

Large-scale three-dimensional Gaussian process extinction mapping

S. E. Sale ¹ and J. Magorrian ²★

¹*Astrophysics Group, School of Physics, University of Exeter, Stocker Road, Exeter EX4 4QL, UK*

²*Rudolf Peierls Centre for Theoretical Physics, Keble Road, Oxford OX1 3NP, UK*

Accepted 2018 August 2. in original form 2018 June 11

ABSTRACT

Gaussian processes (GPs) are the ideal tool for modelling the Galactic interstellar medium, combining statistical flexibility with a good match to the underlying physics. In an earlier paper, we outlined how they can be employed to construct three-dimensional maps of dust extinction from stellar surveys. GPs scale poorly to large data sets though, which put the analysis of realistic catalogues out of reach. Here, we show how a novel combination of the expectation propagation method and certain sparse matrix approximations can be used to accelerate the dust mapping problem. We demonstrate, using simulated *Gaia* data, that the resultant algorithm is fast, accurate, and precise. Critically, it can be scaled up to map the *Gaia* catalogue.

Key words: methods: statistical – dust, extinction.

1 INTRODUCTION

Our Galaxy’s stellar content is the focus of many ambitious, large-scale surveys, the most notable of these being *Gaia* (Gaia Collaboration et al. 2016). The interpretation of the resulting stellar catalogues is hampered by the presence of interstellar dust, which becomes particularly severe near the Galactic plane. Any models for the structure and evolution of the Galaxy constructed from these data must therefore account for the three-dimensional (3D) distribution of extinguishing dust. On the other hand, dust is an important tracer of the densest parts of the Galaxy’s interstellar medium (ISM). Therefore, the 3D dust distribution is of interest in its own right.

In Sale & Magorrian (2014, hereafter [Paper I](#)), we developed a method for mapping interstellar extinction in three dimensions by modelling the (logarithm of the) extinction as a Gaussian process¹ (hereafter GP). [Paper I](#) builds on earlier work by Vergely et al. (2001) and overcomes the most blatant shortcomings encountered in previous work, such as Marshall et al. (2006), Chen et al. (2013), Sale et al. (2014), and Green et al. (2015): it banishes the ‘Fingers of God’ that plague other 3D extinction maps; the treatment of small-scale variations in extinction is much improved, enabling the extinction to a particular point in space to be estimated far more precisely than previously possible. Subsequently, Rezaei Kh. et al. (2017) have described a very similar method.

Modelling the logarithm of extinction as a GP results in a simple mathematical model that exhibits compelling parallels to the true ISM. The extinction to any given point (or set of points) in space is then described by a (multivariate) lognormal distribution, consis-

tent with theoretical arguments (Ostriker, Stone & Gammie 2001; Nordlund & Padoan 1999) that the integrated extinction to any given point is approximately proportional to the product of many small, independent relative compressions or rarefactions of the underlying density field along the line of sight to that point. In any GP, the statistical properties of the fluctuations about the mean are completely defined by the covariance function. In our formulation, the covariance function is directly related to the power spectrum of density fluctuations in the ISM. Our model includes a physical treatment of interstellar turbulence that provides us with a form for the power spectrum of the ISM. Specifically, we adopt that given by Kolmogorov (1941) model of turbulence, though other descriptions (e.g. Goldreich & Sridhar 1995) are available. This turbulent power spectrum is truncated at a spatial scale that corresponds to the energy injection scale in the ISM.

In [Paper I](#), we identified some practical challenges in applying our method to the construction of large-scale extinction maps. The purpose of this paper is to demonstrate how these challenges can be overcome. We begin in Section 2 by reviewing the topic of GP regression. We explain how we use it to map extinction, pointing out the practical difficulties that arise in Section 3. In Section 4, we describe a scheme that uses a combination of two approximate algorithms that make the production of large-scale extinction maps feasible. We verify the validity of our implementation in Section 5 by using it to reconstruct a 3D extinction map from simulated *Gaia* data, before summing up in Section 6.

2 GAUSSIAN PROCESS REGRESSION

We model (the logarithm of) the extinction to points in 3D space as a GP. As for any GP, the statistical properties of this density field are completely described by just two functions: a function

* E-mail: john.magorrian@physics.ox.ac.uk

¹GPs on 2D or 3D spaces are sometimes referred to as ‘Gaussian Random Fields’, indeed this was the terminology we employed in [Paper I](#).

$m(x)$ that gives the expectation value of the field at any position x , and a function $\Sigma(x_1, x_2)$ that returns the covariance of the field values between any pair of points (x_1, x_2) . Our method for mapping extinction in essence reduces to a GP regression problem. In this section, we provide a brief review of GPs and how they are employed in regression problems. Readers interested in a more comprehensive introduction are directed to Rasmussen & Williams (2005).

The defining property of a GP is that, for any set of $N \geq 1$ points $\mathbf{x} = \{x_1, \dots, x_N\}$ in this 3D space, the joint PDF (probability density function) of the values $\mathbf{f} = \{f_1, \dots, f_N\}$ of the GP at those points is Gaussian, namely

$$\mathbf{f} \sim \mathcal{N}(\mathbf{m}, \Sigma), \quad (1)$$

where \mathcal{N} indicates the multivariate Gaussian distribution, in which the mean vector $\mathbf{m}(\mathbf{x})$ and covariance matrix Σ have elements $m_n = m(x_n)$ and $\Sigma_{ij} = \Sigma(x_i, x_j)$, respectively.

Now suppose that we know the values \mathbf{f} of the field at some points \mathbf{x} , but want to infer values \mathbf{f}_* at some other points \mathbf{x}_* . From equation (1), it immediately follows that the joint distribution of \mathbf{f} and \mathbf{f}_* is

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{m}(\mathbf{x}) \\ \mathbf{m}(\mathbf{x}_*) \end{bmatrix}, \begin{bmatrix} \Sigma(\mathbf{x}, \mathbf{x}) & \Sigma(\mathbf{x}, \mathbf{x}_*) \\ \Sigma(\mathbf{x}, \mathbf{x}_*)^T & \Sigma(\mathbf{x}_*, \mathbf{x}_*) \end{bmatrix} \right), \quad (2)$$

in which the sub-matrices $\Sigma(\mathbf{a}, \mathbf{b})$ have elements $[\Sigma(\mathbf{a}, \mathbf{b})]_{ij} = \Sigma(a_i, b_j)$ and so on. Then, using standard properties of Gaussians (e.g. Bishop 2006), the conditional distribution of the \mathbf{f}_* is another Gaussian

$$\mathbf{f}_* | \mathbf{f} \sim \mathcal{N}(\mathbf{m}_{f_* | f}, \Sigma_{f_* | f}) \quad (3)$$

whose mean and covariance are given by

$$\begin{aligned} \mathbf{m}_{f_* | f} &= \mathbf{m}(\mathbf{x}_*) + \Sigma^T(\mathbf{x}, \mathbf{x}_*) \Sigma^{-1}(\mathbf{x}, \mathbf{x}) (\mathbf{f} - \mathbf{m}(\mathbf{x})), \\ \Sigma_{f_* | f} &= \Sigma(\mathbf{x}_*, \mathbf{x}_*) - \Sigma^T(\mathbf{x}, \mathbf{x}_*) \Sigma^{-1}(\mathbf{x}, \mathbf{x}) \Sigma(\mathbf{x}, \mathbf{x}_*). \end{aligned} \quad (4)$$

Therefore, with some simple matrix algebra, we can estimate the values \mathbf{f}_* of the field at the locations \mathbf{x}_* given (perfect) observations of it elsewhere.

This result is easily generalized to the situation in which our observations $\tilde{\mathbf{f}}$ of the underlying values \mathbf{f} are subject to some measurement error. Suppose that the measurement errors for each point are Gaussian and mutually independent, namely

$$\tilde{f}_n = f_n + \epsilon_n, \quad \text{where } \epsilon_n \sim \mathcal{N}(0, \sigma_n), \quad (5)$$

so that

$$\tilde{\mathbf{f}} \sim \mathcal{N}(\mathbf{m}(\mathbf{x}), \Sigma(\mathbf{x}, \mathbf{x}) + \Sigma_m), \quad (6)$$

where Σ_m is a diagonal matrix with the observational uncertainties σ_n along the diagonal. Then the joint distribution of $\tilde{\mathbf{f}}$ and \mathbf{f}_* is

$$\begin{bmatrix} \tilde{\mathbf{f}} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{m}(\mathbf{x}) \\ \mathbf{m}(\mathbf{x}_*) \end{bmatrix}, \begin{bmatrix} \Sigma(\mathbf{x}, \mathbf{x}) + \Sigma_m & \Sigma(\mathbf{x}, \mathbf{x}_*) \\ \Sigma(\mathbf{x}, \mathbf{x}_*)^T & \Sigma(\mathbf{x}_*, \mathbf{x}_*) \end{bmatrix} \right), \quad (7)$$

from which it follows that the posterior distribution of the \mathbf{f}_* is again a Gaussian

$$\mathbf{f}_* | \tilde{\mathbf{f}} \sim \mathcal{N}(\mathbf{m}_{f_* | \tilde{\mathbf{f}}}, \Sigma_{f_* | \tilde{\mathbf{f}}}) \quad (8)$$

whose mean and covariance are given by

$$\begin{aligned} \mathbf{m}_{f_* | \tilde{\mathbf{f}}} &= \mathbf{m}(\mathbf{x}_*) + \Sigma^T(\mathbf{x}, \mathbf{x}_*) [\Sigma(\mathbf{x}, \mathbf{x}) + \Sigma_m]^{-1} (\tilde{\mathbf{f}} - \mathbf{m}(\mathbf{x})), \\ \Sigma_{f_* | \tilde{\mathbf{f}}} &= \Sigma(\mathbf{x}_*, \mathbf{x}_*) - \Sigma^T(\mathbf{x}, \mathbf{x}_*) [\Sigma(\mathbf{x}, \mathbf{x}) + \Sigma_m]^{-1} \Sigma(\mathbf{x}, \mathbf{x}_*). \end{aligned} \quad (9)$$

So, we can estimate the values \mathbf{f}_* of the field at the locations \mathbf{x}_* given noisy observations of it elsewhere. This is known as GP

regression. GP regression is a practical solution in many regression problems, since it does not require one to assume some simple parametrized form for the data.

The implementation of GP regression can typically be divided into two stages. First there is learning/training phase in which inferences are made on any required latent or unknown parameters, especially those that define the mean and/or covariance functions, $m(x)$ and $\Sigma(x, x')$. Then, there is the prediction phase in which equations (8) and (9) are used to predict of the values of the process for locations where observations are not available.

Consider the situation where the hyperparameters, θ , that set the mean and covariance functions, are unknown. Then, in the learning phase, we seek to determine the posterior distribution of the hyperparameters given the observations $\tilde{\mathbf{f}}$,

$$p(\theta | \tilde{\mathbf{f}}) = \frac{p(\tilde{\mathbf{f}} | \theta) p(\theta)}{\int d\theta p(\tilde{\mathbf{f}} | \theta) p(\theta)} \quad (10)$$

where $p(\tilde{\mathbf{f}} | \theta)$ is a multivariate Gaussian PDF, as in equation (6) and $p(\theta)$ some hyperprior. Then, as the prediction phase, this can be combined with the expression (8) for $p(\mathbf{f}_* | \tilde{\mathbf{f}}, \theta)$ to give a posterior on \mathbf{f}_* :

$$p(\mathbf{f}_* | \tilde{\mathbf{f}}) = \int d\theta p(\mathbf{f}_* | \tilde{\mathbf{f}}, \theta) p(\theta | \tilde{\mathbf{f}}). \quad (11)$$

With all but the smallest data sets, the computational cost of calculating of the probabilities in the learning phase of GP regression is dominated by solving for the inverse of the covariance matrix, $[\Sigma(\mathbf{x}, \mathbf{x}) + \Sigma_m]^{-1}$ that appears in equation (9). Given N observations, the CPU time cost of solving for the inverse of the $N \times N$ matrix scales as $\mathcal{O}(N^3)$ and memory requirements as $\mathcal{O}(N^2)$.

3 EXTINCTION MAPPING AS A GAUSSIAN PROCESS REGRESSION PROBLEM

As in Paper I, we assume that $a(x)$, the logarithm of the extinction to the point x , is described by a GP. Then, the joint PDF $p(\mathbf{a} | \mathbf{l}, \mathbf{b}, s, \Theta, \zeta)$ of the log extinctions to many points is a multivariate Gaussian. The parameters Θ control how the mean extinction varies with position $x = (l, b, s)$. We demonstrated in Paper I how to derive the covariance function of extinction. This recognized that the 3D extinction map is a projection of the 3D density map, so that the extinction covariance function is obtained by integrating the density covariance along (any) two sightlines. As a result, the (extinction) covariance function depends on the small-scale turbulent structure of the ISM (ζ) and Θ . Table 1 provides a summary of our notation.

The remaining important ingredients of our model include a Galaxy model β that describes the distribution of the different stellar populations within the Galaxy and of the underlying stellar physics. This supplies us with a likelihood function $p(\tilde{\mathbf{y}} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, s, a, \beta)$ for the stellar observables $\tilde{\mathbf{y}}$ (such as apparent magnitudes and effective temperature) that depend on the distance s and log extinction a to the star, as well as the star's position within the galaxy (because the mixture of stellar populations changes with position). We assume that the observations, $\tilde{\mathbf{l}}$ and $\tilde{\mathbf{b}}$, of the Galactic coordinates of each star are measured exactly, so that we can condition all the following analysis on these values.²

²Formally, we assume the likelihood on the true Galactic coordinates of each star can be approximated by a delta function and then marginalize the true Galactic coordinates.

Table 1. A list of the parameters adopted in our model.

$(\tilde{l}_i, \tilde{b}_i)$	The (l, b) coordinates observed for the i th star in the catalogue.
\tilde{y}_i	All other directly observed quantities (e.g. broad-band fluxes, trigonometric parallax) for the i th star.
s_i	The distance to the i th star in the catalogue.
A_i	The extinction to the i th star in the catalogue.
a_i	The logarithm of extinction to the i th star in the catalogue, i.e. $a_i = \ln A_i$.
z_i	The component of the Gaussian mixture model used to describe the likelihood $p(\tilde{y}_i, s_i, z_i a_i, \tilde{l}_i, \tilde{b}_i, \zeta, \beta)$.
Θ	Parameters that set the large-scale distribution of extinction.
ζ	The model for the small-scale turbulent variations of extinction.
β	The background Galaxy model, including prior on position, metallicity, age, etc.

Therefore, the probability of anything in our model can be obtained simply by marginalizing the innocuous looking

$$p(\tilde{l}, \tilde{b}, s, \tilde{y}, \mathbf{a}, \Theta, \zeta, \beta) = p(\beta)p(\zeta)p(\Theta|\zeta)p(\tilde{l}, \tilde{b}, s|\beta)p(\mathbf{a}|\tilde{l}, \tilde{b}, s, \Theta, \zeta)p(\tilde{y}|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta), \quad (12)$$

in which the first three factors are our model's priors. The fourth, $p(\tilde{l}, \tilde{b}, s|\beta)$, is the spatial density of *all* stars in our assumed Galaxy model β and provides a prior on the position of each observed star. We assume that this can be factorized so that the prior on each star is independent of all other stars:

$$p(\tilde{l}, \tilde{b}, s|\beta) = \prod_{n=1}^N p(s_n|\tilde{l}_n, \tilde{b}_n, \beta)p(\tilde{l}_n, \tilde{b}_n|\beta). \quad (13)$$

The last two factors in equation (12) are more challenging. As discussed above, $p(\mathbf{a}|\tilde{l}, \tilde{b}, s, \Theta, \zeta)$ is just a multivariate Gaussian, the only difficulty being the sheer size of its covariance matrix. The final factor

$$p(\tilde{y}|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta) = \prod_{n=1}^N p(\tilde{y}_n|\tilde{l}_n, \tilde{b}_n, s_n, a_n, \beta) \quad (14)$$

relates each star's observables \tilde{y}_n to its position within the Galaxy and the (log) extinction between us and it.

In Sale & Magorrian (2015, see equations 14, 16, and 21 of that paper), we showed that each star's contribution to the marginal likelihood equation (14) is well approximated by a mixture of Gaussians in $(a, \log s)$. Although it is often possible to identify the components in such mixtures with, e.g. distinct stellar luminosity classes, we stress that this Gaussian-mixture approximation is introduced primarily for computational convenience: we attach no physical significance to the underlying Gaussians. For each star, it is convenient to introduce a latent variable z , which is a vector that indicates which Gaussian the star was drawn from: if there are M Gaussians in the mix representing $p(\tilde{y}|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta)$ then z has M elements, just one of which is equal to 1, the rest being zero (i.e. one hot). Then,

$$p(\tilde{y}, z|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta) = \prod_{m=1}^M [G_m(\log s, a)]^{z_m}, \quad (15)$$

where $G_m(\log s, a)$ is the m th Gaussian, whose mass, mean, and covariance depend on $(\tilde{y}, \tilde{l}, \tilde{b})$.

Expressed in terms of the full set $\mathbf{z} = (z_1, \dots, z_N)$ of membership indicators, the probability of everything in our model becomes

$$p(\tilde{l}, \tilde{b}, s, \tilde{y}, \mathbf{a}, \mathbf{z}, \Theta, \zeta, \beta) \approx p(\beta)p(\zeta)p(\Theta|\zeta)p(\tilde{l}, \tilde{b}|\beta)p(\Theta|\zeta) \times p(\mathbf{a}|\tilde{l}, \tilde{b}, s, \Theta, \zeta)p(\tilde{y}, \mathbf{z}|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta), \quad (16)$$

where

$$p(\tilde{y}, \mathbf{z}|\tilde{l}, \tilde{b}, s, \mathbf{a}, \beta) = \prod_{n=1}^N \prod_{m=1}^{M_n} [G_{nm}(\log s_n, a_n)]^{z_{nm}}. \quad (17)$$

Marginalizing \mathbf{z} in this expression gives the more physically meaningful equation (12), hiding the details of our Gaussian mixture parametrization.

3.1 What do we want to know?

As discussed in Paper I, there are four broad tasks contained under the umbrella of extinction mapping:

- (i) inferring the extinction and/or distance to stars within an observed catalogue;
- (ii) inferring the extinction to some arbitrary locations;
- (iii) determining the large-scale distribution of extinction and, by extension, dust; and
- (iv) constraining the turbulent physics of the ISM by studying how extinction varies on small scales.

All of these tasks are solved by marginalizing equation (16) appropriately. For example, finding distances and extinctions of the stars within the catalogue, task (i), involves the posterior

$$p(s, \mathbf{a}, \Theta|\tilde{y}, \tilde{l}, \tilde{b}, \zeta, \beta) = \frac{1}{p(\tilde{y}, \tilde{l}, \tilde{b}, \zeta, \beta)} \sum_z p(\tilde{l}, \tilde{b}, s, \tilde{y}, \mathbf{a}, \mathbf{z}, \Theta, \zeta, \beta), \quad (18)$$

where the evidence $p(\tilde{l}, \tilde{b}, \tilde{y}, \zeta, \beta)$ in the denominator, which depends only on directly observed variables, is an uninteresting normalization constant unless one wants to compare different assumptions for ζ or β . Marginalizing \mathbf{a} and s from this gives the posterior distribution of the large-scale distribution of dust Θ ,

$$p(\Theta|\tilde{y}, \tilde{l}, \tilde{b}, \zeta, \beta) = \frac{1}{p(\tilde{y}, \tilde{l}, \tilde{b}, \zeta, \beta)} \sum_z \int ds \int d\mathbf{a} p(\tilde{l}, \tilde{b}, s, \tilde{y}, \mathbf{a}, \mathbf{z}, \Theta, \zeta, \beta), \quad (19)$$

which solves task (iii). Given the Gaussian mixture description, we adopt from Sale & Magorrian (2015), the inner integral over \mathbf{a} can be found trivially for any choice of \mathbf{z} and s . Similarly, multiplying equation (19) by the prior $p(\zeta)$ and marginalizing Θ gives the posterior distribution of ζ . For the time being, however, we will ignore the possibility of examining the small-scale characteristics of dust, although we intend to return to this issue in a future paper.

Determining the distributions above essentially constitutes the learning phase of the GP regression. Subsequently in the prediction phase of the GP regression, we can estimate the extinction to some arbitrary points in space, i.e. task (ii). If performed on some regular

grid of locations this constitutes the production of an extinction map. We can extend the probability of everything [equation (16)] to include the PDF of extinctions \mathbf{a}_* to given set of locations $\mathbf{x}_* \equiv (\mathbf{l}_*, \mathbf{b}_*, \mathbf{s}_*)$. The result is

$$\begin{aligned} p(\tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \tilde{\mathbf{y}}, \mathbf{a}, \mathbf{a}_*, \mathbf{z}, \Theta, \zeta, \beta | \mathbf{x}_*) \\ = p(\beta) p(\zeta) p(\tilde{\mathbf{l}}, \tilde{\mathbf{b}} | \beta) p(\Theta | \zeta) \\ \times p(\mathbf{a}, \mathbf{a}_* | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \mathbf{x}_*, \Theta, \zeta) p(\tilde{\mathbf{y}}, \mathbf{s}, \mathbf{z} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{a}, \zeta, \beta), \end{aligned} \quad (20)$$

which differs from equation (16) only in the $p(\mathbf{a}, \mathbf{a}_* | \dots)$ factor. Marginalizing,

$$\begin{aligned} p(\mathbf{a}_* | \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{x}_*, \zeta, \beta) \\ = \frac{1}{p(\tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta)} \int d\Theta \sum_{\mathbf{z}} \int d\mathbf{s} \int d\mathbf{a} p(\tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \tilde{\mathbf{y}}, \mathbf{a}, \mathbf{a}_*, \mathbf{z}, \Theta, \zeta, \beta | \mathbf{x}_*), \end{aligned} \quad (21)$$

where, as in equation (19) above, much of the marginalization can be carried out ‘by hand’.

3.2 Comparison with standard GP regression

Notice that the PDF equation (16) is a Gaussian in \mathbf{a} : if the stellar distances \mathbf{s} and membership indicators \mathbf{z} were known, then our extinction-mapping job would reduce to a straightforward GP regression problem. In practice, however, stellar distances \mathbf{s} are typically very uncertain. One way of dealing with this is to project the distance uncertainties onto extinction by multiplying the distance uncertainties by the mean gradient of extinction with respect to distance, i.e. by using a first-order Taylor expansion. This approach has been used by Vergely et al. (2001) in their extinction maps. McHutchon & Rasmussen (2011) discuss, in significantly greater detail, a similar approach for the general GP regression problem. This first-order Taylor expansion is a good approximation only when the uncertainties on distance are small. In practice this is not the case, especially when the data used to estimate stellar distances is limited to photometry (Green et al. 2014; Sale & Magorrian 2015). Moreover, the mean gradient of extinction with respect to distance is unknown. Therefore, we must instead sample the unknown distances of the stars in the learning phase. In Paper I, this was achieved through the use of MCMC (Markov chain Monte Carlo). A similar approach was subsequently discussed by Cervone & Pillai (2015).

As in standard GP regression, we do not know the precise (log) extinction a to each star, but instead possess only some noisy estimate of it. This is much easier to deal with: as shown in the derivation of equations (8) and (9) above, it is straightforward to account for measurement errors in the values of a GP assuming the uncertainties are normally distributed.

These observational uncertainties are therefore straightforward to deal with, at least in principle. The main *practical* difficulty in applying GP regression to the extinction mapping problem is sheer computational expense. We are forced to infer or sample from the joint distributions that include \mathbf{s} and \mathbf{z} . Each such sample requires the calculation of $p(\mathbf{a} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \zeta)$. For a sample of N stars, this involves finding the inverse of an $N \times N$ matrix, which scales in CPU time as $\mathcal{O}(N^3)$ and in memory as $\mathcal{O}(N^2)$. Given that we now live in an era of large surveys that contain $\sim 10^8$ – 10^9 objects, this is a clear barrier to applying the method of Paper I.

4 ACCELERATED EXTINCTION MAPPING

We have just seen that the fundamental bottleneck in the scheme of Paper I is the calculation of the joint PDF $p(\mathbf{a} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \zeta)$ of

the extinctions to all N stars in the catalogue. Therefore, it makes sense to look for ways of approximating this PDF. Given the large size of modern survey catalogues, any such approximation should not only deliver significant speed and memory savings, but also allow easy parallelization to allow the catalogue to be spread across multiple nodes. In this section, we describe a combination of two such schemes: expectation propagation (EP, Minka 2001) and the partially independent conditional approximation (PIC, Snelson & Ghahramani 2007). Appendix A gives a more detailed overview of these and similar approaches for accelerating GP regression.

In order to simplify the exposition of our approach, we initially assume that the distances \mathbf{s} to the stars, their Gaussian mixture components \mathbf{z} and the dust microphysics ζ are known exactly a priori.

4.1 Acceleration using PIC and EP: simplified case

Let us begin by focusing on the learning phase of GP regression. In the simplified case in which we know \mathbf{s} and \mathbf{z} , this reduces to estimating the posterior distribution of Θ , the hyperparameters that determine the large-scale distribution of dust. Employing the assumptions above allows us to simplify the posterior equation (12), so that

$$p(\mathbf{a}, \Theta | \tilde{\mathbf{y}}) \propto p(\tilde{\mathbf{y}} | \mathbf{a}) p(\mathbf{a} | \Theta) p(\Theta), \quad (22)$$

and, marginalizing,

$$p(\Theta | \tilde{\mathbf{y}}) \propto \int d\mathbf{a} p(\tilde{\mathbf{y}} | \mathbf{a}) p(\mathbf{a} | \Theta) p(\Theta). \quad (23)$$

For clarity, we have suppressed the implicit conditioning on the assumed $(\tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \mathbf{z}, \zeta, \beta)$.

We begin by splitting the N stars into K sub-catalogues, so that

$$\mathbf{a} = \begin{pmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_K \end{pmatrix}, \quad (24)$$

where the k th sub-catalogue contains N_k stars and

$$N = \sum_{k=1}^K N_k. \quad (25)$$

Directly observed stellar quantities, $\tilde{\mathbf{y}}$, are also partitioned into a matching set of K sub-catalogues. Typically, the partitioning is based on the on-sky position of the stars, so that sub-catalogues cover non-overlapping regions of the sky.

As we employ the Gaussian mixture model (GMM) approximation (15) and assume we know \mathbf{s} and \mathbf{z} , the likelihood for each star follows

$$\tilde{a}_n | a_n \sim \mathcal{N}(a_n, \sigma_n), \quad (26)$$

where \tilde{a}_n is the ‘observed’ log-extinction to the n th star, so that $\tilde{\mathbf{a}}$ is a sufficient statistic for $\tilde{\mathbf{y}}$ in this context. Further assuming that observations of one star are independent of observations of all other stars, implies that

$$p(\tilde{\mathbf{y}} | \mathbf{a}) = \prod_{k=1}^K p(\tilde{\mathbf{a}}_k | \mathbf{a}_k) \quad (27)$$

Unfortunately, we cannot directly partition the expression for $p(\mathbf{a} | \Theta)$ – the bottleneck in our method – in this way. Because

extinctions to neighbouring points in space are inevitably correlated, we instead have that

$$p(\mathbf{a}|\Theta) = \prod_{k=1}^K p(\mathbf{a}_k|\mathbf{a}_1, \dots, \mathbf{a}_{k-1}, \Theta), \quad (28)$$

which means that it is impossible to partition a spatially contiguous catalogue into discrete, disjoint, uncorrelated sub-catalogues. We see that the extinctions in the k th sub-catalogue depend on Θ and the extinctions in the preceding $k - 1$ sub-catalogues. Calculating $p(\mathbf{a}|\Theta)$ in this factorized form is no faster than in its basic unfactorized form.

Instead we make the partially independent (training) conditional approximation (PI(T)C; Snelson & Ghahramani 2007, see also Appendix A1), which involves introducing M appropriately chosen *inducing points*, which are spatially distributed throughout the full catalogue. The joint PDF is approximated as

$$p(\mathbf{a}|\mathbf{u}, \Theta) \approx \prod_{k=1}^K p(\mathbf{a}_k|\mathbf{u}, \Theta), \quad (29)$$

where \mathbf{u} gives the values of the GP at the full set of inducing points. By making this PI(T)C approximation, we are assuming that the extinctions in different sub-catalogues are independent when conditioned on \mathbf{u} and Θ . With this factorization, we have a cheap means of approximating $p(\mathbf{a}|\mathbf{u}, \Theta)$. The inclusion of \mathbf{u} induces correlations among the values of \mathbf{a} in different sub-catalogues. We discuss the choice of locations of these inducing points later.

Given the full catalogue of observed extinctions $\tilde{\mathbf{a}}$, the joint posterior distribution of \mathbf{u} and Θ is then

$$\begin{aligned} p(\mathbf{u}, \Theta|\tilde{\mathbf{a}}) &\propto \int d\mathbf{a} p(\Theta) p(\mathbf{u}|\Theta) \prod_{k=1}^K (p(\tilde{\mathbf{a}}_k|\mathbf{a}_k) p(\mathbf{a}_k|\mathbf{u}, \Theta)) \\ &= p(\Theta) p(\mathbf{u}|\Theta) \prod_{k=1}^K \left(\int d\mathbf{a}_k p(\tilde{\mathbf{a}}_k|\mathbf{a}_k) p(\mathbf{a}_k|\mathbf{u}, \Theta) \right). \end{aligned} \quad (30)$$

We next use the EP algorithm (Minka 2001) to construct an approximation of the form

$$q(\mathbf{u}, \Theta) = p(\Theta) p(\mathbf{u}|\Theta) \prod_{k=1}^K q_k(\mathbf{u}, \Theta) \quad (31)$$

to this posterior. The resulting $q_k(\mathbf{u}, \Theta)$ factors are known as *site distributions*. Their functional form is chosen to suit the problem at hand. Here, we assume that each q_k is some multivariate Gaussian.

Below we set out the steps involved in our implementation of EP. The reader interested in *why* it works should consult Appendix A2 and references therein.

The EP algorithm uses a couple of sets of distributions that are constructed from the posterior equation (30). The first is the set of *cavity distributions* (see also equation A22) obtained by omitting the k th site distribution from the approximated posterior equation (31):

$$q_{\setminus k}(\mathbf{u}, \Theta) = \frac{1}{Z_k} p(\Theta) p(\mathbf{u}|\Theta) \prod_{l \neq k} q_l(\mathbf{u}, \Theta), \quad (32)$$

in which Z_k is a normalization factor. As our assumed $q_k(\mathbf{u}, \Theta)$ is Gaussian and the factor $p(\mathbf{u}|\Theta)$ is by definition Gaussian, then as long as our prior $p(\Theta)$ is Gaussian these cavity distributions will themselves be Gaussian. We further define the *reduced cavity*

distributions

$$\begin{aligned} q'_{\setminus k}(\mathbf{u}, \Theta) &= \frac{q_{\setminus k}(\mathbf{u}, \Theta)}{p(\Theta) p(\mathbf{u}|\Theta)} \\ &= \frac{1}{Z_k} \prod_{l \neq k} q_l(\mathbf{u}, \Theta). \end{aligned} \quad (33)$$

The second is the set of *tilted distributions* (see also equation A23)

$$q_{\setminus k}(\mathbf{u}, \Theta) = q'_{\setminus k}(\mathbf{u}, \Theta) p(\Theta) p(\tilde{\mathbf{a}}_k, \mathbf{u}|\Theta). \quad (34)$$

Again, these are Gaussian.

After we make an initial guess for the site distributions $\{q_k\}$, the EP algorithm proceeds as follows:

- (i) construct the sets of cavity and tilted distributions, $q_{\setminus k}$ and $q'_{\setminus k}$;
- (ii) for each $k = 1, \dots, K$, construct an improved estimate, q_k^{new} , of the site distribution q_k by matching the first- and second-order moments of q_k^{new} to those of the tilted distribution $q_{\setminus k}/Z_k$.
- (iii) set $q_k = q_k^{\text{new}}$ and repeat until converged.

The CPU and memory costs of carrying out this algorithm are dominated by the $p(\tilde{\mathbf{a}}_k, \mathbf{u}|\Theta)$ factor in the tilted distribution equation (34). The CPU time needed therefore scales as $\mathcal{O}(N_i + M)^3$ and the memory as $\mathcal{O}(N_i + M)^2$.

With some further manipulation we can improve on this scaling, however. Factorize the expensive $p(\tilde{\mathbf{a}}_k, \mathbf{u}|\Theta)$ factor into the product of Gaussians $p(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta) p(\tilde{\mathbf{a}}_k|\Theta)$ and introduce $q'_{\setminus k}(\mathbf{u}|\Theta)$, defined via

$$q'_{\setminus k}(\mathbf{u}, \Theta) = q'_{\setminus k}(\mathbf{u}|\Theta) p(\Theta). \quad (35)$$

Then, the tilted distribution equation (34) becomes

$$q_{\setminus k}(\mathbf{u}, \Theta) = q_{\setminus k}(\Theta) q'_{\setminus k}(\mathbf{u}|\Theta) p(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta) p(\tilde{\mathbf{a}}_k|\Theta). \quad (36)$$

The product of Gaussians $q'_{\setminus k}(\mathbf{u}|\Theta) p(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ appearing here can be written (e.g. Murphy 2007) as a single Gaussian $\bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ multiplied by a normalization factor $F_k(\tilde{\mathbf{a}}_k, \Theta)$. So, we then have

$$q_{\setminus k}(\mathbf{u}, \Theta) = F_k(\tilde{\mathbf{a}}_k, \Theta) q_{\setminus k}(\Theta) \bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta) p(\tilde{\mathbf{a}}_k|\Theta). \quad (37)$$

We provide explicit expressions for $\bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ and $F_k(\tilde{\mathbf{a}}_k, \Theta)$ in Appendix B.

The CPU and memory costs associated with sampling from the tilted distribution in this form approximately scale as $\mathcal{O}(N_k^3)$ and $\mathcal{O}(N_k^2)$ respectively, an improvement over the formulation above. Moreover, all of the terms in this form have simple and intuitive meanings:

- (i) The cavity distribution $q_{\setminus k}(\Theta)$ is, by design, Gaussian and provides the constraints offered by the prior and the other sub-catalogues on Θ .
- (ii) $\bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ is a Gaussian distribution on \mathbf{u} whose mean and covariance depend on $\tilde{\mathbf{a}}_k$ and Θ . It is the predicted distribution of \mathbf{u} given the observations in the k th sub-catalogue, the information provided by other catalogues through the cavity distribution and the GP prior that depends on Θ .
- (iii) $F_k(\tilde{\mathbf{a}}_k, \Theta)$ measures the agreement between the prediction for \mathbf{u} from the cavity distribution and the prediction implied by $\tilde{\mathbf{a}}_k$.
- (iv) Finally, $p(\tilde{\mathbf{a}}_k|\Theta)$ is the simple GP likelihood obtained by considering the k th sub-catalogue in isolation.

We can also illustrate the significance of the factors in the cavity distribution by considering an alternative approximation. Assume that the extinctions in different blocks are independent, when conditioned on Θ . That is, suppose that, instead of the PITC approxi-

mation (29), we assume that

$$p(\mathbf{a}, |\Theta) \approx \prod_{k=1}^K p(\mathbf{a}_k | \Theta), \quad (38)$$

which is equivalent to a block-diagonal approximation to the covariance matrix, with the off diagonal blocks set to zero. It is also equivalent to the PITC approximation with an empty set of inducing points. Working through a similar derivation to that above, we obtain the tilted distribution

$$q_{\setminus k}(\Theta) = q_{\setminus k}(\Theta) p(\tilde{\mathbf{a}}_k | \Theta). \quad (39)$$

Both factors in this block-diagonal tilted distribution also appear in the PITC tilted distribution (37). Thus, the two additional factors in the PITC tilted distribution (37) provide the approximation to the effect of the off-diagonal blocks in the covariance matrix.

Constructing and sampling from the tilted distribution is the most difficult stage in the application of EP to our simplified extinction mapping. As all the prior distributions and site approximations are Gaussian, multiplying them to obtain the overall joint approximation $q(\mathbf{u}, \Theta)$ and subsequently finding the cavity distributions are both straightforward. In addition, updating the site approximations can be done by simply moment matching to the tilted distributions.

4.2 Acceleration using PIC and EP: general case

We now consider how to carry out the learning phase of our GP extinction mapping in realistic cases, by relaxing some of the assumptions made in Section 4.1. In particular, we will make no assumptions about \mathbf{s} and \mathbf{z} . For the time being we will continue to assume a dust power spectrum (ζ), whilst the assumption that the observed on-sky position of the stars is exactly correct is essentially always reasonable.

The fundamental approach we employ remains unchanged: we use a combination of the PIC approximation and EP. Given that we no longer know the distances to the stars, $\tilde{\mathbf{a}}$ is not a sufficient statistic for $\tilde{\mathbf{y}}$. Instead, given some \mathbf{s} and \mathbf{z} , $(\tilde{\mathbf{a}}, \tilde{\mathbf{s}})$ is a sufficient statistic, where $\log \tilde{s}_n$ is the mean log distance of the component in the GMM approximation to the likelihood of the n th star indicated by \mathbf{z}_n , whilst \tilde{a}_n is the mean log-extinction conditioned on \mathbf{z}_n and \mathbf{s}_n .

We may split $p(\mathbf{s}, \mathbf{a}, \mathbf{z} | \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta, \text{flat})$ into three factors,

$$p(\mathbf{s}, \mathbf{a}, \mathbf{z} | \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta, \text{flat}) = p(\mathbf{z} | \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta, \text{flat}) \times p(\mathbf{s} | \mathbf{z}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta, \text{flat}) \times p(\mathbf{a} | \mathbf{s}, \mathbf{z}, \tilde{\mathbf{a}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta, \text{flat}). \quad (40)$$

As we approximate the left-hand side of this equation with a GMM equation (15), as described in Sale & Magorrian (2015), it follows that the final two terms on the right-hand side can be approximated as Gaussians, where

$$\log s_n \sim \mathcal{N}(\log \tilde{s}_n(\mathbf{z}_n), \sigma_n^{(s)}(\mathbf{z}_n)) \quad (41)$$

$$a_n \sim \mathcal{N}(\tilde{a}_n(\mathbf{s}_n, \mathbf{z}_n), \sigma_n^{(a)}(\mathbf{z}_n)). \quad (42)$$

Here, $\tilde{s}_n(\mathbf{z}_n)$ is the mean distance of the component in the GMM approximation \mathbf{z}_n and $\tilde{a}_n(\mathbf{s}_n, \mathbf{z}_n)$ the mean log-extinction implied by \mathbf{z}_n and \mathbf{s}_n .

4.2.1 Learning phase

In addition to Θ and \mathbf{u} , we now also need to infer \mathbf{s} and \mathbf{z} . Fortunately, both \mathbf{s} and \mathbf{z} are ‘local’ as described in Appendix A2: the observations in the k th sub-catalogue depend only on \mathbf{s}_k and \mathbf{z}_k and

are conditionally independent of $\mathbf{s}_{j \neq k}$ and $\mathbf{z}_{j \neq k}$. Thus, the EP site, cavity and tilted approximations do not depend on \mathbf{s} or \mathbf{z} .

By a similar derivation to that given in Section 4.1, the tilted distributions are

$$q_{\setminus k}(\mathbf{u}, \Theta) = \sum_{\mathbf{z}_k} \int d\mathbf{s}_k q_{\setminus k}(\Theta) \bar{q}_k(\mathbf{u} | \tilde{\mathbf{a}}_k, \mathbf{s}_k, \mathbf{z}_k, \Theta) F_k(\tilde{\mathbf{a}}_k, \mathbf{s}_k, \mathbf{z}_k, \Theta) p(\tilde{\mathbf{a}}_k | \Theta, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) p(\tilde{\mathbf{s}}_k, \mathbf{s}_k, \mathbf{z}_k | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \beta). \quad (43)$$

This is the generalization of the tilted distribution in the simplified case (37). As we do not know \mathbf{s}_k nor \mathbf{z}_k a priori, we have an additional factor that gives their probability, which we then marginalize. We currently employ a Metropolis within Gibbs (Tierney 1994) MCMC algorithm to sample from the intergrand above. The marginalization to obtain samples from $q_{\setminus k}(\mathbf{u}, \Theta)$ is then trivial. A new site approximation is then obtained by moment matching.

4.2.2 Prediction phase: mapmaking

Having an estimate of the joint posterior distribution of $(\mathbf{s}, \mathbf{z}, \mathbf{u}, \Theta)$, we can embark on the prediction phase of GP regression and infer the extinction to a regular grid of points and thereby produce a 3D map of extinction. In Appendix C, we derive an approximation to the predictive distribution,

$$p(\mathbf{a}_{\star, k} | \tilde{\mathbf{y}}, \mathbf{x}_{\star}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \approx \iint d\mathbf{u} d\Theta \sum_{\mathbf{z}_k} \int d\mathbf{s}_k \left(q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta) p(\mathbf{a}_{\star, k} | \mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right), \quad (44)$$

where $q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta)$ is the local unmarginalized tilted distribution for the k th block of data as sampled from in each EP iteration. Note, in particular, that we are now able to carry out the prediction phase, by focusing on each sub-catalogue in turn. This allows trivial parallelization of the training phase and is a direct consequence of combining the PIC approximation with EP. The communication and smoothing between the different sub-catalogues is provided by the shared \mathbf{u} and Θ : the sub-catalogues are computationally independent while remaining statistically dependent.

In practice, we obtain samples from $q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta)$ using MCMC. As a result, estimating the moments of $p(\mathbf{a}_{\star, k} | \tilde{\mathbf{y}}, \mathbf{x}_{\star}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta)$ is easy. Specifically, for each sampled $(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta)$, we then obtain samples of $\mathbf{a}_{\star, k}$ from the Gaussian $p(\mathbf{a}_{\star, k} | \mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta)$. We can then estimate the expectation of $\mathbf{a}_{\star, k}$ by taking the expectation of the samples of all the $\mathbf{a}_{\star, k}$ and similarly for higher moments.

4.2.3 Scaling

If we split a catalogue of N stars into K smaller catalogues, each of size M , then the total CPU cost of the training phase in the EP scheme will scale as $\mathcal{O}(NM^2)$, whilst memory costs will scale as $\mathcal{O}(NM)$. As $M = N/K < N$, this is significantly more favourable than the $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$ CPU and memory scaling of vanilla GP regression. Meanwhile, the CPU costs of the prediction phase now scale as $\mathcal{O}(NM)$ rather than $\mathcal{O}(N^2)$. In addition, both phases are trivially parallelizable, with minimal communication between processes only required when forming the new $q(\mathbf{u}, \Theta)$ at the end of each of the few (typically ≤ 5) EP iterations.

In practice, we can divide large observed catalogues into smaller sub-catalogues of roughly constant size, e.g. by having the sub-

catalogues cover equal areas on the sky. Now if we add more similarly sized sub-regions to the map, the cost of producing larger extinction maps will scale linearly with the number of stars N . When allied to the fact that EP permits trivial parallelization, with minimal communication between nodes, we now have an extinction mapping method that can be applied to larger catalogues, such as those that surveys such as *Gaia* are producing. This vastly accelerated method relies on only two approximations: the PIC approximation and, in using EP, we have approximated the site distributions with a multivariate Gaussian $q_k(\Theta, \mathbf{u})$.

The astute reader will have noticed that with very large catalogues we are essentially presented with two options: increase the number of inducing points or reduce the ratio of inducing points to observations. If the spatial coverage of our data set is growing the second of these two options would result in the inducing points not being able to maintain an adequate spatial sampling relative to the scale of the covariance function. As a consequence, there would be an increased loss of precision in the extinction maps produced and Fingers of God would begin to appear as the correlation between blocks becomes weaker (Vanhatalo, Pietiläinen & Vehtari 2010). Therefore, we seek to avoid it as far as possible. Increasing the number of inducing points would normally result in increased computation cost, just as increasing the number of observations increases the cost of vanilla GP regression. Specifically, the contribution to the CPU time needed to invert the covariance matrix at of the R inducing points, as needed in PIC-based GP regression, scales as $\mathcal{O}(R^3)$. When R is small, this contribution is not significant relative to the $\mathcal{O}(NM^2)$ cost related to dealing with the observations. However, given the cubic scaling, it can quickly become important as the number of inducing points increases. Thus, it would appear to be prohibitive to employ a large number of inducing points.

However, we note that there will be a substantial distance between many of these inducing points and so many pairs of inducing points will be essentially uncoupled. Consequently, it may be possible to employ a scheme that leverages the resulting approximate sparsity in the covariance matrix. One possibility would be to utilize the hierarchical off-diagonal low-rank (HODLR) scheme (Ambikasaran et al. 2016), as described in Appendix A.

Alternatively, it is possible to divide the total set of inducing points into a smaller number of subsets, $\mathbf{u} = \{\mathbf{u}_1, \dots, \mathbf{u}_L\}$. As the values at the inducing points are themselves drawn from the underlying GP, we can then apply the PIC approximation to them, introducing a set of ‘hyperinducing points’ \mathbf{v} . In doing so, we establish a clear hierarchy. At the base, we have our observations, divided into sub-catalogues. A number of these sub-catalogues are tied together by a subset of inducing points under the PIC approximation. Finally, all the subsets are related by again applying the PIC approximation and the hyperinducing points. This describes a three-level hierarchy, though there is no reason it could not be extended to contain more levels as required by the size of the data.

A hierarchical PIC scheme of this form is well matched to computational infrastructure: one can easily imagine computational cluster nodes being hierarchically grouped in a similar manner. Moreover, with HEALPIX (Górski et al. 2005), it is easy to group observations into the required hierarchical structure. Consequently, with a hierarchical PIC-EP scheme, it would be possible to map extinction employing catalogues of essentially limitless size, with the CPU time and memory needed scaling only linearly with the total size of the catalogue and with the option of straightforward parallelization.

5 TESTS WITH SIMULATED GAIA DATA

As in Paper I, we employ simulated data to assess the quality of the inferences made our EP-based algorithm. Clearly, *Gaia* will provide the most significant data set for extinction mapping in the near future. With that in mind, we test our implementation, which we name G-MEAD (GP Mapping of Extinction Against Distance), on synthetic *Gaia* observations produced by the Gaia Object Generator (GOG, Luri et al. 2014). GOG produces synthetic observations, including parallaxes and extinctions, for stars in the Gaia Universe Model (GUM, Robin et al. 2012).

The GUM employs the Drimmel, Cabrera-Lavers & López-Corredoira (2003, hereafter D03) Galactic dust model to describe extinction within the Galaxy. The D03 model includes features such as a spiral arms and normalizes extinction along each sightline to that given by Schlegel, Finkbeiner & Davis (1998). It does not, however, include a direct simulation of interstellar turbulence. As a result, it does not necessarily square with the GP-based statistical model we employ and is not the perfect simulation on which to test our method. However, we still consider the GOG simulations a useful test since they are otherwise a very faithful reproduction of what we might expect from the full Gaia catalogue.

In practical terms, the non-simulation of interstellar turbulence has a number of effects. First, we impose a relatively small value of the ratio between the standard deviation and mean of dust (3D) density. If we set this value to zero we would be implicitly assuming that the distribution of dust (and so extinction) could be entirely described by our chosen mean function. By setting it to a small value we are allowing a little variation from the mean. This is necessary because the normalization of the D03 map to Schlegel et al. (1998) impose some turbulent-like features on the data. An effect of the small value of the variance ratio is that our model is less expressive than it otherwise might be. This means it has reduced ability to describe features that do not appear in our mean function. In addition, our statistical model is more complicated than that behind the data we wish to fit, thus we are at risk of overfitting. To mitigate this, we impose a floor on the covariance arrays of the cavity distributions.

Currently, our implementation contains no treatment for the (typically apparent-magnitude based) incompleteness of catalogues. As demonstrated by Sale (2015) selection effects stemming from this incompleteness can have a severe and pathological impact on extinction maps. Unfortunately, the approach developed in Sale (2015) cannot be directly applied here. Therefore, we sidestep this issue by compiling a volume-limited test catalogue. We achieve this by including only stars with masses greater than $1.2M_\odot$ and distances less than 5000 pc. In addition, we also discard the least informative stars, jettisoning those with relative parallax errors above one half.

Our test catalogue ultimately contains 6349 stars within a 1 sq. deg. area defined by $174^\circ \leq l < 186^\circ$ and $|b| < 2.5$ arcmin. We adopt this ‘letter-box’ area to make it easier to plot maps on paper. We then sub-divide this into 36 sub-catalogues of 20 arcmin \times 5 arcmin.

In addition to mapping the extinction A , we would also like to infer the 3D distribution of the dust density. However, as

$$\frac{dA}{ds} = \rho\kappa, \quad (45)$$

where ρ is the dust density and κ its opacity, it is difficult to extract ρ directly. The dust opacity varies with changes in the dust grain size distribution and with dust composition, which means it is not completely straightforward to unpick its influence. So, we instead

we map the extinction pseudo-density $\rho\kappa$, which can be easily determined.

We adopt a mean function that assumes dust is distributed as an exponential disc to which we add a Gaussian ‘bump’ that roughly approximates a spiral arm. As a result, there are five hyperparameters. Three describe the disc: the dust scale length L , scale height H ,³ and the pseudo-density of extinction at the sun $(\rho\kappa)_0$. A further two describe the bump: its distance s_{bump} and total integrated extinction A_{bump} ; we assume the width of the bump to be fixed with a 400 pc standard deviation. We apply a lognormal hyperprior to each of these five parameters with means of 3000 pc for L , 150 pc for H , 0.75 kpc^{-1} for $(\rho\kappa)_0$, 2 kpc for s_{bump} and 0.1 for A_{bump} . The standard deviation of each hyperprior is set equal to its mean value, so that the hyperprior is largely uninformative.

5.1 Inducing point placement

In order to apply our method, we need to decide how to place the inducing points. It would be possible to include the locations (and number) of the inducing points as a set of parameters that are themselves inferred. But doing this is not only computationally expensive, but can also result in overfitting (Alvarez et al. 2009). Therefore, we opt to place them manually.

We place the inducing points along the sub-catalogue boundaries, following Vanhatalo et al. (2010), who found that this choice minimizes discontinuities. If the number of inducing points were to exceed the number of stars in each sub-catalogue, then calculating the covariance function for the inducing points would become the dominant computational cost. So we prefer to have a number of inducing points that is comfortably smaller than number of stars in each sub-catalogue. Finally, as there are essentially no stars in the catalogue at distances greater than 5 kpc, there is no value in placing inducing points at such distances. Therefore, we opt to place a total of 140 inducing points along the 35 catalogue boundaries at distances of 1, 2, 3, and 4 kpc.

5.2 Results

Fig. 1 shows the result obtained by running G-MEAD on this catalogue. The estimated uncertainties on the extinction map are very small, typically on the order of a few per cent. This is more than an order of magnitude smaller than comparable uncertainties in existing maps such as Sale et al. (2014) and Green et al. (2015). To an extent, this is a result of the simulated data not including a proper description of interstellar turbulence, although we still consider this to be an example of the precisions that may be achievable with *Gaia* data.

The map successfully captures the spiral arm that is present in the D03 model. Its relatively weak strength makes it difficult to pick out in Fig. 1, but when the contribution of the exponential disc is subtracted or when looking at individual sightlines, as in Figs 2 and 3, it is much more obvious.

For completeness, we must clarify exactly what Fig. 1 and similar are showing. For extinction we plot maps of the mean and variance of the *marginal* posterior predictive $p(a_* | x_*, \tilde{y}, \tilde{L}, \tilde{H}, \tilde{\zeta})$ on a regular grid of points. Note that we do not impose the condition that extinction must increase with distance. Estimating the map of density involves finding the joint posterior predictive at two locations for

each grid point: the grid point itself and another point in the same direction but 10 pc further away. Then, the density is simply found by sampling from the difference between the two extinctions and imposing the condition that the extinction to the further point is greater than that to the nearer.

There are no significant ‘Fingers of God’ in the plot of pseudo-density and extinction varies smoothly. However, low-level deviations between sightlines can be seen in the pseudo-density plot. But, we note that the strength of these variations are well below the estimated noise. These slight deviations appear because the density map within the area spanned by a given sub-catalogue is only directly conditioned on the observations of that sub-catalogue; the noise characteristics vary somewhat between sub-catalogues (e.g. as the number of stars and their distance distribution changes) resulting in changes into the width of the typically skewed PDF of pseudo-density and so in the posterior expectations of pseudo-density. These small differences are then enhanced slightly further by sampling noise.

We also make a comparison, in Fig. 4, to the D03 Galactic extinction model, on which the simulations are based, showing an excellent agreement. It is also apparent that, whilst the extinction uncertainties appear reasonable, the uncertainties on pseudo-density have been somewhat overestimated and are actually on the order of 10 per cent. This is an extraordinarily high level of precision, but we again caution that with real-data uncertainties will likely be somewhat higher. The uncertainties on pseudo-density have been overestimated due to sampling noise, which has a far greater effect than on extinction due to the relatively crude numerical means by which we estimate pseudo-densities.

In addition, we can estimate the posterior distributions of the hyperparameters that define the exponential dust disc. In Fig. 5, we plot this distribution for the dust scale length and normalization density. The posterior distribution of the scale length is in near exact agreement with the value used in the D03 model.

Roughly 2 h of time on an Amazon AWS t2.micro instance, with 1 (virtual) CPU, was required to produce the map shown in Fig. 1. Given that the region covered contains approximately 1 in every 10^4 stars in the GOG catalogue, we estimate that on the order of 10^4 CPU hours would be required to map the entire *Gaia* catalogue. This is well within the reach of reasonably sized CPU clusters. By way of contrast, we estimate that roughly 10^{11} CPU years would be required to map a similar sized catalogue using vanilla GP regression, which is therefore, obviously, utterly unachievable.

Our input catalogue contained cuts on the basis of the ‘true’ distances and masses of the stars, discarding nearly 90 per cent of sources. With real *Gaia* data, it will obviously not be possible to select the catalogue on this basis. However, it is clearly the case that some stars convey more information about the 3D distribution of extinction than others. Judicious use of e.g. photometric colour/magnitude cuts should enable us to preferentially select the most informative stars and so discard a large proportion of the *Gaia* catalogue.

5.3 Conditioning on $dA/ds \geq 0$

As discussed above, the maps in Fig. 1 do not enforce the requirement that (monochromatic) extinction must increase with distance.⁴

³We note that by choosing a region in the mid-plane, the dust scale height is almost completely unconstrained.

⁴Formally, this requirement is not absolute, being breached in reflection nebulae for example. However, this contribution typically only occurs over limited wavelength ranges and the pipeline processing of photometric data

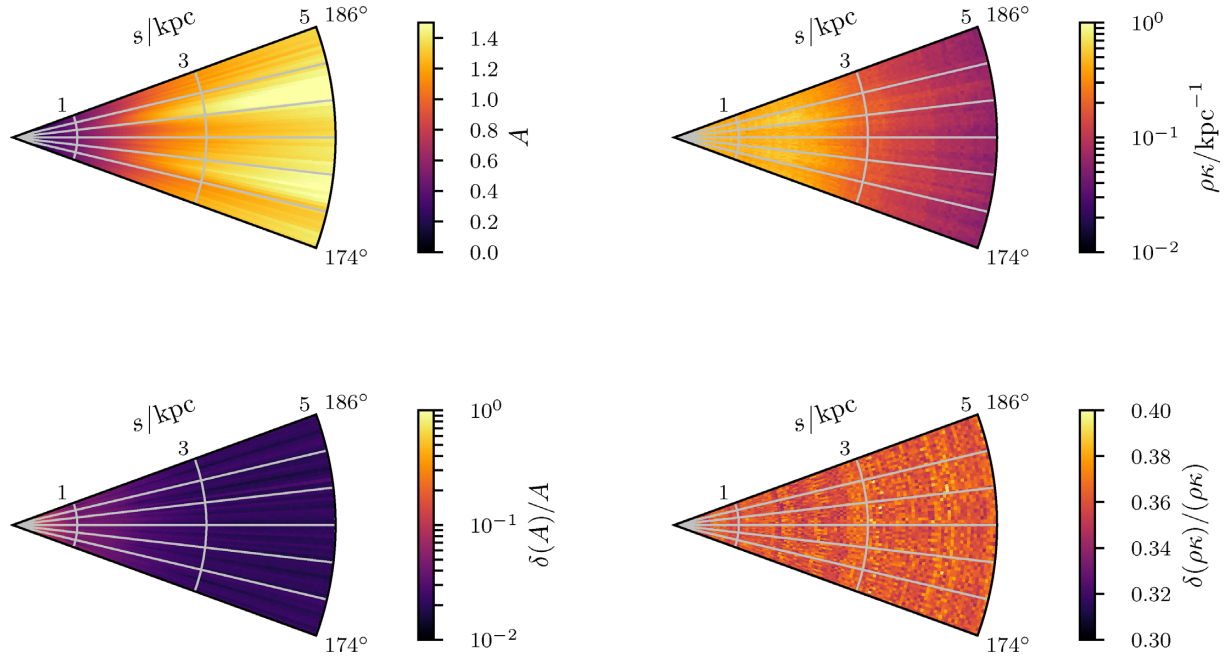


Figure 1. An extinction map of the Galactic anticentre based on GOG (Luri et al. 2014) simulated data. The top left panel shows the estimated extinction, the top right the estimated pseudo-density, whilst the lower two panels show relative uncertainties.




Figure 2. The map of the estimated extinction pseudo-densities with the inferred contribution of the exponential disc of dust having been subtracted. The remaining extinction pseudo-density is due to the spiral arm.




Figure 3. The estimated extinction pseudo-densities along the sightline towards the anticentre. The black crosses show the posterior estimates obtained, whilst the red line shows the pseudo-density from the D03 model.

Were we to condition our maps on $dA/ds \geq 0$, or equivalently $da/ds \geq 0$,⁵ we would be providing extra information that would increase the precision of the maps.

We now consider a line of points along a sightline \mathbf{x}_* . The posterior predictive distribution $p(\mathbf{a}_*|\mathbf{x}_*, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta)$ for the extinction to these points is multivariate normal, following the approximation described in Appendix C. The constraint that $da/ds \geq 0$ truncates this distribution, such that $p(\mathbf{a}_*|\mathbf{x}_*, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, da/ds \geq 0)$ follows a truncated multivariate normal distribution.

We use the method of Pakman & Paninski (2014) to sample from the truncated multivariate normal posterior predictive distribution. Pakman & Paninski (2014) employ a Hamiltonian Monte Carlo approach, with reflection at the truncation boundaries that produces samples in orders of magnitude less time than more straightforward approaches, such as rejection sampling or simple MCMC. Given a sample from the truncated multivariate normal posterior predictive distribution it is then trivial to calculate the moments.

Fig. 6 demonstrates the effect of conditioning on $dA/ds \geq 0$ for the example in the previous section. Each data point in this plot shows the mean and standard deviation of $p(\mathbf{a}_*|\mathbf{x}_*, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, da/ds \geq 0)$, the posterior distribution of extinction conditioned on $dA/ds \geq 0$. The effect of conditioning on $da/ds \geq 0$ is significant, typically providing a twofold improvement in the precision. However, this improvement is limited by the fact that the uncertainty in the unconditioned map is relatively low, due to the lack of small-scale turbulent structure in the simulated ISM. With real data we would expect a larger improvement in the precision.

will often include the subtraction of diffuse emission from extended sources. So, this requirement does hold in practice.

⁵Since we assume $A \geq 0$.




Figure 4. A comparison between the estimated extinction map and the D03 model that is employed in the production of the simulated data. Each panel shows a histogram of the residuals divided by the measured uncertainties. The left-hand panel compares extinctions and the right-hand panel shows pseudo-densities. On each panel, the standard normal distribution has been overplotted with a red line.




Figure 5. The estimated posterior distribution on the disc scale length L and the extinction pseudo-density at the sun. The red line shows the scale length employed in the D03 model.




Figure 6. Posterior estimates with uncertainties of the extinction along the sightline $(l, b) = (174.167^\circ, 0^\circ)$ plotted as black error bars. Note that the small uncertainties make it difficult to resolve some of the error bars. This is compared to the D03 model (red line) that is employed in the production of the simulated data.

6 CLOSING DISCUSSION

In Paper I, we presented a method for mapping extinction in three dimensions using GPs. This statistical description is an excellent

match to the physics of the ISM: it reproduces the observed density and column density distributions, allows the incorporation of a model of interstellar turbulence and neatly delineates the split between the small-scale physics of the ISM that is driven by turbulence and the larger scale physics. Unfortunately, exact GP regression scales poorly to large data sets, with the CPU time needed scaling as $\mathcal{O}(N^3)$. Consequently, processing a Gaia-sized catalogue with the unmodified algorithm of Paper I would have required an unfeasibly long time.

In Section 4, we describe a scheme by which we can dramatically accelerate the inference of GP-based extinction maps. Our approach is based on the combination of two approximations: the PIC approximation (Snelson & Ghahramani 2007) and EP (Minka 2001). The principal benefit of employing these two approximations is that they dramatically decrease the CPU time and memory needed to infer the extinction map, with both becoming an approximately linear function of the number of stars studied. In addition, the method significantly reduces the amount of network communication required if the construction of the map is distributed across multiple computing nodes and enables different portions of the map to be computed asynchronously. Other solutions for accelerating GP regression are available, but the method we have described is fast, trivially parallelizable, and conceptually and computationally straightforward.

We have implemented our scheme in a library G-MEAD. We have shown, using simulated *Gaia* data, that G-MEAD can produce 3D extinction maps with hitherto unmatched precision and accuracy. In addition, these simulations show that our method should be able to easily recognize real features such as spiral arms that have eluded previous maps. When applied to real *Gaia* data the resultant map will not only allow the direct study of the Galactic ISM, but will also support the wider study of the Galaxy, the stars that comprise it and the physical processes at work.

As we have described, it is now feasible to apply our method to a catalogue of a size similar to that which *Gaia* will produce. However, selection effects have a pathological impact on extinction mapping (Sale 2015) and we currently lack a method for dealing with them when mapping extinction with GPs. As a result, the improvements needed to overcome this hurdle will be the subject of a future paper. Once this has been overcome, the GP-based method we have described will be able to produce 3D extinction

maps from *Gaia* that will exhibit unparalleled detail, precision and accuracy.

ACKNOWLEDGEMENTS

We thank James Binney and Ralph Schönrich for comments on an earlier draft of this paper. The research leading to the results presented here was supported by the United Kingdom Science Technology and Facilities Council (STFC, ST/K00106X/1 and ST/M00127X/1), the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 321067.

REFERENCES

- Alvarez M. A., Luengo D., Titsias M. K., Lawrence N. D., 2009, preprint ([arXiv:0912.3268](https://arxiv.org/abs/0912.3268))
- Ambikasaran S., Darve E., 2013, *J. Sci. Comput.*, 57, 447
- Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2016, *IEEE Trans. Pattern Anal. Mach. Intell.*, 38, 252
- Bishop C. M., 2006, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, New York, Inc., Secaucus, NJ, USA
- Cervone D., Pillai N. S., 2015, preprint ([arXiv:1506.08256](https://arxiv.org/abs/1506.08256))
- Chen B. Q., Schultheis M., Jiang B. W., Gonzalez O. A., Robin A. C., Rejkuba M., Minniti D., 2013, *A&A*, 550, A42
- Cornford D., Csato L., Oppen M., 2005, *Geog. Anal.*, 37, 183
- Deisenroth M. P., Ng J. W., 2015, preprint ([arXiv:1502.02843](https://arxiv.org/abs/1502.02843))
- Drimmel R., Cabrera-Lavers A., López-Corredoira M., 2003, *A&A*, 409, 205
- Frey C. B., Osborne M. A., 2013, Technical Report, The Future of Employment: How Susceptible are Jobs to Computerisation?
- Gaia Collaboration et al., 2016, *A&A*, 595, A1
- Gelman A., Vehtari A., Jylänki P., Robert C., Chopin N., Cunningham J. P., 2014, preprint ([arXiv:1412.4869](https://arxiv.org/abs/1412.4869))
- Goldreich P., Sridhar S., 1995, *ApJ*, 438, 763
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Green G. M. et al., 2014, *ApJ*, 783, 114
- Green G. M. et al., 2015, *ApJ*, 810, 25
- Hensman J., Fusi N., Lawrence N. D., 2013, in Nicholson A., Smyth P., eds., *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*. AUAI Press Crovallis, Oregon, p. 282
- Hinton G. E., 2002, *Neural Comput.*, 14, 1771
- Kolmogorov A., 1941, *Akad. Nauk SSSR Dokl.*, 30, 301
- Lindgren F., Rue H., Lindström J., 2011, *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)*, 73, 423
- Luri X. et al., 2014, *A&A*, 566, A119
- Marshall D. J., Robin A. C., Reylé C., Schultheis M., Picaud S., 2006, *A&A*, 453, 635
- McHutchon A., Rasmussen C. E., 2011, in Shaw-Taylor J., Zemel R. S., Bartlett P. L., Pereira F., Weinberger K. Q., eds., *Advances in Neural Information Processing Systems*. Curran, New York, p. 1341
- Minka T. P., 2001, in *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufman, San Francisco, p. 362
- Murphy K. P., 2007, Technical Report, Conjugate Bayesian analysis of the Gaussian distribution. Department of Computer Science, University of British Columbia
- Nordlund Å. K., Padoan P., 1999, in Franco J., Carraminana A., eds., *Interstellar Turbulence*. Cambridge Univ. Press, Cambridge, p. 218
- Ostriker E. C., Stone J. M., Gammie C. F., 2001, *ApJ*, 546, 980
- Pakman A., Paninski L., 2014, *Journal of Computational and Graphical Statistics*, 23, 518
- Parts L., Stegle O., Winn J., Durbin R., 2011, *PLoS Genet.*, 7, e1001276
- Quiñero-Candela J., Rasmussen C. E., 2005, *The Journal of Machine Learning Research*, 6, 1939
- Rasmussen C. E., Williams C. K. I., 2005, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press
- Rezaei Kh. S., Bailer-Jones C. A. L., Hanson R. J., Fouesneau M., 2017, *A&A*, 598, A125
- Robin A. C. et al., 2012, *A&A*, 543, A100
- Rue H., Martino S., Chopin N., 2009, *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)*, 71, 319
- Sale S. E., 2015, *MNRAS*, 452, 2960
- Sale S. E., Magorrian J., 2014, *MNRAS*, 445, 256
- Sale S. E., Magorrian J., 2015, *MNRAS*, 448, 1738
- Sale S. E. et al., 2014, *MNRAS*, 443, 2907
- Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
- Seeger M., Williams C., Lawrence N., 2003, in Bishop C. M., Frey B. J., eds., *Ninth International Workshop on Artificial Intelligence and Statistics*, Society for Artificial Intelligence and Statistics, New Jersey
- Silverman B. W., 1985, *J. R. Stat. Soc. Ser. B (Methodol.)*, 47, 1
- Snelson E., Ghahramani Z., 2006, in *Advances in Neural Information Processing Systems*. MIT Press, Cambridge, MA, p. 1257
- Snelson E., Ghahramani Z., 2007, in Meila M., Shen X., eds., *Thirteenth International Workshop on Artificial Intelligence and Statistics*, Society for Artificial Intelligence and Statistics, New Jersey
- Tierney L., 1994, *Ann. Stat.*, 22, 1701
- Titsias M. K., 2009, in *International Conference on Artificial Intelligence and Statistics*. p. 567
- Tresp V., 2000, *Neural Comput.*, 12, 2719
- Vanhatalo J., Pietiläinen V., Vehtari A., 2010, *Stat. Med.*, 29, 1580
- Vergely J.-L., Freire Ferrero R., Siebert A., Valette B., 2001, *A&A*, 366, 1016
- Whittle P., 1954, *Biometrika*, 41, 434

APPENDIX A: ACCELERATING GP REGRESSION

As a consequence of their versatility, power and widespread adoption across a wide variety of fields there has been significant interest in extending the use of GP regression and inference to larger data sets. In this section, we will briefly review some methods that accelerate GP regression. We are not aware of any up-to-date and totally comprehensive review of such methods, however, Quiñero-Candela & Rasmussen (2005) and Rasmussen & Williams (2005) do examine a many of the methods discussed below in significantly more detail than we can afford them here.

The most trivial approach to accelerating GP regression is to approximate the full solution by using only a small subset of all the available data. If we employ only M of the original N observations, such that $M \ll N$, the CPU cost of the regression is reduced to $\mathcal{O}(M^3)$. However, it is trivially apparent that this is far from optimal, since, by its very nature, it requires discarding most available data. So, it cannot hope to be a competitive solution.

An alternative simple approximation is to divide the entire region studied into a number of small regions and then independently implement GP regression in each, using only the observations within each region. This approximation is sometimes referred to as local GPs. If each local region contains M observations, so that we again have that $M \ll N$, the CPU cost is reduced to just $\mathcal{O}(M^2N)$. There are two significant drawbacks to this approach, however. First, it will introduce discontinuities at the transitions between regions. In extinction mapping, these will manifest as Fingers of God. More importantly, if the size of the regions is small compared to the scale of the covariance kernel then there will be a very significant loss of precision, because when estimating the value of the field at any

given point one is only using a subset of the data that carries relevant information.

To overcome the problem of discontinuities inherent with local GPs, one could employ product of experts (PoE, Hinton 2002), a popular machine learning technique that seeks to simplify complicated inference problems. The concept behind PoE is that one approximates the true probability distribution by the product of a number of simple probability distributions. However, whilst a PoE approximation to GP regression can perform well in regions rich with observations, it produces estimates that are overconfident, particularly in regions where observations are sparse (Deisenroth & Ng 2015).

Bayesian committee machines (BCMs) were introduced by Tresp (2000) as a way of combining information from distinct data sets, with a particular focus on GP regression. They are conceptually similar to PoE, but avoid that approach's problem with overconfidence. However, they have a fundamental drawback in that they cannot deal with uncertainty on the hyperparameters: their entire derivation is conditioned on fixed hyperparameters. Therefore, given that we expect the a priori unknown hyperparameters that describe the large-scale structure of extinction to have a key role in the determination of 3D maps, BCMs cannot directly enable the production of 3D extinction maps within our scheme.

There are a number of methods for GP regression that employ a variational Bayesian approach (e.g. Titsias 2009; Hensman, Fusi & Lawrence 2013). However, these approaches generally assume that the positions of the observations are known exactly. As we do not know the distances of the stars we study, this is not the case with our data and so these methods are not so easily applicable here.

It has long been recognized (e.g. Whittle 1954) that if one has a stationary covariance function and a regular grid of observations then the covariance matrix can be diagonalized through a Fourier transform. The extension of this concept to irregularly grided data and non-stationary covariance functions is provided by the use of Stochastic Partial differential Equations (SPDEs, Lindgren, Rue & Lindström 2011). By combining the use of SPDEs with the Integrated Nested Laplace Approximation (Rue, Martino & Chopin 2009) Lindgren et al. (2011) were able to quickly perform GP regression with large data sets and non-stationary covariance functions. Two of the key stages in the application of SPDEs are the use of a coordinate transformation to shift to a space where the covariance function is stationary and the triangulation of the space such that all observations lie on the vertices of triangles (or tetrahedra in 3D). However, in our case both the covariance function (which depends on the mean function) and the locations of the observations are unknown. Consequently, not only would implementing an SPDE approach be rather complicated in our case, but also much of the CPU time saved by the use of SPDEs would be lost to repeatedly recalculating the coordinate transform and retriangularizing the space.

All the above approaches are approximations to full GP regression and so will inevitably involve some loss of precision. An alternative is discussed by Ambikasaran et al. (2016), who recognize that the covariance matrices of GPs are HODLR, enabling them to be factorized hierarchically following Ambikasaran & Darve (2013). Ambikasaran et al. (2016) show that one can then factorize the covariance matrix in $\mathcal{O}(N \log^2 N)$ time. Subsequently, finding the inverse or determinant of the covariance matrix is very fast and scales as $\mathcal{O}(N \log N)$. In our application, however, the uncertainties in the mean extinction function and the distances to individual stars mean that both the covariance function itself and the covariance matrix are unknown. This would mean that we could not avoid repeatedly

applying the relatively expensive factorization step. Consequently, although this approach is significantly faster than the naive application of GP regression, we found that it was not competitive with the approach we describe in Section 4.

A1 Sparse approximate GP regression

We start by noting that the joint, marginal, and conditional distributions of two sets of points that are drawn from the same GP are all (multivariate) Gaussian. We consider a GP with N observations \mathbf{y} at locations \mathbf{x} that are partitioned into K subsets, $\mathbf{y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_K\}$ and $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$. We further assume, for the sake of simplicity, that all of these subsets contain $M = N/K$ observations where $M \ll N$. We can then use the rules of conditional probability to decompose the GP probability

$$p(\mathbf{y}|\boldsymbol{\theta}) = \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{y}_1, \dots, \mathbf{y}_{k-1}, \boldsymbol{\theta}). \quad (\text{A1})$$

Thus, when conditioned on the first $k-1$ blocks, the k th block follows a multivariate Gaussian PDF with the mean vector and covariance matrix

$$\begin{aligned} \mathbf{m}_{k|1\dots k-1} &= \mathbf{m}_k + \boldsymbol{\Sigma}_{k,1\dots k-1} \boldsymbol{\Sigma}_{1\dots k-1,1\dots k-1}^{-1} (\mathbf{y}_{1\dots k-1} - \mathbf{m}_{1\dots k-1}) \\ \boldsymbol{\Sigma}_{k|1\dots k-1} &= \boldsymbol{\Sigma}_{k,k} - \boldsymbol{\Sigma}_{k,1\dots k-1} \boldsymbol{\Sigma}_{1\dots k-1,1\dots k-1}^{-1} \boldsymbol{\Sigma}_{1\dots k-1,k}. \end{aligned} \quad (\text{A2})$$

Where

$$\begin{pmatrix} \mathbf{y}_{1\dots k-1} \\ \mathbf{y}_k \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mathbf{m}_{1\dots k-1} \\ \mathbf{m}_k \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_{1\dots k-1,1\dots k-1} & \boldsymbol{\Sigma}_{1\dots k-1,k} \\ \boldsymbol{\Sigma}_{k,1\dots k-1} & \boldsymbol{\Sigma}_{k,k} \end{pmatrix} \right), \quad (\text{A3})$$

so that, for example, \mathbf{m}_k and $\boldsymbol{\Sigma}_{k,k}$ are the marginal mean and covariance for the k th block. It is important to note that factorizing the PDF in this manner does not reduce the CPU or memory costs involved in calculating or sampling from $p(\mathbf{y}|\boldsymbol{\theta})$, which retain their $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$ scaling, respectively.

A simple approximation to the factorization above would be to assume that off-diagonal blocks in the covariance matrix are populated exclusively with zeros, then

$$\begin{aligned} p(\mathbf{y}|\boldsymbol{\theta}) &\approx \prod_{k=1}^K p(\mathbf{y}_k | \boldsymbol{\theta}) \\ \mathbf{m}_{k|1\dots k-1} &\approx \mathbf{m}_k \\ \boldsymbol{\Sigma}_{k|1\dots k-1} &\approx \boldsymbol{\Sigma}_{k,k}. \end{aligned} \quad (\text{A4})$$

This approximation is essentially the local GP approximation: when conditioned on the hyperparameters, the values of the GP in different blocks are independent. This would result in significant and attractive computational savings: CPU time and memory size now scale as $\mathcal{O}(M^2 N)$ and $\mathcal{O}(MN)$. However, in practice, this is rarely a sensible approximation and will lead to large discontinuities between blocks and overfitting of hyperparameters.

But, what if we could find a computationally cheap way to approximate the impact of the off-diagonal blocks in the covariance matrix? There exists a family of approaches that seek to obtain a sparse approximation to the full covariance matrix using so-called 'inducing points'. We introduce the values \mathbf{u} of a GP at R inducing points \mathbf{x}_u and note that, by the definition of a GP, both $p(\mathbf{y}|\boldsymbol{\theta})$ and $p(\mathbf{y}, \mathbf{u}|\boldsymbol{\theta})$ follow multivariate Gaussian PDFs. Snelson & Ghahramani (2007) propose the PITC approximation, whereby one assumes that subsets of observations are independent given the value of the GP at the inducing point, i.e.

$$p(\mathbf{y}_k | \mathbf{y}_1, \dots, \mathbf{y}_{k-1}, \mathbf{u}, \boldsymbol{\theta}) \approx p(\mathbf{y}_k | \mathbf{u}, \boldsymbol{\theta}) \quad (\text{A5})$$

We can therefore approximate the factorized PDF above

$$\begin{aligned} p(\mathbf{y}|\boldsymbol{\theta}) &= \int d\mathbf{u} p(\mathbf{u}|\boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{y}_1, \dots, \mathbf{y}_{k-1}, \mathbf{u}, \boldsymbol{\theta}) \\ &\approx \int d\mathbf{u} p(\mathbf{u}|\boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{u}, \boldsymbol{\theta}). \end{aligned} \quad (\text{A6})$$

This approximate likelihood can then be employed to obtain the posterior required in the learning phase of GP regression (10).

$$\begin{aligned} p(\boldsymbol{\theta}|\mathbf{y}) &\approx \frac{p(\boldsymbol{\theta})}{p(\mathbf{y})} \int d\mathbf{u} p(\mathbf{u}|\boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{u}, \boldsymbol{\theta}) \\ p(\mathbf{u}, \boldsymbol{\theta}|\mathbf{y}) &\approx \frac{p(\mathbf{u}, \boldsymbol{\theta})}{p(\mathbf{y})} \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{u}, \boldsymbol{\theta}). \end{aligned} \quad (\text{A7})$$

When conditioned on \mathbf{u} , which has mean \mathbf{m}_u and covariance $\boldsymbol{\Sigma}_{u,u}$, and given the PITC approximation, the means and covariances of the k th block are

$$\begin{aligned} \mathbf{m}_{k|u} &\approx \mathbf{m}_k + \boldsymbol{\Sigma}_{k,u} \boldsymbol{\Sigma}_{u,u}^{-1} (\mathbf{u} - \mathbf{m}_u) \\ \boldsymbol{\Sigma}_{k|u} &\approx \boldsymbol{\Sigma}_{k,k} - \boldsymbol{\Sigma}_{k,u} \boldsymbol{\Sigma}_{u,u}^{-1} \boldsymbol{\Sigma}_{u,k}. \end{aligned} \quad (\text{A8})$$

From which it is clear how the inducing points provide an approximation to the true mean and covariance (A2).

Having employed the PITC approximation to accelerate the learning phase we now turn to the prediction phase. As discussed in Section 2, we can use observations of a GP at some points to constrain it at others. Specifically, the value of the GP at some unobserved position(s) is related to that at other observed positions,

$$p(\mathbf{y}_*|\mathbf{y}) = \int d\boldsymbol{\theta} p(\mathbf{y}_*|\mathbf{y}, \boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathbf{y}), \quad (\text{A9})$$

PITC provides us with a computationally cheap approximation to $p(\boldsymbol{\theta}|\mathbf{y})$. However, we still need to calculate $p(\mathbf{y}_*|\mathbf{y}, \boldsymbol{\theta})$. In vanilla GP regression, if the covariances of \mathbf{y} given $\boldsymbol{\theta}$ are stored during the learning phase, then the CPU time needed to find or sample from $p(\mathbf{y}_*|\mathbf{y}, \boldsymbol{\theta})$ scales as $\mathcal{O}(N^2)$. In most astrophysical applications, this cost will easily become restrictive.

Instead, Snelson & Ghahramani (2007) further propose the PIC approximation, that extends further upon PITC. To illustrate this, we initially consider estimating the value y_* of a GP at a single test point \mathbf{x}_* . If we break \mathbf{y} into blocks as in the PITC approximation, we obtain the factorization

$$p(\mathbf{y}_*, \mathbf{y}|\boldsymbol{\theta}) = p(\mathbf{y}_*|\mathbf{y}_1, \dots, \mathbf{y}_K, \boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{y}_1, \dots, \mathbf{y}_{k-1}, \boldsymbol{\theta}). \quad (\text{A10})$$

As with the exact factorization of $p(\mathbf{y}|\boldsymbol{\theta})$ (equation A1), this factorization provides no computational benefit. The PIC approximation assumes that the value y_* of the GP at the prediction point depends only on one block \mathbf{y}_i of \mathbf{y} , when conditioned \mathbf{u} .

$$p(\mathbf{y}_*|\mathbf{y}_1, \dots, \mathbf{y}_K, \mathbf{u}, \boldsymbol{\theta}) \approx p(\mathbf{y}_*|\mathbf{y}_i, \mathbf{u}, \boldsymbol{\theta}). \quad (\text{A11})$$

Typically, the block chosen will cover the region of space in which \mathbf{x}_* falls. Combining this with the existing approximation (A5) obtains

$$p(\mathbf{y}_*, \mathbf{y}|\mathbf{u}, \boldsymbol{\theta}) \approx p(\mathbf{y}_*|\mathbf{y}_i, \mathbf{u}, \boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\mathbf{u}, \boldsymbol{\theta}), \quad (\text{A12})$$

and so

$$p(\mathbf{y}_*|\mathbf{y}, \mathbf{u}, \boldsymbol{\theta}) \approx p(\mathbf{y}_*|\mathbf{y}_i, \mathbf{u}, \boldsymbol{\theta}). \quad (\text{A13})$$

Calculating or sampling from this is cheap, scaling as $\mathcal{O}(M^2)$ since it is only conditioned on a single block \mathbf{y}_i , in addition to the inducing

points and hyperparameters, rather than all K blocks. Therefore,

$$p(\mathbf{y}_*|\mathbf{y}) \approx \iint d\mathbf{u} d\boldsymbol{\theta} p(\mathbf{y}_*|\mathbf{y}_i, \mathbf{u}, \boldsymbol{\theta}) p(\mathbf{u}, \boldsymbol{\theta}|\mathbf{y}), \quad (\text{A14})$$

where we can employ the PITC approximation (A7) to $p(\mathbf{u}, \boldsymbol{\theta}|\mathbf{y})$.

In general, we wish to predict the values \mathbf{y}_* of the GP at many prediction points \mathbf{x}_* . We can break \mathbf{y}_* into blocks analogously to \mathbf{y} , typically based on the positions \mathbf{x}_* . Then,

$$p(\mathbf{y}_*|\mathbf{y}, \mathbf{u}, \boldsymbol{\theta}) \approx \prod_{k=1}^K p(\mathbf{y}_{*,k}|\mathbf{y}_k, \mathbf{u}, \boldsymbol{\theta}), \quad (\text{A15})$$

therefore,

$$p(\mathbf{y}_*|\mathbf{y}) \approx \iint d\mathbf{u} d\boldsymbol{\theta} p(\mathbf{u}, \boldsymbol{\theta}|\mathbf{y}) \prod_{k=1}^K p(\mathbf{y}_{*,k}|\mathbf{y}_k, \mathbf{u}, \boldsymbol{\theta}). \quad (\text{A16})$$

The CPU time and memory requirements of the training phase of vanilla GP regression are dominated by the need to invert and store an $N \times N$ covariance matrix, resulting in them scaling as $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$, respectively. However, making the PI(T)C approximation of Snelson & Ghahramani (2007) alters our needs so that we instead must invert and store $K M \times M$ covariance matrices. As a result, CPU time and memory requirements instead scale as $\mathcal{O}(NM^2)$ and $\mathcal{O}(NM)$, respectively. Given that $M \ll N$, this potentially results in massive savings. In addition, the PIC approximation reduces the cost of the prediction phase, from $\mathcal{O}(N^2)$ under vanilla GP regression to $\mathcal{O}(NM)$. Therefore, PIC potentially makes GP regression viable in a much wider range of problems than vanilla GP regression can approach.

However, whilst PIC can produce excellent results, it cannot be straightforwardly parallelized without significant interthread communication. Consequently, the size of a region that can be easily studied in this way is in practice often limited to that that can be contained on a single computing node.

PI(T)C is just one of a number of inducing point methods, as reviewed by Quiñero-Candela & Rasmussen (2005). The other methods in this group include Subset of Regressors (Silverman 1985), Projected Latent Variables (Seeger, Williams & Lawrence 2003), and Fully Independent Conditional (Snelson & Ghahramani 2006). However, as Snelson & Ghahramani (2007) demonstrate, PI(T)C is the most sophisticated and generally most accurate of them.

A2 Expectation propagation

EP seeks to make inferences based on large data sets, breaking the data down into smaller sets, which are analysed separately, before combining the results together. EP was originally derived by Minka (2001), whilst Bishop (2006) and Gelman et al. (2014) provide accessible introductions.

We are not aware of any prior use of EP in the astrophysics literature, though Gelman et al. (2014) do briefly examine a highly simplified astronomy-inspired example. However, it has been widely applied in a variety of other fields, including genetics (Parts et al. 2011), geostatistics (Cornford, Csato & Oppen 2005) and in examining future employment prospects (Frey & Osborne 2013).

The goal of EP is to approximate the posterior

$$p(\boldsymbol{\theta}|\mathbf{y}) = \frac{1}{p(\mathbf{y})} p(\boldsymbol{\theta}) \prod_{k=1}^K p(\mathbf{y}_k|\boldsymbol{\theta}), \quad (\text{A17})$$

by a factorized distribution of the form

$$q(\theta) = p(\theta) \prod_{n=1}^K \frac{1}{Z_k} q_k(\theta), \quad (\text{A18})$$

in which the (normalized) fitting functions q_k , also known as site distributions, have some simple parametrized functional form. The idea then is to adjust the parameters describing the $\{q_k\}$ to optimize the agreement with the posterior equation (A17).

The EP algorithm uses the Kullback–Leibler (KL) divergence,

$$\text{KL}(f||g) \equiv - \int f(\theta) \log \left(\frac{g(\theta)}{f(\theta)} \right) d\theta, \quad (\text{A19})$$

to quantify the differences between various functions $f(\theta)$ and $g(\theta)$ related to the posterior equation (A17) and the factorized approximation (A18). To begin, consider

$$\text{KL}(p||q) = \text{KL} \left(\frac{1}{p(\mathbf{y})} p(\theta) \prod_k p(\mathbf{y}_k|\theta) \left\| \frac{1}{Z} p(\theta) \prod_k q_k(\theta) \right. \right), \quad (\text{A20})$$

which would be very difficult to minimize directly. The key idea of EP is to focus on just one factor q_l at a time, replacing all factors $p(\mathbf{y}_k|\theta)$ for $k \neq l$ in the first argument by the corresponding $q_k(\theta)/Z_k$, namely

$$\text{KL} \left(\frac{1}{p(\mathbf{y})} p(\theta) p(\mathbf{y}_l|\theta) \prod_{k \neq l} q_k(\theta) \left\| \frac{1}{Z_l} p(\theta) q_l(\theta) \prod_{k \neq l} q_k(\theta) \right. \right). \quad (\text{A21})$$

So, each q_l is fit just to the corresponding posterior factor $p(\mathbf{y}_l|\theta)$, but weighted to include the other $q_{k \neq l}$ and the prior $p(\theta)$. This is nicely visualized by fig. 2 of Gelman et al. (2014).

Introducing the *cavity* distribution

$$q_{\setminus l}(\theta) \equiv \frac{1}{q_l(\theta)} q(\theta) = \frac{1}{Z_l} p(\theta) \prod_{k \neq l} q_k(\theta) \quad (\text{A22})$$

and the *tilted* distribution

$$q_{\setminus l}(\theta) \equiv p(\mathbf{y}_l|\theta) q_{\setminus l}(\theta), \quad (\text{A23})$$

enables the KL divergence equation (A21) to be rewritten more compactly as

$$\text{KL} \left(\frac{1}{C_l} q_{\setminus l}(\theta) \left\| q_{\setminus l}(\theta) q_l(\theta) \right. \right), \quad (\text{A24})$$

with $C_l \equiv \int q_{\setminus l}(\theta) d\theta$. If the form assumed for the $q_k(\theta)$ (and also for the prior $p(\theta)$) belongs to the exponential family of distributions then the second argument is itself another member of the exponential family and the KL divergence is easy to minimize (e.g. Bishop 2006): just take $Z_l = C_l$ and adjust the parameters of $q_l(\theta)$ so that the sufficient statistics of $\frac{1}{Z_l} q_{\setminus l}(\theta)$ and $q_{\setminus l}(\theta) q_l(\theta)$ match. In particular, if we adopt Gaussian functional forms for the q_k , then the best choice of mean and covariance for the factor q_l is given by matching the first- and second-order moments of $q_l(\theta) q_{\setminus l}(\theta)$ to those of $\frac{1}{Z_l} q_{\setminus l}(\theta)$.

The key detail to appreciate from the above is the central roles the cavity and tilted distributions play. In particular that we can define the problem of finding the best approximation $q(\theta)$ to involve only the cavity, tilted, and site distributions. The EP algorithm then simply refines the site approximations $q_k(\theta)$, so that the divergence equation (A24) is minimized and the best approximation to the underlying distribution is found.

EP can be carried out with these updates to the q_k performed in series or in parallel. The parallel procedure can be summarized as follows (paraphrasing Gelman et al. 2014):

(1) *Initialize*: Partition the data into K groups, so that $\mathbf{y} = \{\mathbf{y}_1, \dots, \mathbf{y}_K\}$. Then approximate each $p(\mathbf{y}_k|\theta)$ with $q_k(\theta)$ chosen from a simple distribution (e.g. Gaussian). Use equation (A18) to construct a first guess to $q(\theta)$.

(2) *Iterate*: Perform the following until convergence is achieved:

(a) For $k = 1, \dots, K$:

(i) Construct the cavity and tilted distributions, $q_{\setminus k}(\theta) = q(\theta)/q_k(\theta)$ and $q_{\setminus k}(\theta) = p(\mathbf{y}_k|\theta) q_{\setminus k}(\theta)$. Generally, the latter will not be available analytically, but one can sample from it using, e.g. MCMC.

(ii) Find a new approximation $q'_k(\theta)/Z'_k$ that minimizes equation (A24) by, e.g. moment matching.

(b) Construct the updated approximation $q'(\theta) = p(\theta) \prod_{k=1}^K \frac{1}{Z'_k} q'_k(\theta)$ and adopt this as the new approximation $q(\theta)$.

(c) Check for convergence. If reached then terminate, so that the final approximation is given by $q(\theta)$

The extension to hierarchical models, such as the one we employ, is straightforward and described by Gelman et al. (2014). An advantage of using EP to study hierarchical models is that it allows ‘local’ parameters to be partitioned, such that they only enter one of the likelihood terms. The only practical difference compared to the non-hierarchical algorithm outlined above occurs when calculating the tilted distribution. If there are additional ‘local’ parameters α_k that sit between θ and the data, then the *local unmarginalized* tilted distribution $q_{\setminus k}(\alpha_k, \theta) = p(\mathbf{y}_k|\alpha_k, \theta) p(\alpha_k|\theta) q_{\setminus k}(\theta)$ is found for each $k = 1, \dots, K$. As before, one typically samples from this distribution numerically. Then, the marginal tilted distribution $q_{\setminus k}(\theta)$ is found by marginalizing α_k . In some cases, this marginalization can be performed analytically. However, more typically the marginalization is performed numerically: if the local unmarginalized tilted distribution is found by MCMC then the marginalization is trivial.

In some respects, it is possible to think of EP as a generalization of BCMs. Specifically, BCMs are equivalent to a single pass of a parallel EP algorithm, conditioned on fixed hyperparameters. However, in contrast to BCMs, EP is suitable for use with hierarchical models with non-fixed hyperparameters and so is a good fit for our purposes.

APPENDIX B: DETAILS OF THE MOMENT MATCHING IN EP

Since \mathbf{a}_k and \mathbf{u} are drawn from a GP, it follows that $p(\mathbf{a}_k, \mathbf{u}|\Theta)$ is Gaussian. We express this Gaussian as

$$\mathbf{a}_k, \mathbf{u}|\Theta \sim \mathcal{N} \left(\begin{pmatrix} \mu_k \\ m \end{pmatrix}, \begin{pmatrix} \Sigma_k & \rho_k \\ \rho_k^T & C \end{pmatrix} \right). \quad (\text{B1})$$

Similarly,

$$\tilde{\mathbf{a}}_i, \mathbf{u}|\theta \sim \mathcal{N} \left(\begin{pmatrix} \mu_i \\ m \end{pmatrix}, \begin{pmatrix} \Sigma_i + \sigma_i & \rho_i \\ \rho_i^T & C \end{pmatrix} \right). \quad (\text{B2})$$

The reduced cavity distribution is a Gaussian, which we may write as

$$q'_{\setminus k}(\mathbf{u}, \Theta) \sim \mathcal{N} \left(\begin{pmatrix} m_k \\ \psi_k \end{pmatrix}, \begin{pmatrix} C_k & \rho_k \\ \rho_k^T & \Omega_k \end{pmatrix} \right). \quad (\text{B3})$$

With these expressions $q'_k(\mathbf{u}|\Theta)$ and $p(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ are given by

$$\begin{aligned} \mathbf{u}|\tilde{\mathbf{a}}_k, \Theta &\sim \mathcal{N}(\mathbf{m}'_k, \mathbf{C}'_k) \\ q'_k(\mathbf{u}|\Theta) &\sim \mathcal{N}(\mathbf{m}'_{j,k}, \mathbf{C}'_{j,k}), \end{aligned} \quad (\text{B4})$$

where

$$\begin{aligned} \mathbf{m}'_k &= \mathbf{m} + \rho_k (\Sigma_k + \sigma_k)^{-1} (\tilde{\mathbf{a}}_k - \mu_k), \\ \mathbf{C}'_k &= \mathbf{C} + \rho_k (\Sigma_k + \sigma_k)^{-1} \rho_k^T, \\ \mathbf{m}'_{j,k} &= \mathbf{m}_{j,k} + \rho_{j,k} \Omega_{j,k}^{-1} (\Theta - \psi_{j,k}), \\ \mathbf{C}'_{j,k} &= \mathbf{C}_{j,k} + \rho_{j,k} \Omega_{j,k}^{-1} \rho_{j,k}^T. \end{aligned} \quad (\text{B5})$$

The Gaussian $\bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta)$ and its normalizing factor F_k are given by

$$\begin{aligned} \bar{q}_k(\mathbf{u}|\tilde{\mathbf{a}}_k, \Theta) &\sim \mathcal{N}(\bar{\mathbf{m}}_k, \bar{\mathbf{C}}_k) \\ \bar{\mathbf{C}}_k^{-1} &= \mathbf{C}_k'^{-1} + \mathbf{C}_k'^{-1}, \\ \bar{\mathbf{C}}_k^{-1} \bar{\mathbf{m}}_k &= \mathbf{C}_k'^{-1} \mathbf{m}'_k + \mathbf{C}_k'^{-1} \mathbf{m}'_{j,k}, \\ F_k(\tilde{\mathbf{a}}_k, \Theta) &= \frac{|\bar{\mathbf{C}}_k|^{\frac{1}{2}}}{|\mathbf{C}_k'|^{\frac{1}{2}} |\mathbf{C}_k'|^{\frac{1}{2}}} \\ &\times \exp \left[-\frac{1}{2} \left(\mathbf{m}_k'^T \mathbf{C}_k'^{-1} \mathbf{m}'_k + \mathbf{m}_{j,k}'^T \mathbf{C}_k'^{-1} \mathbf{m}'_{j,k} - \bar{\mathbf{m}}_k^T \bar{\mathbf{C}}_k^{-1} \bar{\mathbf{m}}_k \right) \right]. \end{aligned} \quad (\text{B6})$$

APPENDIX C: DERIVATION OF THE APPROXIMATE PREDICTION DISTRIBUTIONS

We start by recalling that the 3D map of extinction formally follows the posterior predictive distribution (21). Following some manipulation, not repeated here for the sake of brevity, we can rearrange this to a form similar that of the simpler example that was discussed in Section 4.1,

$$\begin{aligned} p(\mathbf{a}_*|\tilde{\mathbf{y}}, \mathbf{x}_*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) &= \int d\Theta \sum_z \int d\mathbf{s} \left(p(\mathbf{s}, \mathbf{z}|\Theta, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) p(\Theta|\tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \right. \\ &\quad \left. \times \int d\mathbf{a} p(\mathbf{a}, \mathbf{a}_*|\mathbf{s}, \mathbf{z}, \Theta, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \right). \end{aligned} \quad (\text{C1})$$

Now we introduce \mathbf{u} into the integrand and marginalize,

$$\begin{aligned} p(\mathbf{a}_*|\tilde{\mathbf{y}}, \mathbf{x}_*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) &= \iint d\mathbf{u} d\Theta \sum_z \int d\mathbf{s} \left(p(\mathbf{s}, \mathbf{z}|\mathbf{u}, \Theta, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) p(\mathbf{u}, \Theta|\tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \right. \\ &\quad \left. \times \int d\mathbf{a} p(\mathbf{a}, \mathbf{a}_*|\mathbf{s}, \mathbf{z}, \mathbf{u}, \Theta, \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \right). \end{aligned} \quad (\text{C2})$$

By the use of the PIC approximation, we can obtain an approximate factorization of the integrand,

$$\begin{aligned} p(\mathbf{a}_*|\tilde{\mathbf{y}}, \mathbf{x}_*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) &\approx \iint d\mathbf{u} d\Theta \left(p(\mathbf{u}, \Theta|\tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) \right. \\ &\quad \left. \times \prod_{k=1}^K \left[\sum_{z_k} \int d\mathbf{s}_k \left(p(\mathbf{s}_k, \mathbf{z}_k|\mathbf{u}, \Theta, \tilde{\mathbf{y}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right. \right. \right. \\ &\quad \left. \left. \left. \int d\mathbf{a}_k p(\mathbf{a}_k, \mathbf{a}_{*,k}|\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{y}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right) \right] \right), \end{aligned} \quad (\text{C3})$$

where we have split \mathbf{a}_* into blocks corresponding to the sub-catalogues. We can then use the EP approximation $q(\mathbf{u}, \Theta)$ to the posterior $p(\mathbf{u}, \Theta|\tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta)$, as found in the learning phase. In addition, as we employ the GMM approximation, we can replace $\tilde{\mathbf{y}}$ by its sufficient statistics $\{\tilde{\mathbf{a}}, \tilde{\mathbf{s}}\}$ and perform the marginalization

$$\begin{aligned} \int d\mathbf{a}_k p(\mathbf{a}_k, \mathbf{a}_{*,k}|\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) &= p(\mathbf{a}_{*,k}|\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \end{aligned} \quad (\text{C4})$$

analytically. We then obtain

$$\begin{aligned} p(\mathbf{a}_*|\tilde{\mathbf{y}}, \mathbf{x}_*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) &\approx \iint d\mathbf{u} d\Theta \left(q(\mathbf{u}, \Theta) \right. \\ &\quad \left. \times \prod_{k=1}^K \left[\sum_{z_k} \int d\mathbf{s}_k \left(p(\mathbf{s}_k, \mathbf{z}_k|\mathbf{u}, \Theta, \tilde{\mathbf{s}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right. \right. \right. \\ &\quad \left. \left. \left. p(\mathbf{a}_{*,k}|\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right) \right] \right). \end{aligned} \quad (\text{C5})$$

This appears to be a formidable equation to deal with. However, we note that if the EP algorithm has converged then the EP approximation and tilted distributions should be approximately equal, thus

$$\begin{aligned} q(\mathbf{u}, \Theta) &\approx q_{\setminus k}(\mathbf{u}, \Theta) \\ q(\mathbf{u}, \Theta) p(\mathbf{s}_k, \mathbf{z}_k|\mathbf{u}, \Theta, \tilde{\mathbf{s}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) &\approx q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta), \end{aligned} \quad (\text{C6})$$

where $q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta)$ is the local unmarginalized tilted distribution. Therefore,

$$\begin{aligned} p(\mathbf{a}_{*,k}|\tilde{\mathbf{y}}, \mathbf{x}_*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \zeta, \beta) &\approx \iint d\mathbf{u} d\Theta \sum_{z_k} \int d\mathbf{s}_k \left(q_{\setminus k}(\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta) \right. \\ &\quad \left. p(\mathbf{a}_{*,k}|\mathbf{s}_k, \mathbf{z}_k, \mathbf{u}, \Theta, \tilde{\mathbf{a}}_k, \tilde{\mathbf{l}}_k, \tilde{\mathbf{b}}_k, \zeta, \beta) \right). \end{aligned} \quad (\text{C7})$$

This paper has been typeset from a \LaTeX file prepared by the author.