Three-dimensional dust density structure of the Orion, Cygnus X, Taurus, and Perseus star-forming regions^{★,★★}

T. E. Dharmawardena¹ , C. A. L. Bailer-Jones¹, M. Fouesneau¹ , and D. Foreman-Mackey²

¹ Max Planck Institute for Astronomy (MPIA), Königstuhl 17, 69117 Heidelberg, Germany
e-mail: dharmawardena@mpia.de

² Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York, NY 10010, USA

Received 12 May 2021 / Accepted 10 November 2021

ABSTRACT

Interstellar dust affects many astronomical observations through absorption and reddening, yet this extinction is also a powerful tool for studying interstellar matter in galaxies. Three-dimensional (3D) reconstructions of dust extinction and density in the Milky Way have suffered from artefacts such as the fingers-of-god effect and negative densities, and have been limited by large computational costs. Here, we aim to overcome these issues with a novel algorithm that derives the 3D extinction density of dust in the Milky Way using a latent variable Gaussian process in combination with variational inference. Our model maintains non-negative density and hence monotonically non-decreasing extinction along all lines-of-sight, while performing the inference within a reasonable computational time. Using extinctions for hundreds of thousands of stars computed from optical and near-infrared photometry, together with distances based on *Gaia* parallaxes, we employ our algorithm to infer the structure of the Orion, Taurus, Perseus, and Cygnus X star-forming regions. A number of features that are superimposed in 2D extinction maps are clearly debled in 3D dust extinction density maps. For example, we find a large filament on the edge of Orion that may host a number of star clusters. We also identify a coherent structure that may link the Taurus and Perseus regions, and we show that Cygnus X is located at 1300–1500 pc, in line with very-long-baseline interferometry measurements. We compute dust masses of the regions and find these to be slightly higher than previous estimates, likely a consequence of our input data recovering the highest column densities more effectively. By comparing our predicted extinctions to *Planck* data, we find that known relationships between density and dust processing, where high-extinction lines-of-sight have the most processed grains, hold up in resolved observations when density is included, and that they exist at smaller scales than previously suggested. This can be used to study the changes in size or composition of dust as they are processed in molecular clouds.

Key words. methods: numerical – dust, extinction – ISM: clouds – ISM: structure – local interstellar matter – Galaxy: structure

1. Introduction

Interstellar dust affects how we view the Universe by absorbing and scattering starlight, leading to interstellar reddening and extinction. Although interstellar dust encompasses only a small fraction ($\sim 1\%$) of the baryonic matter in our Universe, understanding the properties of dust is crucial for interpreting observations correctly. The absorption of starlight also plays a key role in the thermal balance of the interstellar medium (ISM), driving chemical reactions and enabling gas clouds to collapse. Knowing the spatial distribution of interstellar dust is therefore key both to understanding the ISM as whole and to being able to account for the effects of interstellar extinction.

Past studies have made great strides in mapping the spatial distribution of the ISM in both extinction and dust density. One of the most comprehensive works was by Schlegel et al. (1998), who used IRAS and COBE far-infrared emission to produce an all-sky 2D dust extinction map. This was followed by two decades of improvements fuelled by advances in machine

learning techniques and the application of Bayesian statistics, leading to both new 2D and 3D dust extinction and extinction density maps (e.g. Marshall et al. 2006; Schlafly & Finkbeiner 2011; Sale et al. 2014; Sale & Magorrian 2014, 2018; Lallement et al. 2014, 2019; Green et al. 2015, 2019; Hanson et al. 2016; Rezaei Kh. et al. 2017; Babusiaux et al. 2020; Leike et al. 2020).

Gaussian processes (GPs) have been widely used to study the 3D structure of interstellar dust. Sale & Magorrian (2014) used them to predict extinction in simulated data on individual lines-of-sight and then applied it to IPHAS photometry in a slice of the Galactic plane (Sale et al. 2014). This was followed by Sale & Magorrian (2018), who improved the method to account for correlations between lines-of-sight and applied it to simulated data along individual lines-of-sight.

Green et al. (2019) took a similar approach and produced a 3D extinction map covering the entire sky north of -30° up to 2 kpc. The authors used 2MASS, *Gaia*, and Pan-STARRS data to carry out important sampling on a parameter grid assuming a GP prior on the logarithm of the dust reddening density. These maps have become a go-to resource for estimating interstellar extinction along lines-of-sight as they are easily accessible. The works by Rezaei Kh. et al. (2018a,b) used GP techniques to predict the 3D dust density of the Orion region and the Milky Way disk using *Gaia* DR2, 2MASS, WISE, and APOGEE DR14 data for

* The 3D maps are only available at the CDS via anonymous ftp to [cdsarc.u-strasbg.fr](ftp://cdsarc.u-strasbg.fr) (130.79.128.5) or via <http://cdsarc.u-strasbg.fr/viz-bin/cat/J/A+A/658/A166>

** Our code is available at <https://github.com/Thavisha/Distribution> and the latest results are available from www.mwdust.com

30 000 stars. Leike & Enßlin (2019) and Leike et al. (2020) modelled the dust density as a GP and simultaneously inferred the power spectrum of the dust density. Using this method, they modelled the inner 400 pc molecular clouds coverage using 30 million sources from multiple surveys, including *Gaia*.

Modelling the 3D dust distribution of the Milky Way with GPs is not without its challenges and issues, however. One is the so-called fingers-of-god effect, in which the density distribution is elongated along the line-of-sight due to tangential accuracy being higher than radial accuracy. It can be avoided if high-accuracy distances are available and/or if correlations between points in 3D space are incorporated explicitly rather than as individual lines-of-sight coupled together in the plane of the sky. Another key feature that must be considered is the physical requirement that densities be positive, and hence extinction must be monotonically non-decreasing along any line-of-sight. One way of achieving this is to model the logarithm of the density instead of density or extinction itself. Gaussian processes are computationally intensive, naively scaling with $O(N^3)$ in number of sources. Methods with improved scaling are therefore important for ensuring that a wide range of problems are feasible with minimal tradeoffs between, for example, resolution, map size, and number of sources. Finally, the reproducibility of the results may be hampered by the lack of code availability or non-use of public libraries.

In this paper we present a technique that directly accounts for the 3D correlations between density points and enforces positive density throughout. To accelerate the processing, we use GP latent variable models in combination with variational inference to infer the logarithm of the dust density. We describe our method in Sect. 2 and validate it using simulated data in Sect. 3. Then, using input data from Sect. 4, we apply it to the four well-known star-forming regions (SFRs), Orion, Cygnus X, Taurus, and Perseus, in Sects. 5 and 6 to map their dust in 3D. From this we can compute their total masses (Sect. 7), and in Sect. 8 we compare our results to *Planck* dust emission measures.

2. 3D extinction and density mapping model

2.1. Overview

To ensure that density is always positive, we model the logarithm of the density, and impose a GP prior to account for correlations between the densities at points in three dimensions. While we use the term ‘dust density’ or simply ‘density’ for ease throughout this work, in reality we are actually mapping ‘dust extinction density’ in mag pc^{-1} units. For it to be a traditional density in g cm^{-3} units the dust extinction density needs to be converted using the dust optical depth and dust opacity.

We optimise the hyperparameters for this GP along with several other parameters that approximately condition the GP as described in Sects. 2.3 and 2.4, to infer the logarithm of the density. At the end we draw sample maps from the final conditioned GP to explore the dust density distribution and its uncertainty; these are integrated to predict extinction and its uncertainty.

Our approach requires measurements and uncertainties of the extinctions towards a sample of stars, as well as their positions in three dimensions. Uncertainty on distance is included following the same approach as Rezaei Kh. et al. (2020), which involves incorporating them into the observed extinctions by inflating their uncertainties (see their Sect. 2 and Eqs. (1) and (2)). We utilise GP package *GPyTorch* (Gardner et al. 2018) and probabilistic programming package *Pyro* (Bingham et al. 2019; Phan

et al. 2019) in our implementation. Both packages are built upon *PyTorch* (Paszke et al. 2019), an open-source machine learning framework.

2.2. Latent Gaussian process function

Mathematically, our quantity of interest, dust density (ρ), can be described as a latent quantity or function – a quantity that we do not directly observe. This latent quantity can be transformed by integrating along lines-of-sight to forward model the observed quantity, which is extinction. Dust density is expected to vary in a complex way spatially, which cannot be accurately captured by a parametric function. As a result, the model for the latent function should ideally be flexible and non-parametric, which makes GPs ideal for this purpose.

We modelled the measured extinctions with a Gaussian likelihood. Given that $\mathbf{A}_{\text{mod}} = \int_0^d \rho ds_i$, if we adopted a GP prior on ρ , then the posterior would also be Gaussian in ρ , with analytic solutions for its mean and covariance. This is the approach of Rezaei Kh. et al. (2017) (their Appendix A). However, to enforce non-negative densities, we defined our GP prior on $\log_{10}(\rho)$, with the consequence that the exact posterior no longer has a simple analytic form (see Appendix A). We handled this by using variational inference (see Sect. 2.3) and approximated the posterior in $\log_{10}(\rho)$ as normally distributed. The prior on $\log_{10}(\rho)$ – which we now denote ϕ for brevity – is modelled using a GP with a constant mean (see Chap. 2.2 and Eq. (2.14) in Rasmussen & Williams 2006),

$$P(\phi(x, y, z)) = \mathcal{GP}(c, \alpha K(\mathbf{x}_1 - \mathbf{x}_2)), \quad (1)$$

and elements of the covariance matrix given by a 3D radial basis function (RBF) kernel

$$K(\mathbf{x}_1, \mathbf{x}_2) = \exp\left(-\frac{1}{2}(\mathbf{x}_1 - \mathbf{x}_2)^T \boldsymbol{\gamma}^{-2}(\mathbf{x}_1 - \mathbf{x}_2)\right), \quad (2)$$

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are two position vectors in 3D space, $\boldsymbol{\gamma}$ is a diagonal matrix of scale lengths, α is a scale factor, and c is a constant mean (see Chap. 4.2 Eq. (4.9) and Chap. 2.2 Eq. (2.14) in Rasmussen & Williams 2006)^{1, 2}. We abbreviate the right hand side of Eq. (1) to $\mathcal{GP}(\Theta)$ throughout this paper, where Θ represents the hyperparameters of the GP.

Our GP has five hyperparameters: three physical scale lengths in the three Heliocentric Cartesian coordinates (x, y, z) of physical space; one exponential scale factor; the mean density. These hyperparameters are used to generate the GP prior from which sets of ϕ are predicted. In effect, treating ϕ as a latent variable introduces one free parameter at each point in x, y, z , all coupled together through the GP prior. We evaluate the GP on a grid, which we take to be regular in Galactic (spherical) coordinates (l, b, d) as this eases the implementation of line-of-sight integration. We then convert the coordinates of the centres of the cells to Heliocentric Cartesian coordinates for use in Eq. (1).

To compare to observed extinction (A_{obs}) to what our model density distribution predicts, we must integrate this density distribution along the line-of-sight. We first exponentiate the distribution of ϕ to obtain ρ and then numerically integrate along the lines-of-sight to all stars in our observed dataset. The

¹ Available online at <http://www.gaussianprocess.org/gpml/>

² See the *Gpytorch* documentation for the n-dimensional generalisation of the RBF kernel: <https://docs.gpytorch.ai/en/stable/kernels.html>

numerical integration is described in Sect. 2.5. These integrated densities, that is, the model extinctions (A_{mod}), are compared to our observed extinctions (via the likelihood) in order to optimise the model. The posterior for our model is therefore

$$P(\phi, \Theta | A_{\text{obs}}) \propto P(A_{\text{obs}} | \phi) P(\phi | \Theta) P(\Theta), \quad (3)$$

where the first term on the right does not depend on Θ because A_{obs} is independent of Θ once conditioned on ϕ . As mentioned above, as we have a prior on log density rather than density, we no longer have close form solutions for the integrals over the model densities. We therefore use numerical methods to compute the integrals to achieve a Gaussian approximation of the posterior in ϕ . We do this via variational inference, as explained in the next section. When doing this, we are free to infer the hyperparameters of the GP prior at the same time. As shown in Eq. (3), our model is implicitly hierarchical: the GP is a prior on (the logarithm of) dust density, the plausible values of which are determined by the hyperparameters, and we optimise the dust density as a function of position.

2.3. Variational inference

The above approach to modelling the dust density effectively makes the problem one of inferring the joint distribution of densities at all points, and as a result the problem has a very high-dimensional parameter space. To handle this effectively we use two techniques. The first is to reduce the dimensionality of the problem by conditioning the GP only on a subset of points, known as the inducing points. The locations of these inducing points are parameters that can be optimised along with the other free parameters. Although the use of inducing points lowers the dimensionality, it is still high, because there are three free parameters per inducing point. There is no ‘correct’ choice of the number of inducing points. A rule of thumb is to incorporate as many inducing points as computationally feasible given the available resources (e.g. Wang et al. 2019; Wu et al. 2021; Yi 2020). Results in the literature show that there is minimal gain once the number of inducing points rises above several percent of the full dataset (Wang et al. 2019). We typically use 500–1000 of inducing points in this work, which is as much our computational setup allowed at the time.

To circumvent the still large number of free parameters, we employ variational inference to approximate the target marginalisation integrals. Variational inference replaces the target posterior with an approximate posterior that is easier to work with, and finds the parameters for this approximation that best reproduce the true posterior (Bishop 2006; Blei et al. 2017). This allows the direct computation of the approximate posterior and its gradient with respect to the free parameters, thereby enabling the use of gradient-descent optimisers, as described below in Sect. 2.4.

This approximate posterior is referred to as the variational distribution. As we are using a GP, we use a multivariate Gaussian distribution for this approximate posterior (as described above in Sect. 2.2) to give the joint distribution of ϕ at the inducing points. The mean and covariance of this Gaussian are what we need to infer. In our case, the variational distribution approximates the true (non-Gaussian) posterior of ϕ at the inducing points, analogous to conditioning a GP on a function that is approximating the real (non-Gaussian) distribution, and only conditioning at a subset of the available points (Bishop 2006; Blei et al. 2017).

We want to find that approximate posterior that is most similar to the true one. We achieve this by finding the parameters

of the (variational) distribution that minimises the difference between this distribution and $P(\phi | A_{\text{obs}})$. To minimise this difference, we maximised the evidence lower bound (ELBO). The ELBO is a variational bound, that is to say, it bounds a quantity by a functional over some set of functions (Bishop 2006, Chap. 10). It computes a minimum value for the evidence (the integral of the likelihood over all possible parameter values) given a set of likelihood evaluations.

The motivation for the ELBO is that it is equal to the logarithm of the evidence (which is equal to the expectation of the likelihood over the prior) plus the negative of the Kullback–Leibler (KL) divergence (Bishop 2006; Blei et al. 2017)

$$\text{ELBO}(V) = \mathbb{E}(\log P(A_{\text{obs}} | \phi)) - \text{KL}(V(\phi) \| P(\phi | A_{\text{obs}})), \quad (4)$$

where the KL divergence is a measure of the dissimilarity of two distributions, and V is the variational distribution. In our case the two distributions whose KL divergence is measured are the exact posterior $P(\phi | A_{\text{obs}})$ and the approximation, the variational distribution. This is calculated as:

$$\text{KL}(V(\phi) \| P(A_{\text{obs}}, \phi)) = \int V(\phi) \ln \left(\frac{V(\phi)}{P(\phi | A_{\text{obs}})} \right) d\phi, \quad (5)$$

following Bishop (2006, Chap. 10 Eqs. (10.2)–(10.4)) and Blei et al. (2017, Sect. 2.2 Eqs. (11)–(14)). This can be efficiently minimised using stochastic gradient descent algorithms. The full set of free parameters to be inferred in our model therefore consists of the GP hyperparameters, the locations of the inducing points and the parameters of variational distribution (i.e. the means and covariances).

Once the inference is complete, we have our conditioned GP. From this we can draw samples of ϕ that can be exponentiated to then determine the 16th, 50th, and 84th percentiles of ρ . We can also numerically integrate (as described in Sect. 2.5) the samples of ρ over distance to determine the 16th, 50th, and 84th percentiles of extinction. We used samples to determine confidence intervals as our distributions of ρ and extinction are not Gaussian; if they were, the resulting integration (or sum) would itself be Gaussian; however, we end up needing to sum over log-normal random variates, which do not necessarily have similar convenient properties. Sampling, on the other hand, is computationally efficient.

2.4. Implementation

The GP is implemented using *GPytorch* (Gardner et al. 2018), and *Pyro* (Bingham et al. 2019; Phan et al. 2019) is used for variational inference. The inputs to our algorithm are the positions and extinctions (including uncertainties) of a sample of stars within a region of interest in the Galaxy. We define a ‘training grid’ on which the GP can infer densities. The grid is composed of cells, and every cell is assumed to be uniform, that is, the value of the density is the same at every point within a given cell. The boundaries of this grid are defined in galactic coordinates to enclose the region of interest. The number of cells along each axis of the grid is independent, and can be spaced arbitrarily. The number of cells define the resolution of our training grid and hence the minimum size of the recovered features in the final reconstruction.

We have three sets of free parameters in our model that are optimised. The GP hyperparameters, the locations of the inducing points within our training grid and the means and covariances of the variational distribution. The hyperparameters

of the GP require starting values to begin the optimisation. For the mean of the RBF kernel we estimate the mean density of our region of interest based on prior knowledge and the specific values are given in Sects. 5 and 6. The scale factor of the kernel is initiated from the expected scale (amplitude) of deviations from the mean density in dex (since our model is logarithmic); these values must be below 1 dex because we do not expect variations in dust density of an order of magnitude on scales of 1 pc in the ISM. Finally, we chose the starting scale length based on the typical size of clumps to be recovered. Table 1 below presents the input hyperparameters as well as their conditioned counterparts. We randomly select a subset of the positions of stars in our input sample for the initial locations of the inducing points (however, a user is free to place them in any initial configuration of their choice and can leave these positions fixed throughout the optimisation process if they wish). The variational distribution is initiated by *GPyTorch* with means of zero and variances of 1.

As seen in Eq. (5), estimating ELBO requires integration; this is calculated using 32 Monte Carlo samples at each iteration, and ADaptive Moment Estimation with Weight decay (AdamW; Loshchilov & Hutter 2017), implemented in *Pyro*, is used to maximise ELBO. AdamW has several tunable parameters, including the learning rate, which determines the step size at each iteration. We chose relatively large learning rates to avoid the optimisation becoming stuck in a local minimum or taking too long to optimise, while also avoiding moving around too much and so potentially missing the optimum values while maintaining numerical stability. AdamW is set to terminate training after a fixed number of iterations chosen based on when the variation in the ELBO is less than 1% across the last 10 iterations and the model has converged.

For our training grid, we chose a set of boundaries in all three dimensions that closely encompass our region of interest with padding in l, b, d . This speeds up the run time of the model compared to using a larger region, but requires that we account for the effects of the foreground and background dust that have been omitted. The padding added in the d dimension accounts for the fore and background dust.

In the case of foreground dust, which would directly affect all observed extinctions, the model would increase the density of the closer cells in order to account for the missing foreground cells (i.e. from zero to the lower boundary). To accommodate this we add a set of ‘ghost cells’ to hold this additional dust, similar to an approach used in grid-based hydrodynamics models (e.g. O’Brien & Bussmann 2018). No stars are given as input to these ghost cells, giving the model the freedom to insert foreground dust to explain the extinctions of the stars closest to the lower distance boundary. This avoids biasing the final scale factor and mean density. Fewer ghost cells are required when the foreground extinction is low (for example nearby regions, or at high galactic latitude) while high-extinction lines-of-sight, particularly for distant SFRs in the galactic plane, will require more cells. The exact amount must be tuned on a case-by-case basis.

The case of background regions is slightly more complex. Because our model is conditioned on extinction, and the density at a point is informed by the extinctions of stars at larger distances on nearby lines-of-sight, we require a sufficient number of sources in the background to the star formation region in order to infer the density at a point within the star formation region unambiguously. Restricting the input sources exclusively to the star formation region in which we want to infer the dust density would discard vital information, so sources must be included at greater distances such that they sample the density at a sufficient number of lines-of-sight inside the region of interest. Hence, we

pad the model volume at greater distances. Unlike with the foreground ghost cells, these background cells contain input sources, and so the padding must be increased until a sufficient number of sources are available to encapsulate the underlying structure of the region. Once again, this was determined for each case individually by iteratively increasing the distance until the changes in the structure are no longer significant by eye. Beyond the range probed by the input set of stars, the predicted density will rapidly converge to the mean of the GP prior, as this is the maximum a posteriori value in the absence of data on which to condition. This does not preclude low-extinction lines-of-sight from introducing low dust-density regions in this mean density if the source is sufficiently distant and if there is little dust in front of it, as there is information available here.

Once the GP is conditioned, we are free to predict the density at arbitrary points. We exploit this to produce visualisations of the reconstructed density at higher resolution, making them easier to interpret. We therefore define a grid with two to three times higher resolution than the training grid but the same boundaries, which we refer to as the ‘prediction grid’ hereafter. Although it does not provide any new information compared to the training grid, it is especially beneficial for visualising structures that are small and only a few pixels in size in the training grid.

As our optimisation routine AdamW requires the gradient of the ELBO all quantities are stored as *Torch* tensors or distributions to exploit *PyTorch*’s autograd module. These tools have been scaled to large datasets and give us good performance for our problem.

2.5. Integration along lines-of-sight

The integration of density is approximated as a summation. This is carried out in spherical coordinates with the Sun at the origin to optimise the computations by minimising the number of cell boundary crossings.

We want to compute

$$\mathbf{A}_{\text{mod}} = \int_0^{d_{*,i}} 10^\phi ds_i \quad (6)$$

for each star, where ϕ is the logarithm of extinction density, s_i is the path to the i th star and $d_{*,i}$ is its distance. As ϕ is discretised on a grid, we approximate this integral as the sum

$$\mathbf{A}_{\text{mod}} = \sum_{j=0}^{j_*} 10^{\phi_j} \Delta d_j, \quad (7)$$

where the index j iterates over the grid cells along the line-of-sight, j_* is the index of the cell that contains the i th source, ϕ_j is the density in the j th cell and Δd_j is the length of the path crossing the j th cell. This summation is refined slightly to account for the partial crossing of the final cell on the path, namely,

$$\mathbf{A}_{\text{mod}} = \sum_{j=0}^{j_*-1} (10^{\phi_j} \Delta d_j) + 10^{\phi_{j_*}} (d_* - d_{j_*-1}), \quad (8)$$

where d_{j_*-1} is the distance to the boundary between the $j_* - 1$ th and the j_* th cell.

3. Testing model on simulated data

As an initial test of our model we run it with a set of simulated data with known densities. We generate a 3D density field

Fig. 1. Synthetic extinction estimates used as input for the model, drawn from the simulated density distribution. Each point in the plot corresponds to one simulated star, with a randomly generated (l, b, d) and its corresponding extinction.

comprising a constant background overlaid by three Gaussian overdensities randomly located within a cube. This superposition of clumps on a background provides a simple analogue to the real distribution of matter that is easy to analyse. The background of the density field is set to $1.6 \times 10^{-3} \text{ mag pc}^{-1}$, which gives a total background extinction of 0.5 mag when integrating over the model distance range (see below) providing sufficient extinction to be visible in the model. We generate extinctions towards 120 000 simulated sources within the model volume by integrating along lines-of-sight to a representative sample of coordinates spread randomly throughout the 3D density field. The coordinates are distributed in l, b and d^3 , assuming a uniform space-density of sources. The integration is performed on a grid with boundaries $l = 50^\circ$ to 90° , $b = -10^\circ$ to $+10^\circ$, and $d = 180$ – 500 pc and consists of $l \times b \times d = 100 \times 100 \times 115$ cells. We add zero-mean Gaussian noise to these extinctions with a standard deviation equal to 10% of the true line-of-sight extinction. These noisy extinctions, along with their noise-free 3D positions, serve as the input to the model and are shown in Fig. 1. We only apply the distance uncertainty method of Rezaei Kh. et al. (2020) to our real star formation data and it is not applied here.

Our algorithm was trained on a grid of $n_l \times n_b \times n_d = 20 \times 20 \times 35$ cells³ with 1000 inducing points and the output predictions were made on a grid of $n_l \times n_b \times n_d = 100 \times 100 \times 105$ cells to be able to directly compare to the ground truth maps. The grid boundary coordinates were the same during training and for the prediction. Figure 2 shows the variation and convergence of the ELBO for our model run. The variations in ELBO over the last few iterations are below the 1% level, demonstrating that the model has converged.

As we can see from our comparison of the simulated and reconstructed data, our model reproduces the overall shape of the extinction and density distributions well. In the histogram presented in Fig. 3 the predicted source-by-source extinction follows the same trend as true extinction with only a small fraction ($< 10\%$) of the sources being underestimated. We do see that the

model reconstructions of extinction and density are smoothed out compared to the ground truth, as is expected from a GP method. This smoothing is more evident in the plots of the residuals (Figs. 4 and 5). In the total extinction we see an average underprediction of less than 1 magnitude, and in density it is slightly larger, with a typical (maximal) variation of $\sim 12\%$ ($\sim 20\%$).

The underestimation of extinction, particularly the highest extinctions, is driven by a combination of two effects; the density is smoothed, spreading the same extinction over a larger range of distances, and the vanishingly small probability of having sources that sample the very highest extinctions makes it difficult to measure them. On the other hand, densities are both overestimated in some regions and underestimated in others, which also arises because of smoothing – when some material is moved along the line-of-sight, the clump becomes larger and hence the core will be underestimated and the wings slightly overestimated.

The smoothing performed by the GP inevitably introduces some error into the reconstructed positions of structures in density, particularly along the line-of-sight. However, from the predicted and residual plots of density shown in Fig. 5 we can see that this is typically less than one scale length (i.e. features may be offset from their true positions by up to 20 pc in the final map in this case).

For our density reconstructions we also need to consider the typical uncertainty scales on the values of the density. In Appendix B, we present the density and extinction as a function distance for several lines-of-sight, along with cuts along l and b as a function distance. From these figures we infer a typical uncertainty of $\lesssim 20\%$ in density.

Overall, the size of the uncertainties and errors of the output of our model are small, demonstrating that our algorithm provides good performance in both localising structures and estimating their magnitude, provided that the scale length is smaller than the typical molecular cloud size.

The time taken to train our model depends on several factors. (The run times of our model for our star formation region are given in the upcoming Table 1). The number of cells in the training grid is the largest factor, with the run time empirically scaling slightly faster than linear ($\sim O(n^{1.4})$) with the number of cells n . The run time relative to the number of inducing points, however, scales approximately $O(n^{0.95})$ but consistent with linear scaling. Finally, as the number of sources changes, the run time varies $O(n^{0.8})$. The prediction run time makes up a negligible fraction of the run time taking only several hundred seconds compared to the hours – days required for training depending on the described factors above.

4. Data

In the rest of this work, we use data from the stellar parameter catalogue of Foesneau et al. (2022)⁴. Specifically we use their inferred extinction at 547 nm (A_0) as A_{obs} and their distances as d . Our code is wavelength agnostic, however, and could be used with extinctions at any wavelength from any catalogue.

The Foesneau et al. parameters are inferred from photometry from *Gaia* DR2, 2MASS, and WISE together with *Gaia* parallaxes. Their model simultaneously estimates A_0 , distances, R_0 , effective temperature, luminosity, surface gravity, mass and age for each star independently. The models predict

³ The choice of these values affect the final reconstruction; we recommend using the largest grid that fits in memory during integration on the machine used for training.

⁴ Available at <http://dc.zah.uni-heidelberg.de/>

Fig. 2. Variation in ELBO through model iterations. *Pyro* reports ELBO as a loss function and therefore seeks to minimise $-1 \times \text{ELBO}$ instead of directly maximising ELBO, and the y axis on this plot is therefore positive instead of negative. Minimising $-1 \times \text{ELBO}$ is therefore analogous to minimising the KL divergence.

Fig. 3. Comparison of true and predicted extinctions from the simulated data. *Top*: histograms of true extinctions and predicted extinctions. *Bottom*: residual extinctions, i.e. predicted extinctions minus true extinctions, for the set of simulated sources.

Fig. 4. Comparison of true and predicted extinctions from the simulated data. *Top*: simulated extinction map up to 500 pc. *Middle*: our model reconstructed extinction map up to 500 pc. *Bottom*: residual of model predicted minus simulated extinction.

the reddened spectral energy distributions (SEDs) from the data using PARSEC isochrone models (Chen et al. 2014; Marigo et al. 2013; Rosenfield et al. 2016), the ATLAS9 atmospheric library (Castelli & Kurucz 2003), and Fitzpatrick extinction law (Fitzpatrick 1999). Their estimates result from Markov chain Monte Carlo (MCMC) samples of the posterior parameter distribution. During this procedure, the models are interpolated using a neural network. In the regions we study in this present work,

Fig. 5. Comparison of true and predicted densities from the simulated data. *Top:* simulated density map sampled at the indicated distances. *Middle:* Our model reconstructed density map sampled at the same distances as the simulated map. *Bottom:* residual of model-predicted density minus simulated density.

the typical (median) uncertainties in extinction are 0.34 mag, and the fractional parallax uncertainties are typically between 0.13 and 0.19.

We show their Galactic A_0 map and the regions we study in this work in Fig. 6. We do not make any cuts on the set of sources based on stellar type or stellar parameters; we simply include

Fig. 6. Extinctions as a function of Galactic coordinates from the catalogue of [Fouesneau et al. \(2022\)](#) with the star formation regions analysed in this paper highlighted. *Top:* full Galactic extinction map. *Middle left:* Taurus. *Middle right:* Perseus. *Bottom left:* Orion. *Bottom right:* Cygnus X.

all sources that fall within our region boundaries. Young stellar objects (YSOs) often have circumstellar dust that may be incorrectly attributed to interstellar extinction when that is determined in the Fouesneau et al. catalogue. We did not filter out such objects. However, we expect these objects to be rare, because they are often very red and so would lack a BP measurement, and thus be excluded from the catalogue.

Using this catalogue as our input data provides us with advantages compared with previous Studies. First, one can achieve more reliable estimates of stellar parameters by combining multiple spectroscopic and photometric surveys. This catalogue exploits additional data beyond *Gaia* instead of using *Gaia*-DR2 A_G from [Andrae et al. \(2018\)](#). [Leike & Enßlin \(2019\)](#) used those in their analysis and were limited to the systematics of the GDR2 extinctions. Second, infrared indicators such as the Rayleigh Jeans colour excess method (RJCE) ([Majewski et al. 2011](#)), optimised for particular applications, are less sensitive to low column density than optical bands. These A_k estimates need careful considerations to extrapolate their usage on non-giant stars (e.g. [Rezaei Kh. et al. 2020](#)). Finally, the method in

[Fouesneau et al. \(2022\)](#) jointly estimates distance with the extinction (and other properties). As a result, we do not rely on the inverse parallax as distance measurements ([Bailer-Jones 2015](#)), and we obtain a coherent set of input data.

5. Orion

The Orion SFR is the nearest cluster of young, massive stars ($>8 M_\odot$) at an approximate distance of 400 pc ([Menten et al. 2007](#)). The stellar ages in the Orion SFR range from 2 to 12 Myr and the SFR is separated into three main components, the Orion A and Orion B molecular clouds, and the λ Orionis molecular ring ([Bally 2008](#)). Orion has been studied extensively in literature, making it an ideal real-world test case for our model. Its 3D dust density distribution has been explored by [Rezaei Kh. et al. \(2018b, 2020\)](#), who found a foreground cloud in front of the Orion A molecular cloud at 345 pc. The dust extinction distribution was investigated by [Lombardi et al. \(2011\)](#) (hereafter L11) and [Schlafly et al. \(2015\)](#) (hereafter S15). L11 explores the

Fig. 7. Total extinction of Orion reconstructed by our model at a distance of 550 pc.

distances to several clouds within the Orion SFR using foreground stellar densities and infers masses for them, while S15 explore the 3D structure using dust reddening to study dust ring structures in detail. Attempts to understand the structure of Orion are particularly motivated by the need to study the feedback of massive stars on their environment and the impact this has on the star-forming process (Bally 2008).

5.1. Orion model setup

The model setup including the total number sources from Sect. 4, chosen region size along galactic longitude l , galactic latitude b and distance d along line-of-sight (l bounds, b bounds, d bounds), spacing of cells along l, b, d (n_l, n_b, n_d) and initial input and final conditioned Gaussian kernel hyperparameters are presented in Table 1. We begin testing our model for the Orion SFR by setting the scale lengths of the Gaussian kernel to 10 pc, a value we arrive at based on the distance to the complex and the projected sizes of the molecular clouds. We initialised the mean density to that used by Rezaei Kh. et al. (2017). We set the initial scale factor based on the size of the perturbations expected from the mean background. We then ran the model a number of times, starting the next run from the end point of the previous one, until we found a satisfactory region of parameter space for the scale factor and mean density for the final model run. This was an intuitive process, and resulted in the initial values listed in Table 1.

5.2. Inferred structure of Orion

We summarise the coordinates and the sizes of the main components of Orion identified here in Table 2. Figures 7, 8, and C.1 show our inferred dust densities and the line-of-sight extinctions computed from these. These maps recover the structure of the Orion region well when compared to other literature results, such as L11, S15 and Zucker et al. (2020), as well as the *Planck* dust emission (shown in Fig. D.2; Planck Collaboration X 2016). We do not see evidence for a separate foreground cloud along the line-of-sight to Orion A as suggested by Rezaei Kh. et al. (2020). We do, however, see that Orion A is quite extended, stretching from 350 to 450 pc as first suggested by Großschedl et al. (2018) using YSO data from *Gaia* DR2 and later by Rezaei Kh. et al. (2020) using 3D dust densities. The *Planck* 850 μm and dust intensity maps (see Figs. D.1 and D.2) show a bubble at $l = 215^\circ$ and $b = -20^\circ$ at the edge of Orion A in the direction of the Galactic centre. This bubble is also seen in the infrared extinction by L11, and S15 suggests it lies beyond 550 pc. However, our maps show this bubble is closer to us along the line-of-sight, appearing as an over-density in all three of our maps. It is densest at 405 pc while spanning $380 \leq d \leq 430$ pc.

We recover Orion B and find it to be composed of two segments and to be elongated along the line-of-sight (seen in Fig. 8). The component adjacent to Orion A (Ori B comp 1 in Fig. 8) appears at around 380 pc and disappears at around 400 pc before reappearing at larger distances with the densest region

Fig. 8. 3D dust density distribution of Orion sampled at the indicated distances.

around 430 pc (consistent with [Rezaei Kh. et al. 2020](#)), while the component closer to λ Orionis (Ori B comp 2 in Fig. 8) appears around 380 pc and extends up to 420 pc.

λ Orionis itself is the bubble-like structure whose nearside is around 360 pc and the far side is around 110 pc further along the line-of-sight, placing it closer than previous measurements such as [L11](#) who infer a distance 445 ± 50 pc. Additionally, we recover a peak in density at the centre of the ring (on sky) at a distance of 360–390 pc; this was not seen by [L11](#) or [S15](#) but is clearly visible in the *Planck* dust intensity map shown in Fig. D.2.

Another arc-like feature linking Orion A and Orion B can be seen at $205^\circ \leq l \leq 210^\circ$ and $-19^\circ \leq b \leq -15^\circ$ between 400 and 450 pc, possibly corresponding to NGC2170 and the dust

surrounding the Barnard Loop HII region. The large cloud seen at $205^\circ \leq l \leq 215^\circ$ and $b = -5^\circ$ in the *Planck* 850 μ m map and in extinction by [L11](#) is not visible in our results. This may imply the cloud is located beyond 550 pc, possibly indicating its association with the Canis Major SFR instead of Orion.

Finally, we recover a large filament at $l \approx 186^\circ$ visible at 270–300 pc, even though it is not clearly visible as a coherent structure in extinction. This emphasises the power of 3D reconstructions recovering structure lost in integration to 2D (however it must be noted that if differentiated, 2D extinction maps may also have the potential to show similar localisation). We refer to this filament as the “ λ filament” in Fig. 7 as it is closest to λ Orionis. A similar structure is also visible in the *Planck*

Table 1. Summary of model setup and behaviour for the selected star formation regions.

	Orion	Cygnus X	Perseus	Taurus
l bounds ($^\circ$)	$180 \leq l \leq 217$	$73 \leq l \leq 87$	$154 \leq l \leq 164$	$165 \leq l \leq 180$
b bounds ($^\circ$)	$-25.5 \leq b \leq -3.8$	$-4 \leq b \leq 6$	$-25 \leq b \leq -14$	$-20 \leq b \leq -10$
d bounds (pc)	$250 \leq d \leq 550$	$800 \leq d \leq 2200$	$180 \leq d \leq 500$	$40 \leq d \leq 350$
Stars incorporated from (pc)	270	1200	200	90
Training n_l, n_b, n_d (cells)	40, 40, 55	40, 40, 65	30, 30, 35	30, 30, 35
Training resolution (pc)	2.3	3.5	1.0	0.23
Number of sources	87283	703291	92829	107051
Initial scale length (x, y, z ; pc)	10	15	10	5
Final scale length (x, y, z ; pc)	9.5, 9.3, 9.5	19.1, 19.1, 19.2	12.8, 5.8, 6.4	9.2, 6.7, 4.9
Initial mean density (mag pc $^{-1}$)	4.6×10^{-4}	3.0×10^{-3}	4.6×10^{-4}	1.5×10^{-3}
Final mean density (mag pc $^{-1}$)	1.06×10^{-3}	1.2×10^{-3}	6.0×10^{-4}	1.1×10^{-3}
Initial scale factor (dex)	0.297	0.060	0.297	0.291
Final scale factor (dex)	0.231	0.057	0.293	0.136
Learning rate	0.01	0.01	0.01	0.005
Number of iterations	500	500	1000	1500
Number of inducing points	1000	1000	1000	500
Predicting n_l, n_b, n_d (cells)	100, 100, 105	100, 100, 105	100, 100, 105	100, 100, 105
Predicting resolution (pc)	0.9	1.4	0.3	0.07
Run time (hours) ⁽¹⁾	132	287	39	59

Notes. ⁽¹⁾The run times quoted here are the wall clock times for a machine setup of a single node with $2 \times$ Intel(R) Xeon(R) CPU E5-2680 v3 at 2.50 GHz, no hyperthreading, 24 cores, 180 GB RAM.

maps (Figs. D.1 and D.2). This filament may be host to several small clouds such as LDN 1558 and TGU L156 (L11) as they overlap on sky. However, without confirmed distances to these condensations, we cannot corroborate this.

6. Cygnus X, Perseus, and Taurus

We select three more SFRs to demonstrate the capabilities of our algorithm. These were chosen on the grounds that their 3D dust densities have been less well studied, and because they have different properties from Orion. Once again, we summarise the coordinates and the sizes of the main components of these regions identified through this work in Table 2.

Cygnus X is the most massive SFR within 2 kpc of the Sun, with a total mass of $3 \times 10^6 M_\odot$ (Schneider et al. 2006) and cluster ages of up to 18 Myr (Maia et al. 2016). It is home to the largest number of massive protostars and the largest OB stellar association Cygnus OB2 (Guarcello et al. 2013). While the Orion SFR is located south of the Galactic plane, Cygnus X is in the Galactic plane, so shows higher foreground extinction and crowding. Of the four regions studied here, Cygnus X is the furthest away from us, with estimated distances in the range 1300–2000 pc (Rygl et al. 2010, 2012; Schneider et al. 2006).

The Taurus star formation region is the nearest to us at a distance of 145 pc (Yan et al. 2019) and an age below 5 Myr. Deviating from the common embedded cluster mode star formation seen in other SFRs such as Orion and Perseus, Taurus is the prototype low-mass star formation region ($0.7\text{--}1.0 M_\odot$) where stars appear to form in relative isolation (Kraus et al. 2017). The proximity of Taurus to us also makes it one of the largest SFRs on sky. This, combined with its large population of young stars, makes it an interesting star formation region for analysis.

Finally, we study the Perseus star formation region that neighbours Taurus on sky and is at a distance of 310 pc (Yan et al. 2019). Perseus is home to several regions of active or recent star

formation with typical ages of 1–5 Myr (Pavlidou et al. 2021). It is the smallest of the four regions in angular size presented in this work. There are several filamentary structures spreading towards it from the other nearby star formation regions including Taurus and California. With a mass of around $100 M_\odot$, this SFR is the closest region that is still actively forming low- to intermediate-mass stars (Bally et al. 2008).

The model setup for these three regions is similar to the one we used for Orion, bar a few modifications to match properties of these regions as given in Table 1.

6.1. Cygnus X

The distance to Cygnus X has historically proven difficult to determine, in part because of its location in the Galactic Plane. Distances in the literature span from 1.3 kpc (Rygl et al. 2010) to 2 kpc (Schneider et al. 2006). With our 3D density mapping algorithm we are able to localise the densest regions of Cygnus X to 1.3–1.5 kpc as shown by Fig. 9. Our distance estimate is consistent with the maser parallax distances measured by Rygl et al. (2010) and Rygl et al. (2012). Using CO data in which adjacent regions appear to be interacting, Schneider et al. (2006) inferred that Cygnus X is a monolithic structure at 1.7 kpc. However, we detect small clumps throughout our distance range, corroborating suggestions in the literature that the Cygnus X region may contain both the large SFR (which we place at 1.3–1.5 kpc) and a number of smaller clouds spread along the line-of-sight, as first suggested by Dickel et al. (1969) based on extinction measurements. The clumpy structure and upper distance limit also matches well with Zucker et al. (2020). While Zucker et al. (2020) recover some clumps closer to us than 1200 pc, we do not because we do not have any stars closer than 1200 pc included in our model.

We see a separation in the Cygnus X North and South Clouds in Fig. 9 and find that Cygnus X South is closer to us, with the dense region in Cygnus X South starting at 1300 pc. The south

Table 2. Locations and approximate sizes of the main components identified in the star formation regions.

Component	Central l, b coordinates ($^{\circ}$)	los extent (pc)	Notes
<i>Orion:</i>			
Orion complex	200, -15	280–500	
Orion A	210, -20	340–450	
Orion A bubble	215, -20	380–430	
Orion B comp 1	206, -16	380–400, 430–475	Separated into two segments along los
Orion B comp 2	204, -11	380–420	
λ Ori	195, -12	360–470	
λ Ori centre	195, -12	360–390	
AB arc	210, -15	400–450	
Filament λ	$l = 186$	270–300	Could hold LDN 1558 and TGU L156 clumps
<i>Cygnus X:</i>			
Cygnus X complex	80, 0	1300–1500	
Cyg X North	82, 0	1500	Densest region
Cyg X South	79, -1	1350	Densest region
NAN	0, 85	<800	Located closer than Cygnus X
<i>Perseus:</i>			
Perseus complex	160, -20	300–350	
IC348	160, -18	325	
NGC 1333	158, -20	325	
California filament	158, -16	250–350	
Taurus filament	163, -17	300–330	
Perseus filament	157, -24	300–320	Connecting Perseus and Taurus
<i>Taurus:</i>			
Taurus complex	172, -15	110–190	Densest at 150–180 pc
TMC 1	174, -14	145–190	
TMC 2 main	174, -16	110–145	
L1498 clump	170, -19	195	Part of TMC 2
B215/L1506 extended clump	172, -17.5	195–220	
Taurus filament	165, -17	220–250	Connecting Perseus and Taurus
Arcs 1+2	$b < -14^{\circ}$	190–250	Could hold L1540, L1507 and L1503

cloud is densest at 1350 pc, while the north cloud appears to be densest at around 1500 pc. We also see the presence of the North American Nebula (NAN) in the extinction Fig. 10 at $l = 0^{\circ}, b = 85^{\circ}$. At 600 pc (Schneider et al. 2006) it is much closer to us than the region encompassed by our model and therefore not visible in our 3D density plot. Further, as expected given the high extinction in the region and its placement in the Galactic plane, Cygnus X has the highest mean density of any of the regions we examine.

6.2. Perseus

Our density distribution (Fig. 11) shows that Perseus itself lies between 300 and 350 pc. This is larger than the mean distance of 240 pc determined by Lombardi et al. (2010) (hereafter L10) who used foreground stars to arrive at this distance. However, it agrees well with the mean distance of 310 pc from Yan et al. (2019) (using *Gaia* DR2 extinction and parallaxes) and the range of 294–350 pc found by Zucker et al. (2018), who used photometry, *Gaia* DR2 astrometry, and ^{12}CO data.

We identify the two largest clusters IC348 ($l = 160^{\circ}, b = -18^{\circ}$) and NGC1333 ($l = 158^{\circ}, b = -20^{\circ}$) at the eastern and western edges of the cloud complex at similar distances to one another. We also detect a clump just above IC 348 at $l = 160^{\circ}, b = -17.5^{\circ}$ located at a distance of ~ 306 –330 pc. It is densest at a distance of ~ 325 pc and is clearly detected in

both the *Planck* dust emission in Fig. 12 as well as Zucker et al. (2020) who place it at a distance of 331 ± 16 pc. We do not localise the smaller clumps associated with the Perseus SFR. Although IC348 is also visible in our total extinction map (Fig. 13) NGC1333 is not as distinct, further demonstrating the importance of 3D mapping to fully understand the structure of molecular cloud regions.

Interestingly, we also recover the elongated filaments that appear to extend towards Perseus from the Taurus and California molecular clouds when viewed on sky. We also detect filamentary structure extending from the western edge of Perseus away from the Galactic plane. We localise these filaments in 3D for the first time. These filaments are visible in *Planck* emission (see Figs. 12, D.1, D.2) as well as extinction maps presented by L10, but they have not been localised in 3D before this work. The filament associated with California may in fact be on the edge of Perseus with a distance starting at 250 pc out to nearly 350 pc. The filament between Perseus and Taurus corresponds to a distance of 300 pc extending to ~ 330 pc. The filament furthest from the Galactic plane is at 300 pc but appears to be more diffuse than the other filaments and we cannot clearly associate it with Perseus or any other nearby SFRs. We refer to these filaments as the ‘California filament’, ‘Taurus filament’, and ‘Perseus filament’, respectively, in our figures.

Fig. 9. 3D dust density distribution of Cygnus X sampled at the indicated distances.

6.3. Taurus

Taurus has been separated into two components, TMC 1 and TMC 2, based on their sky positions (Lombardi et al. 2010). In our 3D density distribution plot shown in Fig. 14, the Taurus cloud complex extends from 110–190 pc, with the highest density region at 150–180 pc. This is consistent with Yan et al. (2019), who found a mean distance of 145 pc. The dense cloud at $l = 174^\circ, b = -16^\circ$ in Fig. 15 corresponds to the main component of TMC 2. This component is elongated along the line-of-sight (Fig. 14), starting at 110 pc and continuing to 145 pc; therefore, TMC 2 is closer to us than TMC 1. While the main component of TMC 2 (‘TMC 2 main’ in Figs. 15 and 14) is home to

several smaller clumps we do not separate these clumps in our predictions.

TMC 1, on the other hand, appears at 145 pc and continues to 190 pc, showing that it is distinct from TMC 2. Similar to TMC 2 main, we cannot separate the smaller components of TMC 1 in our maps. This suggests that TMC 1 is a coherent structure containing the clumps.

The overdensity at $l = 170^\circ, b = -19^\circ$ and 195 pc can be associated with the L1498 clump, part of TMC 2. Furthermore, the elongated structure at 195–220 pc that appears to connect TMC 2 main and clump L1498 on sky can be associated with the B215/L1506 extended clump, also part of TMC 2 on sky. While the B215/L1506 is clearly visible in 3D density in Fig. 14 it is

Fig. 10. Total extinction of Cygnus X reconstructed by our model at a distance of 1600 pc.

barely visible in total extinction as seen by Fig. 15; this reiterates the importance of 3D density localisation compared to extinction maps where information is lost in the integration along the line-of-sight.

We also see evidence for a structure at 220–250 pc at $l = 165^\circ, b = -17^\circ$ (‘Taurus filament’ in Fig. 14). This likely corresponds to the other (closer) end of the filament identified in the Perseus data, corroborating this feature. It appears to lie at a distance intermediate to the two SFRs, along a trajectory linking them. Whether this is an independent feature or a bridge between the two regions remains to be seen. In our density figure, we also see two arc-like structures at $\sim b < -14^\circ$ appearing one after the other extending from 201 pc to the end of our distance range. This could be associated with the filamentary structure adjoining TMC 1, and holds the L1540, L1507 and L1503 clumps. As noted in the Sect. 3 there are uncertainties associated with the localisation of structure in all four regions. These uncertainties include shifts in position by approximately one scale length.

7. Total mass

To calculate the dust and total masses of our regions, we take our integrated extinction (calculated up to the upper distance boundary) as described in Sect. 2.5 and convert it to a column density using the dust mass-extinction coefficient (opacity) of $\kappa_0 = 26\,000\text{ cm}^2\text{ g}^{-1}$ calculated at 547 nm. To calculate this

opacity we use an absorption coefficient of $8.5 \times 10^3\text{ cm}^2\text{ g}^{-1}$ and albedo of 0.67 at 547 nm⁵ (Draine 2003a,b). We integrate the total extinction along l and b to derive a dust mass as follows:

$$M_d = d_{\max}^2 \int_{b_{\min}}^{b_{\max}} \cos b \, db \int_{l_{\min}}^{l_{\max}} dl \frac{A_{\text{mod},d_{\max}}(l, b)}{1.086\kappa_0}, \quad (9)$$

where the factor $1.086 = \frac{1}{\ln 2.512}$ converts from magnitudes of extinction to optical depth, $A_{\text{mod},d_{\max}}$ is the total extinction map predicted by our algorithm up to upper distance boundary and D is the distance assumed for the mass calculation, which, in our case, is assumed to be the upper distance boundary of our model setup presented in Table 1, $d_{\max}$.

To derive the total mass we use a gas:dust ratio of 124 (Draine 2003a,b) and multiply the dust mass by this ratio. We estimate the uncertainties in the masses in Table 3 by sampling over the uncertainties in our extinction maps, which in turn come from the finite spread of the Gaussian posteriors in the dust densities. Table 3 presents the calculated masses.

Both L11 and S15 derived the total masses for Orion using the same l and b range as each other. This range is approximately two degrees smaller on each axis compared to our measured region and excludes the filamentary structure we detect. They also assumed smaller upper distance boundaries. In order to be

⁵ Data obtained from <https://www.astro.princeton.edu/~draine/dust/dustmix.html>

Fig. 11. 3D dust density distribution of Perseus sampled at the indicated distances.

able to directly compare our masses to [L11](#) and [S15](#), we carried out the following two steps. First, we recalculated the mass from this work using the same l, b boundaries used in the literature, and second we scaled the literature masses up to the $d_{\max}$ used in this work.

Following step 1, our total mass for Orion recalculated to the smaller l, b range used in [L11](#) and [S15](#) is $735^{+260}_{-177} \times 10^3 M_{\odot}$. This estimate is now directly comparable to the literature masses from [L11](#) and [S15](#) scaled to our distance of 550 pc (step 2), which are $504 \times 10^3 M_{\odot}$ and $444 \times 10^3 M_{\odot}$, respectively. We find our recalculated mass to be 1.4 times larger than [L11](#) and 1.6 times larger than [S15](#).

Our input A_0 values (see Sect. 4) are derived using longer wavelengths than either [L11](#) or [S15](#), and hence are sensitive

to higher total extinction. This means we are more likely to recover higher column densities and therefore greater masses in the present work. Another reason for the variations between the masses are the two different approaches used. While we use a conversion from extinction to mass using κ_0 and a dust:gas ratio, both [L11](#) and [S15](#) use an approach that adopts a hydrogen column density to extinction relationship. The typical (systematic) uncertainties on the conversions in each case are approximately a factor of two (e.g. [Draine 2003a](#)).

In the cases of Perseus and Taurus, we use the same l, b range as [L10](#), but not the same $d_{\max}$ boundaries. Therefore, we can skip step 1 given for Orion and only carry out step 2, where we scale up the literature mass up to our $d_{\max}$ in order to be able to directly compare our masses. Once scaled we derive total

Fig. 12. Comparison to *Planck* dust optical depth maps. *Top four panels:* *Planck* dust optical depth maps for Orion (top left), Cygnus X (top right), Perseus (upper middle left), and Taurus (upper middle right). *Bottom four panels:* ratio of total extinction predicted by our model to *Planck* optical depth normalised by the median of the ratio. Predicted extinction contours are overlaid. Ratio plots are arranged in the same pattern as the *Planck* emission maps.

Fig. 13. Total extinction of Perseus reconstructed by our model at a distance of 500 pc.

Table 3. Inferred masses.

Region	Dust mass ($10^3 M_{\odot}$)	Total mass ($10^3 M_{\odot}$)
Orion	$9.1^{+3.2}_{-2.2}$	1130^{+400}_{-270}
Cygnus	$88.2^{+7.0}_{-6.2}$	10900^{+900}_{-800}
Perseus	$1.5^{+0.1}_{-0.1}$	187^{+17}_{-14}
Taurus	$1.2^{+0.1}_{-0.1}$	149^{+17}_{-14}

masses of $117 \times 10^3 M_{\odot}$ for the Perseus region and a total mass of $110 \times 10^3 M_{\odot}$ for the Taurus region for L10. Our total masses are 1.5 times and 1.3 times larger for Perseus and Taurus, respectively, similar to the variation observed in Orion. The causes of variations discussed above also applies to these two regions as well. A recent publication by Zucker et al. (2021) predicted total masses for Perseus, Taurus, Orion A and B as well as λ Ori and found them to be of the order of $10^4 M_{\odot}$, which is inconsistent with our work as well as the works we compare to. We have not identified a clear reason for this and will be exploring it in a forthcoming paper.

For the Cygnus X region, Schneider et al. (2016) derived a total mass of $7.9 \times 10^6 M_{\odot}$ using CO 2–1 and 3–2 data, which we have scaled up to our $d_{\max}$ of 2200 pc. However, the l, b range employed by Schneider et al. (2016) is about a degree

shifted from ours; thus recalculating our mass to match this l, b range we derive $8.7^{+0.7}_{-0.6} \times 10^6 M_{\odot}$, making the two masses now directly comparable. The two masses are consistent with each other with our mass being 1.1 times larger compared to Schneider et al. (2016). Further, the mass of Cygnus X is more than ten times larger than the masses of all three of our other regions, as expected given that Cygnus X is the most massive SFR within 2 kpc while the other regions are low- or intermediate-mass.

8. Comparison to the *Planck* legacy products

In addition to its all-sky flux maps, the *Planck* sub-millimetre survey has produced maps of various physical quantities, such as dust optical depth, at a resolution of around 10 arcmin (Planck Collaboration X 2016). We compare these here to our dust maps.

A natural product to compare to extinction is the sub-millimetre optical depth of dust at 353 GHz (τ_{353}). Like extinction, this measures the column density, but along the entire line-of-sight instead of just to certain stars. Both of these quantities depend only on the dust column density and the dust properties at their respective wavelengths. Extinction is proportional to the dust cross-sections at optical wavelengths, and τ_{353} is proportional to the dust cross-sections at sub-millimetre wavelengths. This means that taking the ratio of the two cancels out the column density, leaving only the dust properties.

We show the *Planck* τ_{353} maps of our star-formation regions in Fig. 12. It also includes ratio maps of our predicted extinction

Fig. 14. 3D dust density distribution of Taurus sampled at the indicated distances.

to *Planck* τ_{353} . This ratio is proportional to the ratio of the optical to the sub-millimetre dust cross-section. A similar comparison of the predicted extinction to *Planck* sub-millimetre flux and predicted dust emission is shown in Figs D.1 and D.2. In Fig. 12, we can identify in the *Planck* optical depth map the same features we see in our extinction map. The global structure of the clouds is reproduced well, but the structure of the finer details, for example filaments in Taurus TMC 2, are not; our model is not able to resolve such small structures on account of the length scale we have adopted, but instead encloses several features that are distinct in the *Planck* data as a single overdensity. Similar behaviour is seen in the smallest substructure in Perseus as well.

The ratio of extinction to τ_{353} changes significantly between high and low-extinction lines-of-sight, as seen in Fig. 12. Lines-of-sight that intersect with the molecular clouds (dense regions) all have rather low ratios of extinction to sub-millimetre (dust) emission, while the values are much higher, by up to a factor of ten, in the lines-of-sight with low density (diffuse) regions. The most likely physical explanation for the variation in the $\frac{A_0}{\tau_{353}}$ ratio between dense and diffuse regions is a change in the ratio of the cross-sections between the visual and the sub-millimetre between these lines-of-sight, for example as a result of changes in grain size or composition. As grains grow, the ratio of the optical to sub-millimetre cross-sections tends to shrink, because the optical properties change from Rayleigh to Mie scattering and

Fig. 15. Total extinction of Taurus reconstructed by our model at a distance of 350 pc.

eventually to the geometric-optics regime, yet the sub-millimetre opacities stay in the Rayleigh regime. This relationship between opacity and extinction (and hence changes in grain properties) has been described in the near-infrared by, for example, Lombardi et al. (2014) in an unresolved sense when looking only within the cores of molecular clouds. We can extend this relationship to the wider environments of our star formation regions and their surroundings. If the dust properties are forced to stay the same, the observed trends could also be explained if there is additional material at larger distances not traced in our extinction maps, but that is detected by *Planck* in emission; this would, however, require an unlikely preference for additional dust clouds only along lines-of-sight that already have dust clouds near to them.

Following the argument above, the variation in the ratio of optical extinction to sub-millimetre emission allows us to map changes in grain properties in detail on large scales in the ISM and molecular clouds. This probes the dust processing (probably grain growth) that occurs as density increases. As the grains get larger, the opacities change, and the range of variation tells us something about how much they change. However, as the plots are normalised, they probe only the changes in size and not the absolute sizes.

Models of star formation show that grains start off in the diffuse ISM (smaller grains and standard composition, higher ratios) and as the cloud collapses the density increases, allowing them to grow larger or change in composition (lower ratios) (e.g.

Guillet et al. 2007; Hirashita & Kuo 2011). What we see in the plot is that the grain size changes relatively little in the outer layers of the clouds (but still enough to differentiate them from the diffuse material around it) and as you go deeper into the cloud the processing accelerates at higher densities, so that the ratio changes by a much larger factor moving from the outer layers of the cloud into their cores than it does when moving from the diffuse medium to the outer parts of the cloud.

In effect, we are mapping the trend between $A(\lambda)/\tau_{353}$ and R_V identified by Zelko & Finkbeiner (2020), showing that this trend can be used to trace where and when grain processing occurs in detail. The trend persists on much smaller scales than those identified by Schlafly et al. (2016), more similar to the assumption of SFRs as the typical scale.

9. Conclusions

We have introduced an algorithm that reconstructs the 3D dust density distribution of our Galaxy from the extinction to star sources along multiple lines-of-sight. Based on latent variable GPs and variational inference techniques, our algorithm enforces positive density and thus monotonically non-decreasing extinction along any given line-of-sight. Our approach scales favourably in computation time and memory up to large samples and distances. It is built on publicly available software for GPs, machine learning, and Bayesian inference, thereby ensuring

portability and reproducibility while maximising the stability of the code.

We applied our algorithm to four well-known SFRs in the Milky Way: Orion, Cygnus X, Perseus, and Taurus. Our predicted 3D dust densities improve on previous work by recovering features in 3D density, for example filaments in Orion and condensations in Taurus, that are otherwise lost in 2D extinction maps. Our maps shed light on debates over the structure of Cygnus X, suggesting that the main SFRs are at 1300–1500 pc and that there is a collection of more distant clumps that likely contributes to its on-sky projection. We find that the filament seen in extinction between Taurus and Perseus is a unified structure that lies upon a line joining the two regions and may connect the two clouds. We suggest that several clusters in the Orion region outside the main SFRs may be components of a large filament at around 300 pc from us; distance measurements to these clusters are required to confirm this.

Based on the predicted extinction maps, we estimate masses for the clouds. Our estimates are comparable to those in the literature, albeit higher by a factor of 1.1–1.5, with a typical uncertainty of 10%, but 35% in the case of Orion. This may be a result of the inclusion of WISE photometry in the determination of the input extinctions, which improved the recovery of the highest column densities.

Comparing our results with the *Planck* maps, we see that lines-of-sight that intersect dense clouds are systematically different from those that do not, most likely as a result of differences in grain size or composition; dense regions have much smaller ratios between optical and sub-millimetre cross-sections, as expected if the average grain size is larger or the composition is different in dense regions. This shows that the previously observed trend towards a lower ratio of optical extinction to sub-millimetre optical depth when passing from diffuse regions to denser regions also holds in a resolved sense.

The flexibility of our method means that it can easily be extended. For example, future work could use multiple kernels to account for different structures in the ISM, enabling structures of variable sizes to be recovered more effectively, something that is particularly relevant when mapping a large fraction of the Galaxy in one go. Alternatively, the model could take into account the wavelength dependence of extinction (R_0) on map changes in the grain-size distribution in three dimensions.

Acknowledgements. This project is partially funded by the Sonderforschungsbereich SFB 881 “The Milky Way System” of the German Research Foundation (DFG). T.E.D. wishes to thank Thomas Müller at the Center for Astronomy Education and Outreach, Haus der Astronomie for putting together the www.mwdust.com website.

References

- Andrae, R., Fouesneau, M., Creevey, O., et al. 2018, *A&A*, **616**, A8
- Babusiaux, C., Fourtune-Ravard, C., Hottier, C., Arenou, F., & Gómez, A. 2020, *A&A*, **641**, A78
- Bailer-Jones, C. A. L. 2015, *PASP*, **127**, 994
- Bally, J. 2008, *Overview of the Orion Complex*, ed. B. Reipurth, 4, 459
- Bally, J., Walawender, J., Johnstone, D., Kirk, H., & Goodman, A. 2008, *The Perseus Cloud*, ed. B. Reipurth, 4, 308
- Bingham, E., Chen, J. P., Jankowiak, M., et al. 2019, *J. Mach. Learn. Res.* **20**, 973
- Bishop, C. 2006, *Pattern Recognition and Machine Learning* (New York: Springer-Verlag)
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. 2017, *J. Am. Stat. Assoc.*, **112**, 859
- Castelli, F., & Kurucz, R. L. 2003, *Proc. IAU Symp.*, **210**, poster A20
- Chen, Y., Girardi, L., Bressan, A., et al. 2014, *MNRAS*, **444**, 2525
- Dickel, H. R., Wendker, H., & Bieritz, J. H. 1969, *A&A*, **1**, 270
- Draine, B. T. 2003a, *A&A*, **41**, 241
- Draine, B. T. 2003b, *ApJ*, **598**, 1017
- Fitzpatrick, E. L. 1999, *PASP*, **111**, 63
- Fouesneau, M., Andrae, R., Dharmawardena, T., et al. 2022, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202141828>
- Gardner, J., Pleiss, G., Weinberger, K. Q., Bindel, D., & Wilson, A. G. 2018, in *Advances in Neural Information Processing Systems*, eds. S. Bengio, H. Wallach, H. Larochelle, et al. 31 (Curran Associates, Inc.)
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, *ApJ*, **810**, 25
- Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *ApJ*, **887**, 93
- Großschedl, J. E., Alves, J., Meingast, S., et al. 2018, *A&A*, **619**, A106
- Guarcello, M. G., Drake, J. J., Wright, N. J., et al. 2013, *ApJ*, **773**, 135
- Guillet, V., Pineau Des Forêts, G., & Jones, A. P. 2007, *A&A*, **476**, 263
- Hanson, R. J., Bailer-Jones, C. A. L., Burgett, W. S., et al. 2016, *MNRAS*, **463**, 3604
- Hirashita, H., & Kuo, T.-M. 2011, *MNRAS*, **416**, 1340
- Kraus, A. L., Herczeg, G. J., Rizzuto, A. C., et al. 2017, *ApJ*, **838**, 150
- Lallement, R., Vergely, J. L., Valette, B., et al. 2014, *A&A*, **561**, A91
- Lallement, R., Babusiaux, C., Vergely, J. L., et al. 2019, *A&A*, **625**, A135
- Leike, R. H., & Enßlin, T. A. 2019, *A&A*, **631**, A32
- Leike, R. H., Glatzle, M., & Enßlin, T. A. 2020, *A&A*, **639**, A138
- Lombardi, M., Lada, C. J., & Alves, J. 2010, *A&A*, **512**, A67
- Lombardi, M., Alves, J., & Lada, C. J. 2011, *A&A*, **535**, A16
- Lombardi, M., Bouy, H., Alves, J., & Lada, C. J. 2014, *A&A*, **566**, A45
- Loshchilov, I., & Hutter, F. 2017, ArXiv e-prints [arXiv:1711.05101]
- Maia, F. F. S., Moraux, E., & Joncour, I. 2016, *MNRAS*, **458**, 3027
- Majewski, S. R., Zasowski, G., & Nidever, D. L. 2011, *ApJ*, **739**, 25
- Marigo, P., Bressan, A., Nanni, A., Girardi, L., & Pumo, M. L. 2013, *MNRAS*, **434**, 488
- Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., & Picaud, S. 2006, *A&A*, **453**, 635
- Menten, K. M., Reid, M. J., Forbrich, J., & Brunthaler, A. 2007, *A&A*, **474**, 515
- O’Brien, A., & Bussmann, M. 2018, *Comput. Fluids*, **165**, 43
- Paszke, A., Gross, S., Massa, F., et al. 2019, in *Advances in Neural Information Processing Systems 32*, eds. H. Wallach, H. Larochelle, A. Beygelzimer, et al. (Curran Associates, Inc.), 8024
- Pavlidou, T., Scholz, A., & Teixeira, P. S. 2021, *MNRAS*, **503**, 3232
- Phan, D., Pradhan, N., & Jankowiak, M. 2019, ArXiv e-prints [arXiv:1912.11554]
- Planck Collaboration X. 2016, *A&A*, **594**, A10
- Rasmussen, C. E., & Williams, C. K. I. 2006, *Gaussian Processes for Machine Learning* (Cambridge, MA: MIT press)
- Rezaei Kh., S., Bailer-Jones, C. A. L., Hanson, R. J., & Fouesneau, M. 2017, *A&A*, **598**, A125
- Rezaei Kh., S., Bailer-Jones, C. A. L., Hogg, D. W., & Schultheis, M. 2018a, *A&A*, **618**, A168
- Rezaei Kh., S., Bailer-Jones, C. A. L., Schlafly, E. F., & Fouesneau, M. 2018b, *A&A*, **616**, A44
- Rezaei Kh., S., Bailer-Jones, C. A. L., Soler, J. D., & Zari, E. 2020, *A&A*, **643**, A151
- Rosenfield, P., Marigo, P., Girardi, L., et al. 2016, *ApJ*, **822**, 73
- Rygl, K., Brunthaler, A., Menten, K. M., et al. 2010, in *10th European VLBI Network Symposium and EVN Users Meeting: VLBI and the New Generation of Radio Arrays*, **10**, 103
- Rygl, K. L. J., Brunthaler, A., Sanna, A., et al. 2012, *A&A*, **539**, A79
- Sale, S. E., & Magorrian, J. 2014, *MNRAS*, **445**, 256
- Sale, S. E., & Magorrian, J. 2018, *MNRAS*, **481**, 494
- Sale, S. E., Drew, J. E., Barentsen, G., et al. 2014, *MNRAS*, **443**, 2907
- Schlafly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, **737**, 103
- Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2015, *ApJ*, **799**, 116
- Schlafly, E. F., Meisner, A. M., Stutz, A. M., et al. 2016, *ApJ*, **821**, 78
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
- Schneider, N., Bontemps, S., Simon, R., et al. 2006, *A&A*, **458**, 855
- Schneider, N., Bontemps, S., Motte, F., et al. 2016, *A&A*, **587**, A74
- Wang, K. A., Pleiss, G., Gardner, J. R., et al. 2019, ArXiv e-prints, [arXiv:1903.08114]
- Wu, L., Miller, A., Anderson, L., et al. 2021, *Proc. Machine Learning Research*, **130**, 2926
- Yan, Q.-Z., Zhang, B., Xu, Y., et al. 2019, *A&A*, **624**, A6
- Yi, W. 2020, Sparse and Variational Gaussian Process (SVGP) – What To Do When Data is Large, <https://web.archive.org/web/20210626082229/https://towardsdatascience.com/sparse-and-variational-gaussian-process-what-to-do-when-data-is-large-2d3959f430e7?gi=c8a1f699647f>
- Zelko, I. A., & Finkbeiner, D. P. 2020, *ApJ*, **904**, 38
- Zucker, C., Schlafly, E. F., Speagle, J. S., et al. 2018, *ApJ*, **869**, 83
- Zucker, C., Speagle, J. S., Schlafly, E. F., et al. 2020, *A&A*, **633**, A51
- Zucker, C., Goodman, A., Alves, J., et al. 2021, *ApJ*, **919**, 35

Appendix A: The approximate posterior

We use two approximations to compute our posterior. The first is that the posterior distribution over $\log_{10}(\rho) = \phi$ can be approximated as a Gaussian. The second is to compute the covariances of the GP prior only using a subset of points at which the prior is defined.

The first approximation arises as follows. We adopted a Gaussian likelihood for the observed extinctions, A_{obs} , which have uncertainties σ_{obs} , namely,

$$\ln P(A_{\text{obs}} | A_{\text{mod}}) = -\frac{1}{2} \sum \left(\frac{A_{\text{obs}} - A_{\text{mod}}}{\sigma_{\text{obs}}} \right)^2 + \text{const.}, \quad (\text{A.1})$$

where A_{mod} are the corresponding extinctions predicted by the model, and the sum is taken over the measured extinctions for different lines-of-sight.

Our goal is to infer the densities, $\rho(s)$, at various distances s along the different lines-of-sight. For one line-of-sight, the integral of these dust densities gives the extinction to a star at distance d ,

$$A_{\text{mod}} = \int_0^d \rho_{\text{mod}}(s) ds. \quad (\text{A.2})$$

Thus the likelihood can be written

$$\ln P(A_{\text{obs}} | \rho_{\text{mod}}) = -\frac{1}{2} \sum \left(\frac{A_{\text{obs}} - \int_0^d \rho_{\text{mod}} ds}{\sigma_{\text{obs}}} \right)^2 + \text{const.} \quad (\text{A.3})$$

In principle, we could use this to solve for ρ_{mod} , for which we require some regularisation of these model densities because they are otherwise degenerate. Specifically, if we adopted a Gaussian prior on ρ_{mod} we would have a Gaussian posterior in ρ_{mod} and so a straightforward closed-form solution (see Sects. 2.2 and 2.3 of [Rezaei Kh. et al. 2017](#)).

However, we also want to impose non-negativity on ρ , and so we chose to infer $\phi = \log_{10} \rho$ instead of ρ . The likelihood is

$$\ln P(A_{\text{obs}} | \rho_{\text{mod}}) = -\frac{1}{2} \sum \left(\frac{A_{\text{obs}} - \int_0^d 10^\phi ds_i}{\sigma_{\text{obs}}} \right)^2 + \text{const.}, \quad (\text{A.4})$$

which makes it explicit that the posterior will not be Gaussian in ϕ when we have a Gaussian prior in ϕ . We therefore have two options for computing the posterior: we could solve for the distribution of ϕ at all points in space by sampling the full posterior (e.g. using MCMC), or we could approximate the posterior with something more tractable. We chose the second option, and approximate the posterior as a Gaussian using variational inference. That is, when we compute the likelihood from Eq. A.4 and combine it with the prior, we assume that the resulting posterior has a Gaussian distribution. This is our first approximation.

Ideally we would calculate the covariances of the GP using all the grid points, but for computational time reasons we do this instead only at a subset of points, known as the inducing points. Using an iterative procedure we then find the set of inducing points, as well as the mean and covariance of the distribution at these points that best approximates the exact posterior over all the data. This is our second approximation. This procedure is implicit to *Gpytorch*.

Once these two approximations have been made, the Gaussian approximation of the posterior can be written

$$P(\phi | \mathbf{x}, \mathbf{x}_i, \phi_i) = \mathcal{N}(\phi | \boldsymbol{\mu}_\phi, \boldsymbol{\Sigma}_\phi), \quad (\text{A.5})$$

where ϕ_i is the set of log densities $\log_{10} \rho$ at the inducing points whose means and covariances are those of the variational distribution (see Sect. 2.3), $\mathbf{x}_i$ is the set of positions of the inducing points, and $\mathbf{x}$ is the set of other positions at which the posterior is computed (e.g. the grid points). The mean of the posterior is

$$\boldsymbol{\mu}_\phi = K_x^T K_{x_i}^{-1} \phi_i, \quad (\text{A.6})$$

where $K_{x_i} = K(x_i, x_i)$ (as defined in Eq. 2) is an $M \times M$ matrix where M is the number of inducing points and $K_x = K(x_i, x)$ is an $M \times N$ matrix where $N = n_l \times n_b \times n_d$ is the total number of points in the grid. The covariance of the posterior is

$$\boldsymbol{\Sigma}_\phi = K_{xx} - K_x^T K_{x_i}^{-1} K_x, \quad (\text{A.7})$$

where $K_{xx} = K(x, x)$ is an $N \times N$ matrix.

Appendix B: Additional figures showing the behaviour of the model with simulated data introduced in Sect. 3

In this appendix we present the additional figures that demonstrate the behaviour of our model using the simulated data and their model predictions introduced in Sect. 3. First we show the simulated versus model predicted extinctions and densities as a function of distance for a selected set of lines-of-sight. This is followed by simulated and model-predicted densities as a function of l and b for the same set of distance points as Fig 5.

Appendix B.1: Extinction and density along selected lines-of-sight for the simulated and model predicted data introduced in Sect. 3

In this subsection we present the extinctions and densities along a set of selected lines-of-sight for our simulated data in comparison to our model predicted data of the simulations. From these figures we infer a typical uncertainty of $\lesssim 20\%$ in density.

Appendix B.2: Simulated and model-predicted densities along the l and b axes for the same distance points as Fig. 5 in Sect. 3

Similar to the previous subsection, here we present the simulated and model-predicted densities, however this time along the l and b axes for specific distances. The combination of these two subsections allows us to infer a typical uncertainty of $\lesssim 20\%$ in predicted density.

Appendix C: 3D extinction distributions predicted by our algorithm

Here we present the integrated extinction to the given distance points predicted by our model for the SFRs Orion (Sect. 5), Cygnus X, Perseus and Taurus (Sect. 6).

Appendix D: Planck 850 μm flux and 550 μm dust intensity maps

Following from Sect. 8, in this appendix we show the Planck 850 μm flux and 550 μm dust intensity maps for our four SFRs. We also show the normalised ratio of Planck data over model predicted total integrated extinction. The overlaid contours are our total model predicted extinctions up to the upper distance boundary introduced by Figs. 7, 10, 13, and 15.

Fig. B.1: Line-of-sight extinctions and densities of the simulated data. *Top*: Density and extinction sampled along selected lines-of-sight for the same simulated data presented in Sect. 3. *Middle*: Model reconstructed densities and extinctions sampled along the same lines-of-sight. *Bottom*: Residual of model-predicted density minus simulated density.

Fig. B.2: Profiles of simulated and model-predicted densities along b for the same distance points as Fig. 5 in Sect. 3. *Top:* $l = 55^\circ$. *Middle:* $l = 65^\circ$. *Bottom:* $l = 80^\circ$.

Fig. B.3: Profiles of simulated and model-predicted densities along l for the same distance points as Fig. 5 in Sect. 3. *Top*: $b = -5^\circ$. *Bottom*: $b = 5^\circ$.

Fig. C.1: 3D extinction distribution of Orion up to the indicated distances.

Fig. C.2: 3D extinction distribution of Cygnus X up to the indicated distances.

Fig. C.3: 3D extinction distribution of Perseus up to the indicated distances.

Fig. C.4: 3D extinction distribution of Taurus up to the indicated distances.

Fig. D.1: Comparison to Planck 850 μm maps. *Top four panels:* Planck 850 μm maps for Orion (top left), Cygnus X (top right), Perseus (upper middle left), Taurus (upper middle right). *Bottom four panels:* Ratio of total extinction predicted by our model to Planck 850 μm normalized by the median of the ratio. Predicted extinction contours are overlaid. Ratio plots are arranged in the same pattern as the Planck emission maps.

Fig. D.2: As Fig. D.1, showing Planck dust intensity.