived reddening estimates use only intrinsically blue stars. We consider intrinsically redder stars—and therefore less reddened stars—in our analysis than are actually present in our sample of SEGUE-observed stars. In order to simulate the effect of excluding M dwarfs, we modify the luminosity function prior to assign zero probability for $M_r > 6$. When we account for this effect, the distribution of the difference between our and the SEGUE reddening estimates is unbiased. To illustrate this, instead of presenting single reddening difference for each star, we show 100 samples of the distribution of reddening differences between our Bayesian reddening estimate and the SEGUE-derived reddening estimate. The resulting residuals are shown as a function of SFD reddening in the right panel of Figure 12. The residuals are unbiased at the 0.01 mag level, with a scatter of $E(B - V) = 0.13$ mag. This result is indicative of the accuracy we achieve for high-signal-to-noise detections, as most SEGUE targets are well above the detection limit in PS1.

8. CONCLUSION

We have presented a general method for deriving a 3D map of Galactic reddening from stellar photometry. Our technique is based on grouping stars into pixels, determining the joint posterior of distance and reddening for each star, and then determining the most probable reddening–distance relation in each pixel. We have shown that this method correctly

recovers the distance–reddening relationship for simulated lines of sight. We have additionally shown by comparison with SEGUE-derived reddenings, that for high-SNR detections, our Bayesian reddening estimates are unbiased at the 0.01 mag level, with a scatter of ~ 0.13 mag in $E(B - V)$. Based on comparisons with mock catalogs, in highly reddened regions of the Galaxy, our distance inferences have typical uncertainties of +47% on the high end, and –21% on the low end. In high-Galactic latitude regions with low reddening, our distances have typical uncertainties of +52% on the high end, and –38% on the low end. These uncertainties may be reduced by feeding back information on reddening as a function of distance, derived from all the stars along the line of sight. A subsequent paper will present the results of applying the techniques developed here to construct a 3D reddening map covering the $\delta > -30^\circ$ sky.

In addition to determining the dust density in the nearby Galaxy, our method can be used to determine the distribution of stars in the Galactic plane. Earlier optical studies of the distribution of stars in the Galaxy traditionally consider only high-latitude stars, where the correction for dust extinction is straightforward (e.g., Jurić et al. 2008). Infrared surveys of the plane are less sensitive to dust extinction but their wavelength coverage also makes them less sensitive to intrinsic stellar type, rendering photometric distances uncertain. Our technique provides distances to stars throughout the Galactic plane, enabling future studies of the distribution of stars in the disk.

The technique described in this paper is not limited to PS1 photometry. Inclusion of information from the 2MASS J , H , and K_s bands, WISE bands, as well as SDSS u -band photometry will improve our distance and reddening estimates. In addition, kinematic information, such as proper motion, may be incorporated into our framework in order to allow a more precise determination of stellar distances.

Upcoming surveys will also dramatically enhance our ability to measure the distances and reddenings to stars in the Galaxy. The LSST (Ivezic et al. 2008) will provide deeper photometry spanning a similar set of filters as those used in SDSS and PS1, providing photometry for the sky south of $\delta < +34.^\circ 5$. In the nearer future, the Dark Energy Survey will survey a complementary 5000 deg^2 of sky to PS1, in a similar filter set (The Dark Energy Survey Collaboration 2005). The *Gaia* mission (Lindegren et al. 1994), meanwhile, will provide multiband photometry and low-resolution spectroscopy alongside parallax distance measurements and proper motions for one billion stars. *Gaia*’s parallax distances, in particular, will break many of the degeneracies in our model for $r_{p1} \lesssim 20$ stars, while its proper motions will aid in inferring the population each star belongs to. These new data sets will increase the power of our method to determine Galactic reddening and structure.

The Pan-STARRS1 (PS1) surveys have been made possible through contributions of the Institute for Astronomy, the University of Hawaii, the PS1 Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, Queens University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronau-

tics and Space Administration under grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation under grant No. AST-1238877, the University of Maryland, and Eotvos Lorand University (ELTE). Gregory M. Green and Douglas P. Finkbeiner are partially supported by NSF grant AST-1312891. The computations in this paper were run on the Odyssey cluster supported by the FAS Science Division Research Computing Group at Harvard University.

APPENDIX A

HARMONIC MEAN ESTIMATE OF THE BAYESIAN EVIDENCE

The harmonic mean approximation, developed in Gelfand & Dey (1994), allows one to compute the Bayesian evidence using samples returned from a Markov chain Monte Carlo (MCMC) simulation. For a model with parameters θ and data D , Bayes’ rule tells us that

$$p(\theta|D) = \frac{p(D|\theta)p(\theta)}{p(D)}. \quad (\text{A1})$$

We wish to compute the evidence $p(D)$, often denoted by Z . Multiplying each side of the above by an arbitrary function $\phi(\theta)$, rearranging terms and taking the integral over all θ ,

$$\frac{1}{p(D)} \int d\theta \phi(\theta) = \int d\theta p(\theta|D) \frac{\phi(\theta)}{p(D|\theta)p(\theta)}. \quad (\text{A2})$$

The right-hand side is simply the expectation value of

$$\frac{\phi(\theta)}{p(D|\theta)p(\theta)} \quad (\text{A3})$$

for samples drawn from the posterior density $p(\theta|D)$. This is convenient, since MCMC methods draw a set of samples from the distribution $p(\theta|D)$. If the integral of $\phi(\theta)$ is normalized to unity, then

$$\frac{1}{p(D)} \equiv \frac{1}{Z} \approx \left\langle \frac{\phi(\theta)}{p(D|\theta)p(\theta)} \right\rangle_{\text{chain}}. \quad (\text{A4})$$

This estimate has finite variance as long as ϕ has steeper wings than $p(\theta|D)$ (Robert & Wraith 2009). We choose ϕ to be constant within an ellipse centered on a point of high density within the chain, and zero outside (Robert & Wraith 2009). The ellipse is aligned with the principle axes of the covariance matrix, in order to ensure that it contains only well-sampled regions of parameter space.

We find a point of high density by first centering an ellipse on a random point in the chain. Because the points in the chain are sampled proportionately to the posterior probability density, this point is already likely to lie in a well-sampled region of parameter space. We then find the mean of the points in the chain falling in the ellipse, and move the center of the ellipse to that position in parameter space. One can iterate this procedure several times to settle into a densely sampled region of parameter space. The size of the ellipse we use to define $\phi(\theta)$ is chosen such that a preset fraction of samples in the chain are enclosed. For our calculations, we iterate five times to find a dense region of parameter space and scale the ellipse such that it contains 5% of samples.

APPENDIX B

SEGUE-DERIVED REDDENINGS

Here, we review a method for calculating stellar reddenings on the basis of SDSS photometry and SEGUE predicted intrinsic stellar colors. As explained in Schlafly & Finkbeiner (2011), since the extinction in an individual band X is given by

$$A_X = R_X E(B - V), \quad (\text{B1})$$

colors transform as

$$E(X - Y) = A_Y - A_X = (R_Y - R_X)E(B - V). \quad (\text{B2})$$

If we have only one color, $X - Y$, we can therefore estimate reddening as

$$E(B - V) = \frac{E(X - Y)}{R_Y - R_X}. \quad (\text{B3})$$

Our goal is to extend this formula to allow the use of multiple colors, possibly with strong covariance. Since Schlafly & Finkbeiner (2011) only predict the colors of stars, and not their overall magnitudes, we work in color space. We denote the intrinsic stellar colors as \mathbf{c}_i , and the reddened colors as \mathbf{c}_r .

In a multidimensional color space, Equation (B2) becomes

$$\mathbf{c}_r - \mathbf{c}_i = \mathbf{R} E(B - V), \quad (\text{B4})$$

where \mathbf{R} has one component per color $X - Y$, given by $R_{XY} \equiv R_Y - R_X$. The estimated intrinsic colors and observed reddened colors are Gaussian random variables, with covariances Σ_i and Σ_r , respectively. We denote the estimated intrinsic colors as \mathbf{c}'_i , and the observed reddened colors as \mathbf{c}'_r . The likelihood of these two quantities taking on a particular set of values is given by

$$p(\mathbf{c}'_i, \mathbf{c}'_r | E(B - V), \mathbf{c}_r) = \mathcal{N}(\mathbf{c}'_r | \mathbf{c}_r, \Sigma_r) \mathcal{N}(\mathbf{c}'_i | \mathbf{c}_i, \Sigma_i) \quad (\text{B5})$$

$$= \mathcal{N}(\mathbf{c}'_r | \mathbf{c}_r, \Sigma_r) \mathcal{N}(\mathbf{c}'_i | \mathbf{c}_r - \mathbf{R} E(B - V), \Sigma_i). \quad (\text{B6})$$

In the second step, we have replaced \mathbf{c}_i using Equation (B4). Using the symmetry of the Gaussian distribution,

$$p(\mathbf{c}'_i, \mathbf{c}'_r | E(B - V), \mathbf{c}_r) = \mathcal{N}(\mathbf{c}_r | \mathbf{c}'_r, \Sigma_r) \mathcal{N}(\mathbf{c}_r - \mathbf{R} E(B - V) | \mathbf{c}'_i, \Sigma_i) \quad (\text{B7})$$

$$= \mathcal{N}(\mathbf{c}_r | \mathbf{c}'_r, \Sigma_r) \mathcal{N}(\mathbf{R} E(B - V) - \mathbf{c}_r | -\mathbf{c}'_i, \Sigma_i). \quad (\text{B8})$$

If we assume a flat prior on $E(B - V)$ and \mathbf{c}_r , then the above is proportional to the posterior probability density $p(E(B - V), \mathbf{c}_r | \mathbf{c}'_i, \mathbf{c}'_r)$. We could, in practice, frame the question in terms of the intrinsic colors \mathbf{c}_i , and put priors on them based on our Galactic and stellar model, but we wish to avoid tying our SEGUE-derived reddenings in any way to our Bayesian photometric reddening estimates. Now, integrating over \mathbf{c}_r , we obtain a convolution of two Gaussians, which is itself a Gaussian distribution:

$$p(E(B - V) | \mathbf{c}'_i, \mathbf{c}'_r) \propto \int d\mathbf{c}_r \mathcal{N}(\mathbf{c}_r | \mathbf{c}'_r, \Sigma_r) \times \mathcal{N}(\mathbf{R} E(B - V) - \mathbf{c}_r | -\mathbf{c}'_i, \Sigma_i) \quad (\text{B9})$$

$$= \mathcal{N}(\mathbf{R} E(B - V) | \mathbf{c}'_r - \mathbf{c}'_i, \Sigma_r + \Sigma_i). \quad (\text{B10})$$

The probability density function of $E(B - V)$ is thus a ray taken through a multivariate Gaussian. It can be shown that the above is also Gaussian, with mean and standard deviation given by

$$\langle E(B - V) \rangle = \frac{(\mathbf{c}'_r - \mathbf{c}'_i)^T (\Sigma_r + \Sigma_i)^{-1} \mathbf{R}}{\mathbf{R}^T (\Sigma_r + \Sigma_i)^{-1} \mathbf{R}}, \quad (\text{B11})$$

$$\sigma_{E(B-V)}^2 = [\mathbf{R}^T (\Sigma_r + \Sigma_i)^{-1} \mathbf{R}]^{-1}. \quad (\text{B12})$$

We plug intrinsic stellar colors appropriate for the SSPP stellar parameters into \mathbf{c}'_r and use the observed SDSS colors for \mathbf{c}'_i .

One final note is that Schlafly & Finkbeiner (2011) used the SSPP stellar types to derive estimates of the mean and covariance of the magnitudes, rather than the colors. We will now show how to obtain the covariance of the colors from the covariance matrix of the magnitudes. We denote the covariance matrix of the magnitudes as Σ_{ij} , where i and j label passbands. We label the color $m_i - m_j$ as c_{ij} . We then call the covariance matrix of the colors $\Sigma'_{ij,kl}$ and c_{kl} . By expanding out

$$\Sigma'_{ij,kl} = \langle c_{ij} c_{kl} \rangle - \langle c_{ij} \rangle \langle c_{kl} \rangle \quad (\text{B13})$$

in terms of magnitudes, one obtains

$$\Sigma'_{ij,kl} = \Sigma_{ik} - \Sigma_{il} - \Sigma_{jk} + \Sigma_{jl}. \quad (\text{B14})$$

The common choice of colors is to set the i th color to $m_i - m_{i+1}$. Plugging $j = i + 1$ and $l = k + 1$ into the above, we find that the covariance of the i th color with the k th color is given by

$$\Sigma'_{i,k} = \Sigma_{i,k} - \Sigma_{i,k+1} - \Sigma_{i+1,k} + \Sigma_{i+1,k+1}. \quad (\text{B15})$$

REFERENCES

- Allende Prieto, C., Sivarani, T., Beers, T. C., et al. 2008, *AJ*, **136**, 2070
 Annis, J., Soares-Santos, M., Strauss, M. A., et al. 2011, arXiv:1111.6619v2
 Bailer-Jones, C. A. L. 2011, *MNRAS*, **411**, 435
 Belokurov, V., Zucker, D. B., Evans, N. W., et al. 2006, *ApJL*, **642**, L137
 Berry, M., Ivezić, Ž., Sesar, B., et al. 2011, *ApJ*, **757**, 166
 Bond, N. A., Ivezić, Ž., Sesar, B., et al. 2010, *ApJ*, **716**, 1
 Bovy, J., Rix, H.-W., & Hogg, D. W. 2012a, *ApJ*, **751**, 131
 Bovy, J., Rix, H.-W., Hogg, D. W., et al. 2012b, *ApJ*, **755**, 115
 Bovy, J., Rix, H.-W., Liu, C., et al. 2012c, *ApJ*, **753**, 148
 Bressan, A., Marigo, P., Girardi, L., et al. 2012, *MNRAS*, **427**, 127
 Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, **345**, 245
 Chabrier, G. 2001, *ApJ*, **554**, 1274
 Drew, J. E., Greimel, R., Irwin, M. J., et al. 2005, *MNRAS*, **362**, 753
 Fitzpatrick, E. L. 1999, *PASP*, **111**, 63
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2012, *PASP*, **125**, 306
 Gelfand, A. E., & Dey, D. K. 1994, *J. R. Stat. Soc.*, **56**, 501
 Gelman, A., & Rubin, D. B. 1992, *StaSc*, **7**, 457
 Goodman, J., & Weare, J. 2009, *Commun. Appl. Math. Comput. Sci.*, **5**, 65
 Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
 Hodapp, K. W., Kaiser, N., Aussel, H., et al. 2004, *AN*, **325**, 636
 Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv:1008.4686
 Ivezić, Ž., Sesar, B., Jurić, M., et al. 2008a, *ApJ*, **684**, 287
 Ivezić, Z., Tyson, J. A., Acosta, E., et al. 2008, arXiv:0805.2366
 Jurić, M., Ivezić, Ž., Brooks, A., et al. 2008, *ApJ*, **673**, 864
 Kaiser, N., Burgett, W., Chambers, K., et al. 2010, *Proc. SPIE*, **7733**, 77330E
 Kruschke, J. 2011, *Doing Bayesian Data Analysis: A Tutorial Introduction with R* (Amsterdam: Academic Press)
 Lallement, R., Vergely, J.-L., Valette, B., et al. 2013, *A&A*, **561**, A91
 Lallement, R., Welsh, B. Y., Vergely, J. L., Crifo, F., & Sfeir, D. 2003, *A&A*, **411**, 447
 Law, D. R., & Majewski, S. R. 2010, *ApJ*, **714**, 229
 Lee, Y. S., Beers, T. C., Sivarani, T., et al. 2008a, *AJ*, **136**, 2022
 Lee, Y. S., Beers, T. C., Sivarani, T., et al. 2008b, *AJ*, **136**, 2050
 Lindegren, L., Perryman, M. A., Bastian, U., et al. 1994, *Proc. SPIE*, **2200**, 599

- Magnier, E. 2006, in *The Advanced Maui Optical and Space Surveillance Technologies Conference*, ed. S. Ryan (Maui: Maui Economic Development Board)
- 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*, ed. W. J. Jin, I. Platais, & M. A. C. Perryman (Cambridge: Cambridge Univ. Press), 553
- Marshall, D., Robin, A., Reyl  , C., Schultheis, M., & Picaud, S. 2006, *A&A*, 453, 635
- Onaka, P., Tonry, J. L., Isani, S., et al. 2008, *Proc. SPIE*, 7014, 70140D
- Perryman, M. A. C., Lindegren, L., Kovalevsky, J., et al. 1997, *A&A*, 323, 49
- Robert, C. P., & Wraith, D. 2009, arXiv:0907.5123
- Robin, A. C., Reyl  , C., Derri  re, S., & Picaud, S. 2003, *A&A*, 409, 523
- Sale, S. E. 2012, *MNRAS*, 427, 2119
- Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, 737, 103
- Schlafly, E. F., Finkbeiner, D. P., Juri  , M., et al. 2012, *ApJ*, 756, 158
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Sesar, B., Juri  , M., & Ivezi  ,   . 2011, *ApJ*, 731, 4
- Simon, J. D., & Geha, M. 2007, *ApJ*, 670, 313
- Stubbs, C. W., Doherty, P., Cramer, C., et al. 2010, *ApJS*, 191, 376
- The Dark Energy Survey Collaboration. 2005, arXiv:astro-ph/0510346
- Tonry, J., & Onaka, P. 2009, in *Advanced Maui Optical and Space Surveillance Technologies Conference*, ed. S. Ryan (Maui: Maui Economic Development Board)
- Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, *ApJ*, 750, 99
- Vergely, J.-L., Freire Ferrero, R., Siebert, A., & Valette, B. 2001, *A&A*, 366, 1016
- Warnes, G. 2001, *The Normal Kernel Coupler: An adaptive Markov Chain Monte Carlo Method for Efficiently Sampling from Multi-modal Distributions*, Tech. Rep. 395 (Seattle, WA: Department of Statistics, Univ. Washington)
- 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