Mem. S.A.It. Vol. 86, 568

Memorie della

Galactic 3D extinction maps

S.E. Sale[1,2]

[1] Astrophysics Group, School of Physics, University of Exeter, Stocker Road, Exeter EX4 4QL, UK, e-mail: ssale@astro.ex.ac.uk
[2] Rudolf Peierls Centre for Theoretical Physics, Keble Road, Oxford OX1 3NP, UK

Abstract. I review the history of 3D extinction mapping, discussing the success and examining the shortcomings of existing maps. In addition, identifying where different authors have adopted common methodological approaches and where they differ, so as to better understand how different approaches relate to each other.
I will also look to the future, considering the impact Gaia will have, whilst also pondering the technical and methodological challenges that must be overcome if we are to fully take advantage of the wealth of data such surveys promise to provide.

Key words. dust – extinction – ISM: structure

1. Introduction

Despite being a relatively minor contributor to the mass of the ISM, interstellar dust plays an important role in the physics and chemistry of the ISM. Moreover, the extinction of starlight by dust has a significant confounding impact on observations that complicates the study of individual stars and stellar populations. Without an accurate 3d map of extinction, observations from surveys such as Gaia will be unable to determine either the absolute luminosity of stars, or where the effects of extinction will diminish the number of visible stars – making it impossible to understand the structure, dynamics and history of the Galaxy. At the same time, we should also view the extinction of starlight by dust as a significant opportunity: since extinction affects the light from stars distributed throughout the Galaxy it offers a direct route to studying the elusive 3d structure of the ISM. Consequently, there is a long history of studies attempting to determine the 3d distributions of interstellar dust and its resulting extinction, stretching back to Trumpler (1930) and van de Kamp (1930).

In general terms, one maps extinction by taking some ‘input catalogue’ of data and then employing that to build up a picture of how extinction builds up along different lines of sight. The data used can take any form, as long as it is well characterised. However, the sheer volume of photometric data and so its ability to more finely sample the ISM means it is normally preferred. Though the impending availability of Gaia astrometry promises a step change in extinction mapping.

There are a number of specific aims one could adopt for 3d extinction mapping. The most obvious is to infer the extinction to stars that feature in the ‘input catalogue’. As demonstrated by Sale & Magorrian (2014), when mapping extinction in 3d one can estimate the extinction (and other properties such as distance) to a particular star, more precisely than if it were studied in isolation. That this should be the case is not surprising. In the framework of extinction mapping one estimates the extinction to a star conditioned on observations of many stars. Such an estimate will inherently be more precise than one conditioned on observations of just the one star. In practice, the information gleaned from other stars helps narrow the likely combinations of distance and extinction available to the star in question.

One might also be interested in the estimating the extinction to some arbitrary point or points in space. Given all UV, optical and near infra-red observations of the Galaxy are affected by extinction, this is often a crucial task in understanding observations. In particular, our ability to fit models of the Galaxy's stellar density to observations relies on accurate 3d extinction maps (Bovy et al. 2015; Farnhill et al. 2015). Without an understanding of how extinction accumulates in 3d we do not know over what ranges of distance we should expect to see any type of star and so cannot make direct comparison to models. There is already a pressing need for 3d extinction maps (e.g. Schönrich et al. 2015). However, this need will grow greatly as data from Gaia become available.

In addition, perhaps one is interested in the extinction to a particular object that, for whatever reason, is not in the input catalogue. Alternatively, given an independent estimate of the extinction to some object and a prediction of the distance-extinction relationship along the line of sight to it, one can obtain an estimate of the object's distance, sometimes called an 'extinction distance'. Such an approach is often undertaken when more direct methods to estimate distance are unavailable, but the extinction can be found, as is the case with planetary nebulae (e.g. Giammanco et al. 2011).

Determining the large scale distribution of extinction is also an important task. Prosaically, doing so will potentially unveil significant dusty features in the Galaxy, such as spiral arms. It will also shed light on the large scale relationship between dust, the wider ISM and stars.

Finally, mapping extinction in 3d can potentially enable us to examine the distribution of dust on small scales, i.e. $\lesssim$ 100 pc. On these scales the diffuse ISM is largely dominated by turbulence and so any ability to probe this regime would shed light on the applicability of models of turbulence, such as those of Kolmogorov (1941) or Goldreich & Sridhar (1995).

So, we can summarise the possible tasks encompassed within 3d extinction mapping as a desire to make some inference about:

1. The extinction to objects in the input catalogue.
2. The extinction to arbitrary points in space.
3. The large scale distribution of extinction.
4. The small scale distribution of extinction.

Unsurprisingly, given the growing need, in recent years there has been a significant effort devoted to producing 3d extinction maps. The many approaches differ in their details, for example relying on different data. However, there are also significant areas of commonality. In the sections that follow, I will classify maps and techniques that exist to illustrate where different authors have followed the same path and where they diverge.

It is worth noting that there exist many 2d maps of extinction that won't be discussed here. Most prominent amongst these is the Schlegel et al. (1998) map of total Galactic extinction. Other popular applications of 2d extinction mapping include studying the Galactic bulge/bar (Gonzalez et al. 2012) and molecular clouds (Lombardi & Alves 2001). However, whilst such these maps are extremely useful, e.g. for studying extra Galactic objects (in the case of Schlegel et al. 1998), they have limited utility when studying the wider Galaxy and the stars distributed within it.

2. Voxelised methods

The first class of extinction maps that we will consider are those that divide the volume of the Galaxy into many smaller bins. In general, these first divide the 2d sky into a number of 'pixels'. The gridding may be aligned with Galactic coordinates (most maps), use HEALPix (e.g. Green et al. 2015) or may simply divide the sky up arbitrarily (Neckel &

Table 1. An incomplete list of 3d extinction maps created with individual observation matching type methods that cover substantial and contiguous portions of the sky.

Source	Cat. Size	Sky coverage	App. depth/kpc
Fernie (1962)	1 500	Galactic Plane	1
Neckel (1966)	4 700	all sky	2
Scheffler (1966)	4 700	$\lvert z\rvert < 75$ pc	2
Fitzgerald (1968)	7 835	all sky	3
Lucke (1978)	4 000	all sky	2
Krautter (1980)	unclear	Galactic Plane	3
Neckel & Klare (1980)	11 000	$\lvert b\rvert \lesssim 5°$	3
Perry & Johnston (1982)	3 450	all sky	0.3
Ducati (1986)	3 713	all sky	0.5
Pandey & Mahra (1987)	462 open clusters	$\lvert z\rvert < 200$ pc	2
Arenou et al. (1992)	58 000	all sky	1
Jones et al. (2011)	56 000	SDSS footprint	1
Berry et al. (2012)	7×10^7	SDSS footprint	2.5
Gontcharov (2012)	7×10^7	all sky	1
Sale et al. (2014)	4×10^7	$30° < l < 215°$, $\lvert b\rvert < 5°$	4 – 10
Chen et al. (2014)	3×10^7	$140° < l < 240°$, $-60° < b < 40°$	4
Green et al. (2015)	8×10^8	$\delta > -30°$	2 – 10

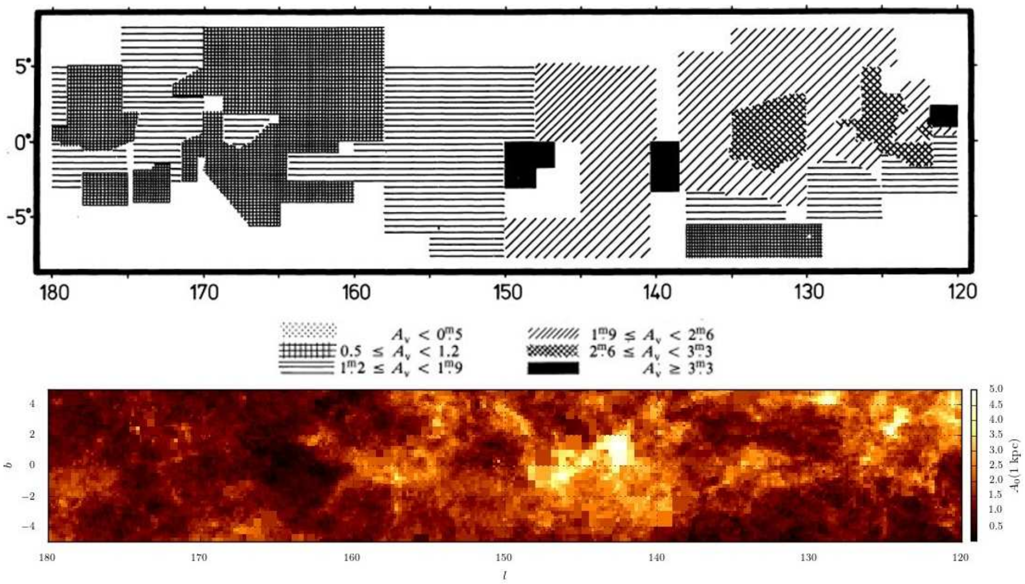

Fig. 1. A comparison between the extinction maps of Neckel & Klare (1980) (top) and Sale et al. (2014) (bottom). Both maps show the total extinction to a distance of 2 kpc. Top panel reproduced with permission from Astronomy & Astrophysics, ©ESO.

Klare 1980). Generally, the maps then create 'voxels' (or 3d pixels) by dividing each on sky pixel into a number of distance bins. Alternatively, each line of sight may instead be divided into a number of extinction bins, whose distance is then inferred (e.g. Marshall et al. 2006). Or, a functional form is instead fit to the distance extinction relationship for each

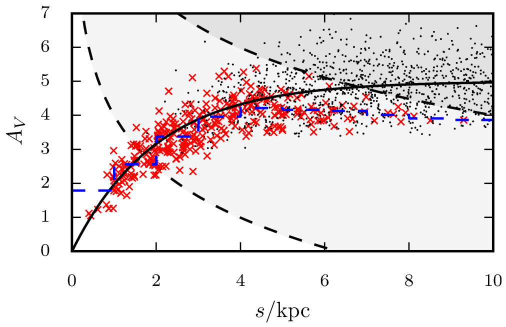

Fig. 2. A pseudo-photometric catalogue represented in distanceextinction space. Red crosses represent stars that appear in the catalogue, black points those that are too faint and so are not included in the catalogue. The two black dashed lines delimit the regions where the catalogue is 100 per cent complete (bellow the lower line) and 100 per cent incomplete (above the higher line). It can be clearly seen that the effect of the faint magnitude limit is to preferentially exclude more extinguished stars. The mean distanceextinction relationship found by taking the mean of the extinctions of observed stars is distance bins (blue dashed line) can be compared to the true mean distanceextinction relationship (black solid line). Figure reproduced from Sale (2015) with permission from MNRAS (Oxford University Press).

pixel (e.g. Arenou et al. 1992). Typically, most maps will proceed by considering one on sky pixel at a time and then build up a distance extinction relationship based on the binning they have applied along the line of sight.

There is one clear question that is invoked by all of these methods: how should the sky be voxelised? This is a question that has no clear answer. In some respect larger voxels are better since they contain more stars and so should reduce the impact of observational uncertainties. However, the ISM contains small scale structure, and having larger voxels makes it more difficult to resolve this detail.

Within this voxelised framework there are two common approaches to producing maps. Either, a map is made such that it agrees with the observations of each individual star, or the produced map aims to reproduce the characteristics of the entire population of observed stars.

2.1. Individual observation matching

The general concept behind mapping when matching to individual observations is simple. If we know the distances and extinctions to a catalogue of stars, we can then build up a picture of how extinction accumulates with distance. However, it is in the detail that the task becomes complicated. Specifically, dealing with the substantial observational uncertainties and biases is where most of the difficulty lies.

Perhaps as a result of its conceptual simplicity, this form of 3d extinction mapping has been relatively popular. A chronological list of maps that have adopted this approach is given in Table 1. It is apparent that more recent maps have benefited from the vast growth of available data; Green et al. (2015) employ a catalogue that is 5 orders of magnitude larger than that used by Fitzgerald (1968). There have also been significant advantages in methodology. Whereas Neckel & Klare (1980) fit their distance extinction relationships 'by-eye', more recent approaches have converged on the use of hierarchical Bayesian models, as discussed in Sale (2012) and applied in Sale et al. (2014) and Green et al. (2015).

Fig. 1 deomnstrates the impact of these advances by comparing slices through the Neckel & Klare (1980) and Sale et al. (2014) maps, that show the total extinction out to a distance of 2 kpc are. Both maps exhibit some common features, for example the relative lack of extinction around $l \sim 165°$. However, the result of over 30 years of progress is immediately apparent. The map of Sale et al. (2014) clearly reveals a turbulent ISM, with structures such as filaments readily apparent. Such detail is not available in the map of Neckel & Klare (1980).

Generally, recent approaches in this class employ optical photometry, which, especially in the era of modern ccd surveys, is plentiful. In addition, as the effects of extinction are relatively strong at optical wavelengths, reasonably precise estimates of the extinction to each star can be found.

The effects of catalogue completeness/ the catalogue selection function on 3d extinction mapping are pathological: for example, faint magnitude limited samples will preferentially exclude more extinguished and so fainter stars, thus biasing maps towards lower extinctions. As visualised in Fig. 2, this effect can be significant. Despite the fact that this effect has long been known (e.g. Neckel 1966), many maps fail to include a treatment for it, whilst others proceed by discarding data. A full treatment for the effects of selection functions, that employs all available data and produces unbiased maps is given in Sale (2015) and implemented in Sale et al. (2014).

The ISM is known to exhibit structure all the way down to significantly sub-AU scales (Spangler & Gwinn 1990). This reality is fundamentally at odds with a mapping approach that forms voxels that may be many hundreds of pc in size. As a result, it is unsurprising to learn that there exists substantial variation in extinction to points within a single voxel. This sub-resolution variation impairs the ability of voxelised methods to precisely predict extinction to particular points in space. Indeed, Sale et al. (2014) found that this differential extinction was the dominant source of uncertainty in the estimation of extinction to any given point within their map.

In the map of Sale et al. (2014) (Fig. 3) there are a number of azimuthal discontinuities, or fingers of God. Similar features are also visible in essentially all other maps within this class, including in Fig. 16 of Green et al. (2015). These fingers of God are clearly unphysical and consequently Fig. 3 does not resemble our expectations of how such a slice through a 3d map should appear. These features can only exist because these methods treat different sightlines independently, whilst they are exaggerated by the fact that measured extinction is strongly correlated between voxels along a line of sight.

We can consider the usefulness of this approach to extinction mapping by considering the aims enumerated in section 1. While this approach does make useful inferences about both the large scale distribution of extinction (task 3) and the extinction to object in the input catalogue (task 1). It struggles to make useful predictions of extinction to arbitrary points in space (task 2), owing to its problems with sub-resolution structure. It also offers no direct means to study the small scale structure of the ISM (task 4).

2.2. Population synthesis

A second class of voxelised maps seeks to produce model observations of a synthetic population of stars that in some statistical sense agree with real observations. The most prominent and indeed prototypical, example of a population synthesis map is that of Marshall et al. (2006), who map the inner Galactic plane ($|l| < 100°, |b| < 10°$). They employ 2MASS photometry, selecting a catalogue of late type giants for each on sky pixel. Employing intrinsically bright giant stars ensures that the map covers a significant range ($\sim$ 10 kpc), whilst they are less sensitive to the poorly constrained finer structure of the Galactic plane than OB stars. Similar methods have subsequently been employed by other authors, though with improved data (Chen et al. 2013; Schultheis et al. 2014) and using the photometry of all stars rather than a subset.

Conceptually the difference between this approach and that of the previous section is that one is no longer interested in the distances and extinctions of each individual star and only seeks a map that can reproduce the characteristics of the wider population. However, doing so requires a direct reliance on a Galaxy model. Therefore, the validity of these maps rests on the accuracy of the Galaxy model employed. If, for example, the assumed stellar density distribution were significantly wrong, then the estimated distribution of dust along each line of sight would also differ significantly from reality. Contrast this with what happens in the individual observation matching methods described above. Those methods typically (implicitly or explicitly) take some prior on the distance distribution of stars, but if the reality is different then the data naturally over-ride the prior and adjust the distribution of stars along the line of sight to more closely match reality. By employing a sample of more evolved giant stars, Marshall et al. (2006) reduce their vulnerability to problems of this type, since their distribution is smoother and better understood

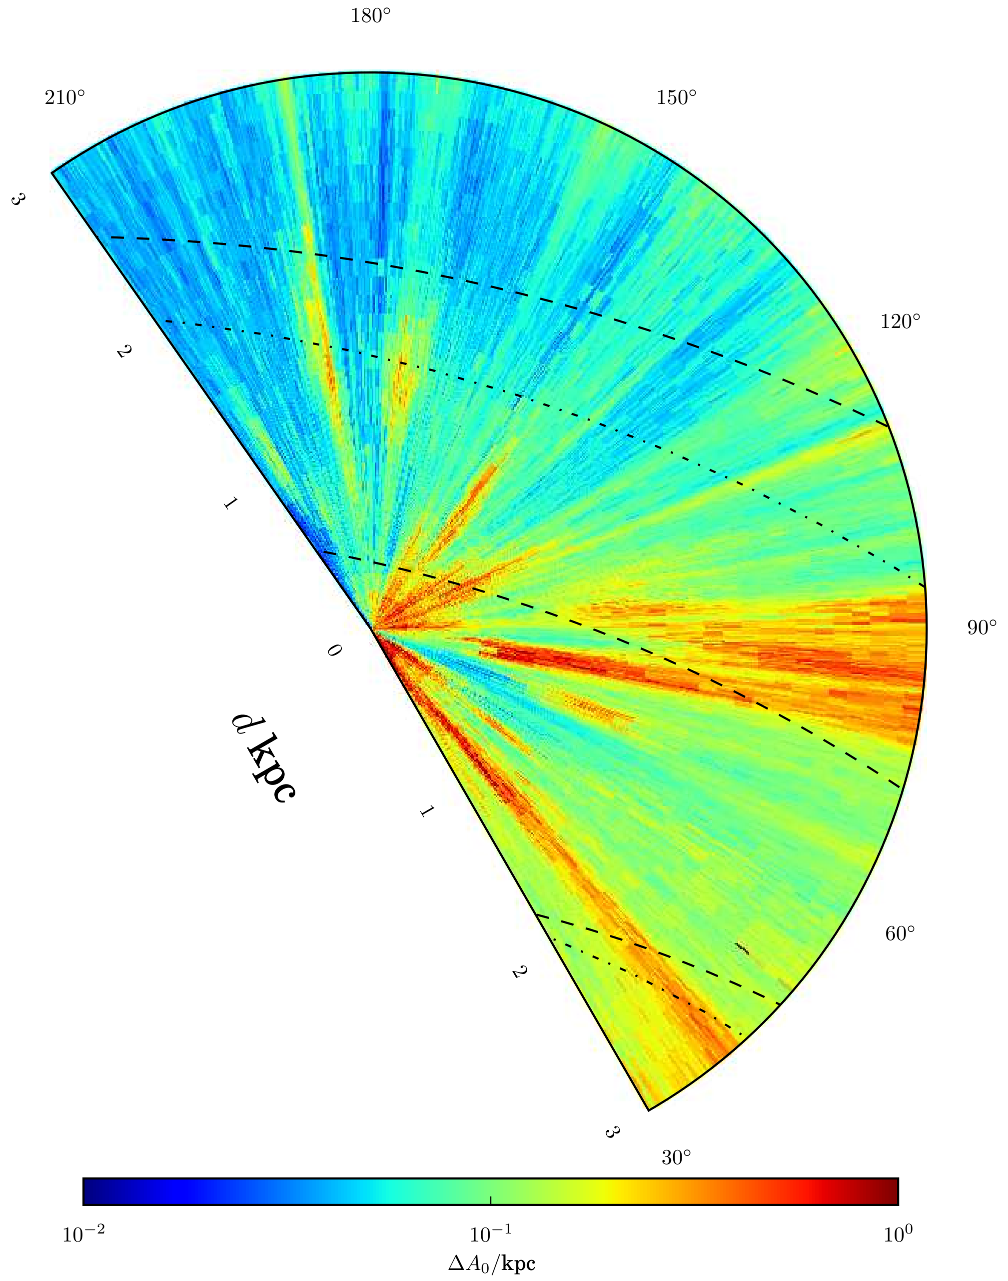

Fig. 3. The Sale et al. (2014) map of extinction at b = 0. The sun lies at the plots origin with the Galactic centre off the bottom of the plot. The dashed lines denote the position of the Sagittarius, local and Perseus spiral arms given by Reid et al. (2014), whilst the dot-dashed lines correspond to the Sagittarius and Perseus arms of Vallée (2013). Figure reproduced from Sale et al. (2014) with permission from MNRAS (Oxford University Press).

Table 2. A list of 3d extinction maps created with population synthesis methods that cover substantial and contiguous portions of the sky.

Source	Cat. Size	Sky coverage	App. depth / kpc
Marshall et al. (2006)	unclear	$\|l\| < 100°, \|b\| < 10°$	10
Chen et al. (2013)	unclear	$\|l\| < 10°$, $-10° < b < 5°$	10
Schultheis et al. (2014)	unclear	$\|l\| < 10°$, $-10° < b < 5°$	10

than the high mass young stars that would otherwise dominate their maps at any significant distance.

All the methods listed in table 2 rely on the use of near-IR photometry. At these wavelengths the effects of extinction are weaker and so the catalogues of the 2MASS and VVV surveys are deep enough to probe the most extinguished regions of the bulge. Conversely, estimating the extinction to individual stars is then more difficult, though this is overcome by considering observations of all the stars within each on-sky pixel.

There are some distinct advantages to producing 3d extinction maps via population synthesis. The effect of the survey selection function is far more easily dealt with: it can be applied to the synthetic observations of the populations within the Galaxy model in a way that mimics its effect on real observations. Additionally the computational cost can be significantly reduced, since there is no longer a contribution to the likelihood from each individual star. However, population synthesis maps also exhibit 'fingers of God; compare Fig. 9 of Marshall et al. (2006) or Fig. 16 of Chen et al. (2013) to Fig. 3 above.

Returning to the list of key tasks in section 1, population synthesis mapping does not directly address task 1, since it marginalises over the properties of individual stars. As with the individual observation matching approach, applications to date also offer no means to study the small scale structure of the ISM (task 4), whilst the ability to infer the extinction to arbitrary points in space is impaired by the existence of significant sub-voxel structure. It does however, enable the large scale structure of extinction to be studied (task 3)

3. Gridless methods

Clearly, the alternative to voxelsing space when producing a 3d extinction map is to simply not do so. However, one then has the problem of not only of determining how likely some combination of extinctions to a range of stars is, but also estimating the extinction to the infinite number of points in space where there is no star and so we lack a direct measurement of the extinction.

Gaussian processes provide the solution to both these problems. In a Gaussian process (GP) the value at some finite number of points follows a joint Gaussian distribution. Thus, not only is the probability of extinction to a collection of points for which we have observations defined. But, we can also regress within these points to make predictions at any unobserved location(s). GPs are each described by a mean function and a covariance function that links the distributions at any two points by providing the values in the covariance matrix of the relevant multivariate Gaussian distribution. One way to view GPs is as the limit of an infinite number of infinitely small voxels, whose values are related by the covariance function.

GPs, can also provide a convincing approximation to the physics of the ISM. Column density PDFs are thought to be largely lognormal in the diffuse ISM (Ostriker et al. 2001; Vázquez-Semadeni & García 2001) and so maps of the logarithm of extinction match this description. Moreover, as the covariance function is the fourier pair of the power spectrum, we can include a physical model for turbulence, e.g. that of Kolmogorov (1941), in the definition of the GP.

A key benefit of GP extinction mapping is that it does not result in fingers of God, as can be seen in Fig. 4. This is as the covariance function acts to ensure that relation-

Table 3. A list of 3d extinction maps created with gridless (i.e. Gaussian process) methods that cover substantial and contiguous portions of the sky.

Source	Cat.	Sky coverage	App. depth / kpc
Chen et al. (1998)	434 clusters	all sky	2
Vergely et al. (2010)	6 993	all sky	0.3
Lallement et al. (2014)	22 883	all sky	1

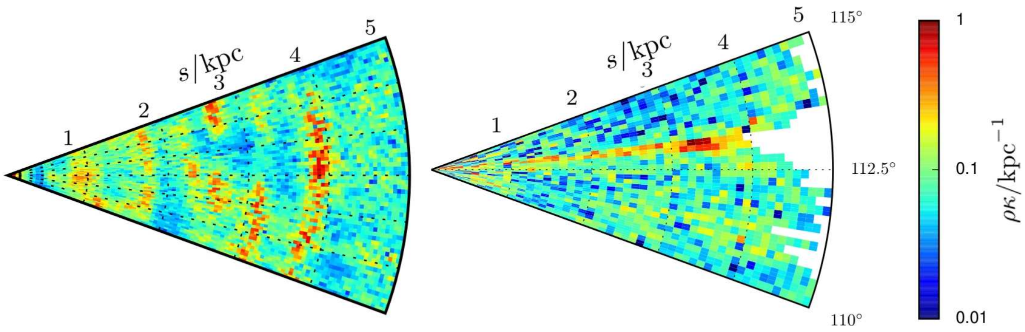

Fig. 4. Maps of dust density in the direction of IRAS 23151+5912. The right hand map is taken from Sale et al. (2014), whilst that on the left was made using the method of Sale & Magorrian (2014). Notice, in particular, the absence of fingers of God in the left hand map.

ship between extinctions along nearby lines of sight is related in a physically believable way. In addition we no longer have to worry about setting the resolution of the map. In a sense it is now set naturally by the availability of the data, since in regions with lots of data the map will have higher precision and be able to pick out smaller scale structure than in more sparsely observed regions. As a result, we can now estimate the extinction to single points in space with more precision than previously possible (Fig, 5), since the closest stars provide stronger constraints than those further away, even within regions that could be contained within a single on-sky pixel in voxelised maps.

If we return to the list of aims from section 1, it is apparent that GP based extinction mapping, in principal, can satisfy all of the aims. In particular, we now have some grasp on the small scale physics. This is as the probability of a map depends on the covariance function, which itself follows directly from the model of the turbulent ISM we employ.

However, implementing GPs on large scales is widely known to be difficult, due to the fact that the CPU cost of doing so normally scales with the cube of the number of observations. This has limited progress, with the maps of Chen et al. (1998), Vergely et al. (2010) and Lallement et al. (2014) the only published examples to date. These are limited to a relatively small volume around the sun and assume some large scale distribution of extinction. Sale & Magorrian (2014) have extended upon their approach , enabling the larger scale structure to be inferred and including improved treatments for distance uncertainties and incompleteness. Work to scale this up to large datasets is ongoing (Sale & Magorrian in prep.).

4. Model based maps

A final alternative is to employ some model of Galactic extinction to produce a 3d map. Some of the earliest 3d maps took the form of simple models for the distribution of dust, whose parameters were constrained by observations (e.g. Parenago 1945). More recently, Drimmel et al. (2003), scale a model of Galactic extinction to match the Schlegel et al. (1998) map of total Galactic extinction. While, Mendez & van Altena (1998), Chen et al. (1999) and

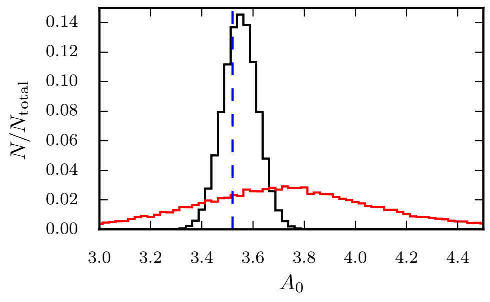

Fig. 5. Estimated posterior probability distributions of extinction to a single point in space from simulations in Sale & Magorrian (2014). The 'true' extinction to this point is $A_0 = 3.52$, as indicated by the vertical dashed blue line. The black histogram shows the posterior distribution found by the method of Sale & Magorrian (2014), whilst the red histogram shows a comparable estimate obtained using the method of Sale (2012). Note that we should not expect that the mean of the posteriors should exactly agree with the true value. The detail of interest is the relative widths of the two posterior distributions (i.e. the much reduced uncertainty in the result from Sale & Magorrian 2014). Figure reproduced from Sale & Magorrian (2014) with permission from MNRAS (Oxford University Press).

Amôres & Lépine (2005) also produce a 3d maps. Despite their strong reliance on models, these maps often prove useful. Particularly in the high latitude sky where the dust can be reasonably assumed to be heavily concentrated locally to us. However, closer to the Galactic midplane, where the distribution of dust along the line of sight is more complicated, these maps diverge from more empirical ones (Sale et al. 2009; Bovy et al. 2015).

5. Looking forward

The upcoming releases of Gaia data promise to revolutionise 3d extinction mapping. In particular, since the relatively crude estimates we currently have for the distances to stars are a significant source of uncertainty in extinction maps, Gaia parallaxes will be of significant use. Moreover, the existence of Gaia data itself creates an impending need for high quality 3d extinction maps, since these will be required to properly exploit Gaia (Bovy et al. 2015).

However, there are a number of potential hurdles that must be overcome if we are to produce 3d extinction maps of the required quality. From a technical standpoint, producing voxelised maps with Gaia sized catalogues is currently feasible. However, producing a Gaussian process based map with a catalogue of this size is far more challenging. Fortunately, this is a problem faced by a number of fields and there already exist useful approaches (e.g. Quiñonero-Candela & Rasmussen 2005). Efforts are already underway to apply methods taken from outside astrophysics to the peculiarities of the problem in hand (Sale & Magorrian in prep.).

Almost all recent 3d extinction maps have assumed a fixed wavelength dependence of dust opacity, as typically parametrised by R_V, generally adopting either the Cardelli et al. (1989), Fitzpatrick (1999) or Fitzpatrick (2004) $R_V = 3.1$ reddening law. However, it has long been known that the wavelength dependence of dust opacity does vary between sightlines (Cardelli et al. 1989). But attempts to measure either the 2d (Schlafly et al. 2010) or 3d (Zasowski et al. 2009) variation of R_V have been limited. Moreover, it appears that the canonical extinction laws of Cardelli et al. (1989) or Fitzpatrick (2004) do not always provide a good fit to data and that it may be necessary to use more than one parameter to describe the extinction law (Nataf et al. 2015).

Modern extinction maps tend to be based either on a set of theoretical isochrones or some empirical determination of the colours and absolute magnitudes of stars along the main sequence and giant branch. It is not clear that available theoretical isochrones have the accuracy needed for extinction mapping, whilst empirical isochrones do not provide a good sampling of astrophysical parameters such as metallicity. However, we can hope that problems here will be remedied by the availability of Gaia data. Though, it is worth noting that efforts to employ Gaia data to infer empirically based isochrones will themselves be reliant on estimates of extinction.

While Gaia will have a significant impact across much of the sky, as an optical survey it cannot penetrate highly extinguished regions. In particular, it will offer little assistance in mapping extinction within the bulge/bar. So extinction maps of the bulge/bar will have to continue to rely on existing photometry (e.g. VVV and GLIMPSE) and spectroscopic (APOGEE) surveys, at least until the putative launch of JASMINE, not expected until the 2020s.

6. Conclusions

3d extinction maps are a key tool in the study of our Galaxy. They not only provide a 3d perspective on the ISM, but they are also a key requirement in the study of the Galaxy's stellar populations. In particular, the imminent arrival of Gaia data provides both a pressing need and clear opportunity for the production of maps. One one hand, Gaia data will potentially enable to production of far more precise and accurate maps than hitherto possible. While on the other hand, we are dependant on high quality 3d extinction maps in order to realise the full potential of Gaia for studying our Galaxy.

There have been many 3d maps of extinction produced to date, employing a variety of data and methods and looking at different regions of the sky. This contribution has attempted help classify these maps, discussing what common ground exists between different approaches and where they differ.

References

Amôres, E. B. & Lépine, J. R. D. 2005, AJ, 130, 659
Arenou, F., Grenon, M., & Gomez, A. 1992, A&A, 258, 104
Berry, M., Ivezić, Ž., Sesar, B., et al. 2012, ApJ, 757, 166
Bovy, J., et al. 2015, `arXiv:1509.06751`
Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, ApJ, 345, 245
Chen, B., Figueras, F., Torra, J., et al. 1999, A&A, 352, 459
Chen, B., et al. 1998, A&A, 336, 137
Chen, B.-Q., Liu, X.-W., Yuan, H.-B., et al. 2014, MNRAS, 443, 1192
Chen, B. Q., Schultheis, M., Jiang, B. W., et al. 2013, A&A, 550, A42
Drimmel, R., Cabrera-Lavers, A., & López-Corredoira, M. 2003, A&A, 409, 205
Ducati, J. R. 1986, Ap&SS, 126, 269
Farnhill, H. J., et al. 2015, MNRAS submitted
Fernie, J. D. 1962, AJ, 67, 224
Fitzgerald, M. P. 1968, AJ, 73, 983
Fitzpatrick, E. L. 1999, PASP, 111, 63
Fitzpatrick, E. L. 2004, in Astrophysics of Dust, eds. A. N. Witt, G. C. Clayton, & B. T. Draine (ASP, San Francisco), ASP Conf. Ser., 309, 33
Giammanco, C., Sale, S. E., Corradi, R. L. M., et al. 2011, A&A, 525, A58
Goldreich, P. & Sridhar, S. 1995, ApJ, 438, 763
Gontcharov, G. A. 2012, Astronomy Letters, 38, 87
Gonzalez, O. A., Rejkuba, M., Zoccali, M., et al. 2012, A&A, 543, A13
Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, ApJ, 810, 25
Jones, D. O., West, A. A., & Foster, J. B. 2011, AJ, 142, 44
Kolmogorov, A. 1941, Akademiia Nauk SSSR Doklady, 30, 301
Krautter, J. 1980, A&A, 89, 74
Lallement, R., Vergely, J.-L., Valette, B., et al. 2014, A&A, 561, A91
Lombardi, M. & Alves, J. 2001, A&A, 377, 1023
Lucke, P. B. 1978, A&A, 64, 367
Marshall, D. J., et al. 2006, A&A, 453, 635
Mendez, R. A. & van Altena, W. F. 1998, A&A, 330, 910
Nataf, D. M., Gonzalez, O. A., Casagrande, L., et al. 2015, `arXiv:1510.01321`
Neckel, T. 1966, Zeitschrift fur Astrophysik, 63, 221
Neckel, T. & Klare, G. 1980, A&AS, 42, 251
Ostriker, E. C., Stone, J. M., & Gammie, C. F. 2001, ApJ, 546, 980
Pandey, A. K. & Mahra, H. S. 1987, MNRAS, 226, 635
Parenago, P. P. 1945, Popular Astronomy, 53, 441
Perry, C. L. & Johnston, L. 1982, ApJS, 50, 451

Quiñonero-Candela, J. & Rasmussen, C. E. 2005, The Journal of Machine Learning Research, 6, 1939
Reid, M. J., Menten, K. M., Brunthaler, A., et al. 2014, ApJ, 783, 130
Sale, S. E. 2012, MNRAS, 427, 2119
Sale, S. E. 2015, MNRAS, 452, 2960
Sale, S. E., Drew, J. E., Barentsen, G., et al. 2014, MNRAS, 443, 2907
Sale, S. E., Drew, J. E., Unruh, Y. C., et al. 2009, MNRAS, 392, 497
Sale, S. E. & Magorrian, J. 2014, MNRAS, 445, 256
Scheffler, H. 1966, Zeitschrift fur Astrophysik, 63, 267
Schlafly, E. F., Finkbeiner, D. P., Schlegel, D. J., et al. 2010, ApJ, 725, 1175
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Schönrich, R., Aumer, M., & Sale, S. E. 2015, ApJ, 812, L21
Schultheis, M., Chen, B. Q., Jiang, B. W., et al. 2014, A&A, 566, A120
Spangler, S. R. & Gwinn, C. R. 1990, ApJ, 353, L29
Trumpler, R. J. 1930, PASP, 42, 214
Vallée, J. P. 2013, International Journal of Astronomy and Astrophysics, 3, 20
van de Kamp, P. 1930, AJ, 40, 145
Vázquez-Semadeni, E. & García, N. 2001, ApJ, 557, 727
Vergely, J.-L., et al. 2010, A&A, 518, A31
Zasowski, G., Majewski, S. R., Indebetouw, R., et al. 2009, ApJ, 707, 510