Charting nearby dust clouds using *Gaia* data only[★]

R. H. Leike^{1,2} and T. A. Enßlin^{1,2}

¹ Max Planck Institute for Astrophysics, Karl-Schwarzschildstraße 1, 85748 Garching, Germany
e-mail: reimar@mpa-garching.mpg.de

² Ludwig-Maximilians-Universität, Geschwister-Scholl Platz 1, 80539 Munich, Germany

Received 22 January 2019 / Accepted 18 July 2019

ABSTRACT

Aims. Highly resolved maps of the local Galactic dust are an important ingredient for sky emission models. Over almost the whole electromagnetic spectrum one can see imprints of dust, many of which originate from dust clouds within 300 pc. Having a detailed 3D reconstruction of these local dust clouds enables detailed studies, helps to quantify the impact on other observables, and is a necessary milestone of larger reconstructions, as every sightline for more distant objects will pass through the local dust.

Methods. To infer the dust density we use parallax and extinction estimates published by the *Gaia* collaboration in their second data release (DR2). We model the dust as a log-normal process using a hierarchical Bayesian model. We also nonparametrically infer the kernel of the log-normal process, which corresponds to the physical spatial correlation power spectrum of the log-density.

Results. Using only data from *Gaia* DR2, we reconstruct the 3D dust density and its spatial correlation spectrum in a 600 pc cube centered on the Sun. We report a spectral index of the logarithmic dust density of 3.1 on Fourier scales with wavelengths between 2 and 125 pc. The resulting 3D dust map as well as the power spectrum and posterior samples are publicly available for download.

Key words. dust, extinction – local interstellar matter – methods: data analysis – solar neighborhood

1. Introduction

Emission and extinction by Galactic dust is a prominent astronomical foreground at many wavelengths. Therefore, knowing its distribution on the 2D sky and in 3D is essential for many astronomical observations. However, dust is also interesting to be studied on its own, as it provides information on the physical conditions in the interstellar medium and informs us about star forming regions. Dust has been mapped out by surveys for a long time, the first notable contribution being [Burstein & Heiles \(1978\)](#). The dust reconstruction by these latter authors, as in most reconstructions of the dust distribution so far, was focused on mapping the dust in 2D on the sky. This can be done by looking at the sky in infrared (IR) wavelengths, where it is dominated by dust emission. However, when mapping out dust using IR emission one is biased by the radiation field, as dust emission is not only proportional to dust density but also to the amount of starlight absorbed by the dust. Furthermore, dust maps that were produced by mapping IR light might contain extended IR sources that are not within our galaxy, as was shown by [Chiang & Ménard \(2019\)](#). On the other hand, a hypothetical cold dust cloud cannot be seen in IR, leading to systematic errors in the analysis of distant targets, for example quasars or the cosmic microwave background (CMB).

For accurate analyses of objects in our Galactic vicinity, it is vital to have a 3D dust map as a foreground model, which informs us about regions that cannot be observed, or can only be observed with less fidelity, due to dust obscuration. The first nonparametric reconstruction of Galactic dust in 3D was published by [Arenou et al. \(1992\)](#). Since then there have been many attempts to chart the dust density in 3D in increasing resolution

and accuracy and for an ever greater part of our galaxy ([Vergely et al. 1997, 2001](#); [Chen et al. 2018](#); [Rezaei Kh et al. 2017, 2018a,b](#); [Gontcharov 2012, 2017](#); [Sale et al. 2014](#); [Sale & Magorrian 2018](#)). A large driving force for 3D dust reconstruction and astronomy in general comes from large surveys like 2MASS ([Skrutskie et al. 2006](#)), Pan-STARRS ([Kaiser et al. 2002](#)) and SDSS/APOGEE ([Albareti et al. 2017](#)) and WISE ([Wright et al. 2010](#)). These surveys provide photometric measurements and some of them provide spectra for thousands of stars, from which the calculation of photometric distances is possible. There are two 3D dust reconstructions based on these data sets that are closely related to the approach taken in this paper: in [Lallement et al. \(2018\)](#) reddening data from 71 000 sources was used to perform a 3D reconstruction using Gaussian process regression. The resulting dust map covers a 4 kpc square of the Galactic plane and 600 pc in perpendicular direction with a voxel size of (5 pc)³. In [Green et al. \(2018\)](#) a 3D dust map was produced by combining the star data of Pan-STARRS and 2MASS, binning it in angular and distance bins, and performing independent Bayesian reconstructions per angular bin. The result is a dust map that covers three quarters of the sky to a distance up to 2 kpc. This reconstruction shows artificial radial structures called the fingers of God effect, in analogy to the well-known phenomenon in cosmology ([Jackson 1972](#); [Tully & Fisher 1978](#)). One way to mitigate this effect is to use higher-accuracy parallax information.

A prominent new survey is performed by the *Gaia* collaboration ([Brown et al. 2016](#)). In its second data release (DR2, [Gaia Collaboration 2018](#)) accurate parallaxes for roughly 2 billion stars are published. The provided catalog also contains estimates of extinction coefficients for a subset of about 88 million stars, using spectral information of the three energy bands of the *Gaia* satellite. Due to the limited spectral information, the accuracy of the extinction coefficients estimated for individual sightlines

[★] The 3D dust absorption maps are also 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/631/A32>

is quite low. For this reason it is recommended by Andrae et al. (2018) to not use the information of individual sightlines but only the joint information of several sightlines. Even though the data quality of individual sightlines is rather low, the sheer amount of data points and the accuracy of the parallaxes outweigh this limitation, as our work shows.

Three-dimensional dust reconstructions have never been performed using *Gaia* data only. Instead, *Gaia* data have only been used for accurate parallax measurements and the more accurate spectral information of other surveys was used (Lallement et al. 2018, for example).

In this paper we present a 3D dust reconstruction using *Gaia* DR2 data only. The results of the reconstruction are provided at the CDS and online¹, and can be used under the terms of the ODC-By 1.0 license. The inference of the unknown dust density from the extinction data is performed by a critical filter, a method for Gaussian process regression with nonparametric kernel learning, first published in Enßlin & Frommert (2011). While the statistical model used here is, apart from minor details, equivalent to the model introduced in that paper, the algorithm to arrive at approximate posterior summary statistics is quite different. The relevant numerical method used in this paper is outlined in Knollmüller & Enßlin (2019), which describes a general method to derive posterior summary statistics for high-dimensional Bayesian inference problems. For a theoretical discussion of the underlying inference framework of information field theory we refer to Enßlin (2019). The algorithm was implemented using the Python package NIFTy5, which is the newest version of the software package NIFTy (Selig 2013; Steininger et al. 2017; Arras et al. 2019a)². Even though mathematical theory, statistical motivations, and numerical details are distributed over the aforementioned papers, this paper is entirely self-contained by describing the whole method.

In Sect. 2 we discuss which part of the *Gaia* data we used. We introduce our statistical model of the interstellar dust density as well as of the measurement in Sect. 3. In Sect. 4 we present a test application of the algorithm using synthetic data, verifying the predictive power of the algorithm. The main results of the dust density reconstruction using *Gaia* data are presented in Sect. 5. This section also contains a brief recommendation on how to use our results. Our dust reconstruction is compared to other 3D dust density reconstructions in Sect. 6 and in Sect. 7 we summarize the findings of this paper.

2. Data

We used the data from the *Gaia* DR2 catalog by Gaia Collaboration (2018) to reconstruct the Galactic dust in the nearby interstellar medium. From the *Gaia* data archive we extract the parallaxes, the *G*-band extinction, the latitude and longitude, as well as their respective uncertainties. A plot of the full *Gaia* extinction data set can be seen in Fig. 1.

We select sources according to the following criteria:

1. the above mentioned data are available for the source;
2. the parallax ϖ is inside a 600 pc cube around the Sun;
3. the relative parallax error is sufficiently low, $\varpi/\sigma_\varpi > 5$;
4. Priam flag 0100001 or 0100002.

Fig. 1. *Top panel:* natural logarithm of the thermal dust emission map produced by Planck Collaboration I (2019). The color-scheme was saturated to visually match that of the *Gaia* dust extinction data shown in the *middle panel*. *Bottom panel:* subset of *Gaia* extinction data within a 600 pc cube centered on the Sun; the data used in this paper. The scale is a natural logarithm of the extinction data in magnitudes. Data points in the same direction were averaged.

The last two criteria are suggested by Andrae et al. (2018). There are about 3.7 million stars selected by these criteria. Figure 1 shows a sky average of the data points used in the reconstruction. In this data plot one can observe structures present also in other dust maps, for example the *Planck* dust map (Fig. 1).

3. Model

3.1. Algorithm

The algorithm is derived from Bayesian reasoning. In Bayesian reasoning, information that some data d provides about a

¹ DOI 10.5281/zenodo.2577337 or at <https://wwwmpa.mpa-garching.mpg.de/~ensslin/research/data/dust.html>

² The version of NIFTy used for this reconstruction is available at <https://gitlab.mpcdf.mpg.de/ift/NIFTy>

quantity of interest s is calculated according to Bayes theorem:

$$P(s|d) = \frac{P(d|s)P(s)}{P(d)}. \quad (1)$$

We note that the quantity of interest can be a (possibly high-dimensional) vector; in our case it is the dust density for every point in space (256^3 degrees of freedom after discretization). There are three main ingredients necessary for the inference of the quantity of interest s :

1. The likelihood $P(d|s)$ of the data d given a realization of the quantity of interest s . We describe our likelihood in Sect. 3.2.

2. The prior $P(s)$ describing the best available knowledge about the quantity of interest s in absence of data. We describe our prior in Sect. 3.3.

3. An inference algorithm that yields a statistical summary of $P(s|d)$ given the joint distribution $P(d, s) = P(d|s)P(s)$ of d and s . We use the inference algorithm described in Knollmüller & Enßlin (2019).

The main quantity of interest s is the logarithmic G -band dust extinction cross-section density $s = \ln(\alpha\rho/\text{pc})$, henceforth referred to as the logarithmic dust density. Hereby, ρ denotes the actual dust mass density and α the average G -band dust cross section per mass. The value of α is uncertain, which is why we report extinction densities, also referred to as dust pseudo-densities.

We approximate the posterior with a Gaussian

$$Q(s) = \mathcal{G}(s - m, D) \quad (2)$$

$$= \frac{\exp\left(-\frac{1}{2}(s - m)^\dagger D^{-1}(s - m)\right)}{|2\pi D|^{\frac{1}{2}}}, \quad (3)$$

by adopting a suitable mean m and uncertainty dispersion D . The approximation is obtained by minimizing the Kullback–Leibler divergence (Kullback & Leibler 1951),

$$\text{KL}(Q, P) = \int dQ \ln \frac{Q}{P}, \quad (4)$$

with respect to the parameters of Q . This approach is known as variational Bayes (Nasrabadi 2007) or Gibbs free energy approach (Enßlin & Weig 2010). The approach we take in finding the unknown approximate posterior mean m and covariance D of Eq. (3) is described in detail in Knollmüller & Enßlin (2019). It can be summarized as follows:

1. Calculate the negative log-probability $-\log(P(s, d))$ for the problem, disregarding normalization terms like the evidence $P(d)$.

2. Perform coordinate transformations of the unknown quantities until those are a priori Gaussian distributed with unit covariance (Kucukelbir et al. 2017; Knollmüller & Enßlin 2018).

3. Choose the class of approximating distributions to be Gaussian with variable mean m and covariance $D = (\mathbb{1} + M_m)^{-1}$, where M_m is the Fisher information metric at the current m . Here $\mathbb{1}$ is the contribution of the prior which was transformed in step 2 to have unit covariance. This uncertainty dispersion is a lower bound to the uncertainty (Cramér 1946; Rao 1947) and has been shown to be an efficient technique to take cross-correlations between all degrees of freedom into account without having to parameterize them explicitly (Knollmüller & Enßlin 2019).

4. Minimize Eq. (4) with respect to m using the Newton Conjugate gradient as a second-order scheme with the covariance D of Q as curvature. The expectation value with respect to Q is hereby approximated through a set of samples drawn from

the approximating distribution Q . Second-order minimization by preconditioning with the inverse Fisher metric is also called natural gradient descent (Amari 1997) in the literature.

3.2. Likelihood

The likelihood $P(d|s)$ can be split into two parts: the first states how the true extinction depends on the dust density and the second states how the actual data is distributed given the true extinction on that line of sight. The first part, which we call the response R , states how the unknown dust extinction density ρ imprints itself on the data. The extinction of light on the i th line of sight L_i is given by the line integral

$$(A_G)_i = [R(\rho)]_i = \int_{L_i} dl \alpha \rho(l). \quad (5)$$

Here α is the average dust cross section per unit of mass and the line of sight $L_i = L_i(\omega)$ is dependent on the true parallax ω . As noted in Sect. 3.1, the value of α is uncertain and we reconstruct the dust extinction density $s = \ln(\alpha\rho/\text{pc})$ instead.

The extinction is additive because the extinction data are given in the magnitudes scale, which is logarithmic. The true parallax ω of the star is uncertain. We assume the true parallax ω to be Gaussian distributed around the published parallax $\tilde{\omega}$ with a standard deviation equal to the published parallax error σ_ω :

$$P(\omega|\tilde{\omega}, \sigma_\omega) = \mathcal{G}(\omega - \tilde{\omega}, \sigma_\omega^2). \quad (6)$$

The parallaxes of *Gaia* DR2 were shown to be Gaussian distributed with incredible reliability by Luri et al. (2018). However, it was also noted by the same authors that there can be outliers. By restricting ourselves to close-by sources for which G -band extinction values are published, we expect to have cut out most of the outliers.

We do not reconstruct the actual positions of the stars in our reconstruction, and therefore we have to marginalize them out to obtain the response:

$$\begin{aligned} P((A_G)_i|\tilde{\omega}_i, \sigma_{\omega_i}, \rho) &= \\ &= \int d\omega_i P((A_G)_i, \omega_i|\tilde{\omega}_i, \sigma_{\omega_i}, \rho) \\ &= \int d\omega_i P((A_G)_i|\omega_i, \tilde{\omega}_i, \sigma_{\omega_i}, \rho) P(\omega_i|\tilde{\omega}_i, \sigma_{\omega_i}, \rho) \\ &= \int d\omega_i P((A_G)_i|\omega_i, \rho) P(\omega_i|\tilde{\omega}_i, \sigma_{\omega_i}) \\ &= \int d\omega_i \delta\left((A_G)_i - \int_{L_i(\omega_i)} dl_i \alpha \rho(l_i)\right) P(\omega_i|\tilde{\omega}_i, \sigma_{\omega_i}) \\ &\approx \delta\left((A_G)_i - \int d\omega_i \int_{L_i(\omega_i)} dl_i \alpha \rho(l_i) P(\omega_i|\tilde{\omega}_i, \sigma_{\omega_i})\right) \\ &= \delta\left((A_G)_i - \int d\omega_i \int_{L_i(0)} dl_i \alpha \mathbb{1}_{[0, \frac{1}{\omega}]}(l_i) \rho(l_i) \mathcal{G}(\omega_i - \tilde{\omega}_i, \sigma_{\omega_i}^2)\right) \\ &= \delta\left((A_G)_i - \int_{L_i(0)} dl_i \alpha \rho(l_i) \text{sf}_{\mathcal{G}}\left(\frac{\frac{1}{l_i} - \tilde{\omega}_i}{\sigma_{\omega_i}}\right)\right). \end{aligned} \quad (7)$$

Here

$$\text{sf}_{\mathcal{G}}(x) = 1 - \int_{-\infty}^x dt \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}t^2\right) \quad (8)$$

denotes the survival function of a standard normal distribution and

$$\mathbb{1}_{[a,b]}(x) = \begin{cases} 1 & \text{for } x \in [a, b] \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

denotes the indicator function for the closed interval $[a, b]$. We note that we made an approximation where we replace the true extinction with the expected extinction. As a consequence the lines of sight are smeared out by the parallax uncertainty in our approximation. This smoothing can be regarded as a first-order correction for the uncertainty of the parallax and was already used by Vergely et al. (2001). A fully Bayesian analysis would treat the true parallax as unknown and infer these along the other unknowns, but this is beyond the scope of this work.

For the algorithm, the integral in Eq. (7) is discretized into a weighted sum, such that each voxel contributes to the line integral over L_i with weight equal to the length of the line segment of L_i within that voxel. That weight is further discounted by the probability $P(l|\omega, \sigma_\omega)$ of that voxel being on the line of sight. Applying the response R thus takes $\mathcal{O}(N_{\text{data}}N_{\text{side}})$ operations, where $N_{\text{data}} = 3\,661\,286$ is the number of data points used in the reconstruction and $N_{\text{side}} = 256$ is the number of voxels per axis. Due to the large number of data points, evaluating the response on a computer turns out to be numerically expensive. The inference algorithm (see Sect. 3.1) is a minimization for which this response has to be evaluated many times. To make the dust inference feasible in a reasonable amount of time we restricted our reconstruction to a 600 pc cube centered on the Sun.

The second part of the likelihood states how the published data $\widetilde{A}_G$ is distributed given the true extinctions $(A_G)_i$. We use the data likelihood recommended by the *Gaia* collaboration in Andrae et al. (2018). This likelihood assumes the data $\widetilde{A}_G$ to be distributed according to a truncated Gaussian with a global variance $N = (0.46 \text{ mag})^2 \mathbb{1}$. This leads to the likelihood

$$P(\widetilde{A}_G|\rho) = \prod_i \frac{\mathcal{G}(\widetilde{A}_{Gi} - R(\rho)_i, N_{ii})}{\text{cdf}_{\mathcal{G}(R(\rho)_i, N_{ii})}(A_G^{\max}) - \text{cdf}_{\mathcal{G}(R(\rho)_i, N_{ii})}(A_G^{\min})} \quad (10)$$

$$\text{for } d \in [A_G^{\min}, A_G^{\max}]. \quad (11)$$

Here $\text{cdf}_{\mathcal{G}(R(\rho)_i, N_{ii})}$ denotes the cumulative density function of a normal distribution with mean $R(s)_i$ and variance N_{ii} . We took the boundaries of the truncated Gaussian to be $A_G^{\min} = 0$ and $A_G^{\max} = 3.609 \text{ mag}$ as recommended in Andrae et al. (2018).

3.3. Prior

We assume the dust density to be a positive quantity that can vary over orders of magnitude. The dust is assumed to be spatially correlated, statistically homogeneous, and isotropic. The statistical model is constructed to be as general as possible with these two properties in mind.

To reflect the positivity and to allow variations of the dust density by orders of magnitude we assume the dust density ρ to be a priori of log-normal distribution with

$$\alpha\rho = \rho_0 \exp(s), \quad (12)$$

$$\text{where } s \sim \mathcal{G}(s, S) \quad (13)$$

is assumed to be Gaussian distributed with Gaussian process kernel S . Here, $\rho_0 = 1/1000 \text{ pc}$ is a constant introduced to give ρ the correct unit and to bring it to roughly the right order of magnitude. By using an exponentiated Gaussian process we

Fig. 2. Graphical representation of the data model for our reconstruction. The logarithmic dust density $\log(\rho)$ is a Gaussian process with a smooth Gaussian process power spectrum $S_{kk'} = 2\pi\delta(k - k')P_s(k)$. The true extinctions A_G are directly dependent on the dust density ρ on each line of sight. The measured extinctions $\widetilde{A}_G$ are assumed to be distributed around the true extinctions A_G following a truncated Gaussian distribution as described in Sect. 3.2.

allow the dust density to vary by orders of magnitude while simultaneously ensuring that it is a positive quantity. In Eq. (13), S is the prior covariance. If we assume no point or direction to be special a priori, then according to the Wiener-Khinchin theorem S can be fully characterized by its spatial power spectrum $S_{kk'} = 2\pi\delta(k - k')P_s(k)$. We nonparametrically infer this power spectrum $P_s(k)$ as well. There are two main motivations to reconstruct the power spectrum. From a physical perspective the power spectrum provides valuable insight into the underlying processes. From a signal-processing point of view, many linear filters can be identified with a Bayesian filter that assumes a certain prior power spectrum. The optimal linear filter is obtained when the power spectrum used for the filter is exactly equal to the power spectrum of the unknown quantity (Enßlin & Frommert 2011). However, the power spectrum of the unknown quantity is usually also unknown, and therefore this must be reconstructed well. While this argumentation holds for linear filters, certainly many aspects of it carry over to nonlinear filters such as the reconstruction performed in this paper.

Figure 2 depicts the hierarchical Bayesian model for the extinction data $\widetilde{A}_G$ resulting from the logarithmic dust density $\ln(\rho)$, which itself is shaped by the power spectrum $P_s(k)$.

Our statistical model for the power spectrum $P_s(k)$ is a falling power law with Gaussian distributed slope and offset modified by differentiable nonparametric deviations. It is, apart from minor details³, an integrated Wiener process (Doob 1953) on log-log-scale.

This is realized by the following formula,

$$\sqrt{P_s(k)} = \exp\left((\phi_m\sigma_m + \bar{m})\log(k) + \phi_y\sigma_y + \bar{y} + \mathbb{F}_{\log(k)t}^{-1}\left(\frac{a}{1 + t^2/t_0^2}\tau(t)\right)\right), \quad (14)$$

where ϕ_m , ϕ_y , and $\tau(t)$ are the parameters to be reconstructed, $\sigma_m = 1$, $\bar{m} = -4$, $\sigma_y = 2$, $\bar{y} = -16$, $\sigma_y = 3$, $a = 11$, $t_0 = 0.2$ are fixed hyperparameters, $\mathbb{F}_{\log(k)t}^{-1}$ denotes the inverse Fourier transform on log-scale, and $V = (600 \text{ pc})^3$ is the total volume of the reconstruction. These hyperparameters settings were determined by trial and error such that data measured from a prior sample have roughly the same order of magnitude as the actual data and such that the dust density varies by more than one order of magnitude in prior samples.

In our reconstruction, the parameter τ for the smooth deviations of the log-log power spectrum was discretized using 128 pixels. The mathematical motivation to take Eq. (14) as a generative prior for power spectra is discussed in Arras et al.

³ The amplitude model given by Eq. (14) is not exactly equivalent to an integrated Wiener process, but is equivalent to it in the limit of vanishing cutoff $t_0 \rightarrow 0$ while still allowing a numerically stable transformation of the prior to a white Gaussian. See Arras et al. (2019b) for details.

Fig. 3. Several prior samples of the logarithmic spatial correlation power spectrum in units of pc^3 .

(2019b). As a rule of thumb, k -modes for which the data constrains the power spectrum very well will be recovered in great detail due to the nonparametric nature of the model. For k -modes on which the data provide little information, the power spectrum will be complemented by the prior which forces it into a falling power law whenever the data are not informative. If the actual physical process deviates strongly from a falling power law for the unobserved k -modes, the prior might artificially suppress or amplify the posterior uncertainty of the result, possibly biasing the uncertainty quantification. Figure 3 shows a few examples of prior samples of power spectra using our choice of hyperparameters. While the individual samples might not appear qualitatively different, it should on the one hand be noted that any kind of power spectrum is representable with our model given enough data and on the other hand that the figure depicts the power spectrum of the log-density on log-log scale. A small deviation in this figure can have a huge impact on the actual statistics. Reconstructing the power spectrum is equivalent to reconstructing the correlation kernel. We show our reconstructed normalized kernel as well as the one assumed by Lallement et al. (2018) in Fig. 4. Certain biases can appear when using a fixed kernel, for example introducing a characteristic length scale of the order of the FWHM of the kernel.

Putting together likelihood and prior, the overall joint information Hamiltonian for our parameters ξ , τ , and ϕ is

$$P(d, \xi, \tau, \phi) = \text{TG}_{A_G^{\min}, A_G^{\min}, 0.46, d}(R(\alpha\rho)) \mathcal{G}((\xi, \phi, \tau)^T, \mathbb{1}), \quad (15)$$

$$\text{where } [R(\alpha\rho)]_i = \int_{L_i} dl \alpha \rho(l), \quad (16)$$

$$\alpha\rho = \frac{1}{1000} \exp\left(\mathbb{F}_{xk}^{-1} \sqrt{P_s(k)} V(\phi, \tau) \xi_k\right), \quad (17)$$

$$\begin{aligned} \sqrt{P_s(k)}(\phi, \tau) = \exp\bigg(& (\phi_m \sigma_m + \bar{m}) \log(k) + \phi_y \sigma_y + \bar{y} \\ & + \mathbb{F}_{\log(k)t}^{-1} \left(\frac{a}{1 + t^2/l_0^2} \tau_t \right) \bigg), \end{aligned} \quad (18)$$

$$\text{TG}_{x_{\min}, x_{\max}, \sigma, \bar{x}}(x) = \prod_i \frac{\mathcal{G}(\bar{x}_i - x_i, \sigma^2)}{\text{cdf}_{\mathcal{G}(x_i, \sigma^2)}(x_{\max}) - \text{cdf}_{\mathcal{G}(x_i, \sigma^2)}(x_{\min})}$$

for $x \in [x_{\min}, x_{\max}]$, (19)

Fig. 4. Log-normal process normalized two-point correlation reconstructed by our method (solid line) and imposed in the reconstruction by Lallement et al. (2018) (dashed line). One can see that the dust is assumed to be strongly correlated at a distance scale of up to about 30 pc. This plot shows normalized 1D cuts through the 3D Fourier transform of the log-normal spatial correlation power spectrum shown in Fig. 7.

which is a truncated Gaussian, and $V = (600 \text{ pc})^3$. The application, calculation of the gradient, and the application of the Fisher metric of Eq. (15) scales almost linearly with the number of voxels N_{side}^3 . More specifically, it takes $\mathcal{O}(N_{\text{data}} N_{\text{side}} + N_{\text{side}}^3 \log N_{\text{side}})$ operations to evaluate Eq. (15), where $N_{\text{data}} = 3\,661\,286$ is the number of data points used in the reconstruction and $N_{\text{side}} = 256$ is the number of voxels per axis.

4. Simulated data test

4.1. Data generation

In this section, a test on simulated data is presented. This test enables us to compare the results of the reconstruction to a known ground truth. As ground-truth for the dust-density, a cube of interstellar medium simulated by the SILCC collaboration (Walch et al. 2015) was used, more specifically from the magneto-hydrodynamic simulation of the interstellar medium *B6-1 pc* at 50 Myr published by Girichidis et al. (2018). This simulation result spans a cube of $(512 \text{ pc})^3$ in size. We computed our synthetic ground truth differential absorption ρ_{mock} from the gas density of the simulation ρ_{sim} via

$$\rho_{\text{mock}}(x - (150, 150, 0)^T) = \sqrt[3]{\rho_{\text{sim}}\left(\frac{512x}{600}\right)} 10^{17} \frac{\text{cm}^3}{\text{g}} \frac{1}{\text{pc}}. \quad (20)$$

Thus, we stretch the 512 pc simulated cube to the 600 pc of our reconstruction, scale it with a constant factor, and shift it by 150 pc. The shift is performed in order to have an underdense region at the center. We also take the third root of the gas density in Eq. (20). There are two reasons for this. A practical motivation for taking the third root is that it reduces the dynamic range. If one does not do this, the sky will be dominated by one very small, but very strongly absorbing blob. A more physical motivation is that very dense regions lead to star formation. These forming stars again reduce the density by blowing the material out of these regions. This feedback mechanism was not included in the simulation by Girichidis et al. (2018) but was shown to have a strong impact on the gas density in a followup simulation by Haid et al. (2018). The third root can be seen as a very

Fig. 5. Results of our test using simulated data. The rows show integrated dust extinction for sightlines parallel to the z -, x -, and y -axis, respectively. *First column:* test reconstruction and *second column:* ground truth synthetic extinction.

crude way of reducing the density in these overdense regions. To obtain the synthetic data from the ground truth differential extinction cube ρ_{mock} , the following operations were performed:

1. Sampling ground truth parallaxes $\omega_i \sim \mathcal{G}(\omega_i - \tilde{\omega}_i, \sigma_i^2)$ according to the parallax likelihood published by the *Gaia* collaboration.

2. Integrating the dust density from the center of the cube to the location of the sampled star location $1/\omega$ using the full resolution of 512^3 voxels⁴.

⁴ We note that the simulation from which the data are used was performed on an adaptive grid. The full resolution of 512^3 is only realized in the high-density regions.

3. Sampling an observed extinction magnitude according to the truncated Gaussian likelihood described in Sect. 3.2.

4.2. Results

We were able to recover a slightly smeared out version of the original synthetic extinction cube. Figure 5 shows integrations with respect to the x -, y -, and z -axis of the synthetic extinction cube and the reconstructed extinction cube are shown. This visually confirms the reliability of the reconstruction.

For a more quantitative analysis, we compared the reconstructed differential extinctions with the ground-truth voxel-wise. More specifically, we computed the uncertainty-weighted

Fig. 6. Normalized histogram of the deviation of the reconstruction from the true solution, in sigmas (gray curve). The black curve is the probability density function of a standard normal distribution, which is plotted as a reference. We note that the values were clipped to the range from -10 to 10 , i.e., the bump in the gray curve at -10 corresponds to outliers that can be up to 250 sigma. These outliers correspond to about 0.15% of all voxels.

residual r

$$r = \frac{\rho_{\text{reconstruction}} - \rho_{\text{ground truth}}}{\sigma_{\rho}}, \quad (21)$$

where $\rho_{\text{reconstruction}}$ and σ_{ρ} are the posterior mean and standard deviation computed from the approximate posterior samples. The ground truth differential extinction was recovered well within the recovered approximate posterior uncertainty, apart from outliers which make up about 0.15% of the voxels. See Fig. 6 for a histogram of the uncertainty-weighted residual. Overall, the reconstruction seems very reliable on a qualitative and quantitatively level within the uncertainty for most of the voxels.

5. Results from *Gaia* data

We reconstruct the dust density in a 600 pc cube using 256^3 voxels, resulting in a resolution of $(2.34 \text{ pc})^3$ per voxel. For our reconstructed volume we also infer the spatial correlation power spectrum of the log-density; see Fig. 7.

In Fig. 8b one can see a projection of the reconstructed dust onto the sky in Galactic coordinates. Figure 9b shows the corresponding expected logarithmic dust density.

In Fig. 8a the projection of the dust reconstruction on the Galactic plane is shown. This view is especially interesting to study the dust morphology as this projection introduces no perspective-based distortion. It is especially suited to identifying underdense regions such as the local bubble in high resolution. A logarithmic plot of the projection on the Galactic plane can be seen in Fig. 9a. We show integrated dust density for sightlines parallel to the x -, y -, and z -axis in Fig. 13.

We provide posterior uncertainty estimation via samples. One should note that these uncertainties might be underestimated due to the variational approach taken in this paper. One can see a map of the expected posterior variance of the sky projection in Fig. 10a and in the plane projection in Fig. 11a. A slice through the galactic plane and corresponding uncertainties can be seen in Fig. 14.

Fig. 7. Log-normal process spatial correlation power spectrum inferred in our reconstruction (solid line) as well as the imposed power spectrum of Lallement et al. (2018) (dashed line). The shaded area around the solid line indicates 1σ error bounds. The y -axis is in units of pc^3 . The functions can be interpreted as the a priori expected value of $|\mathbb{F}\ln(\rho)|^2 / V$, where V is the volume on which the density ρ is defined, and $\mathbb{F}$ is the Fourier transform. The region between $0.0008/\text{pc}$ and $0.426/\text{pc}$ is almost power-law-like with a slope of 3.1 , the spectral index of the power law.

The results of the reconstruction are provided online⁵ and at the CDS, and can be used under the terms of the ODC-By 1.0 license. Proper attribution should be given to this paper as well as to the *Gaia* collaboration (Gaia Collaboration 2016).

Now, we give an overview of known systematic effects and advice on how to use the provided dust map. We do not recommend using the outer 15 pc of the reconstruction. Periodic boundary conditions were assumed for algorithmic reasons, which leads to correlations leaking around the border of the cube. The inferred prior correlation kernel (Fig. 4) suggests that correlations are vanishing after 30 pc .

Furthermore, we provide posterior samples. When performing a further analysis of our reconstruction we recommend doing so for every sample in order to propagate errors.

It was observed in Sect. 4.2 that there can be a small number of outliers, that is, differential dust extinction values that are much larger than the reconstructed value, by amounts that cannot be explained by the reconstructed uncertainty.

Finally, we anticipate a perception threshold that leads to the absence of extremely low-density dust clouds. The two main reasons for this are that the truncated Gaussian likelihood provides less evidence in the regime where the extinction is close to zero and that variational Bayesian schemes are known to underestimate errors. Studying a larger volume will shed further light on this subject, as sightlines for more distant stars still provide information about nearby dust clouds. One should note that the overall *Gaia* extinction data provide 20 times more sightlines than were used in this reconstruction.

6. Discussion

Here we discuss qualitative, quantitative, and methodological differences to other dust-mapping efforts. Table 1 shows

⁵ <https://wwwmpa.mpa-garching.mpg.de/~enssln/research/data/dust.html>, or by its doi:10.5281/zenodo.2577337

Fig. 8. *Left column:* integrated dust extinction from -300 to 300 pc for sightlines perpendicular to the Galactic plane. The image covers a 600 pc cube centered around the Sun. The units are e -folds of extinction. The coordinates are Galactic Cartesian coordinates. The Sun is at coordinate (0, 0), the Galactic center is located towards the left of the plot, and the Galactic West is located towards the top. *Right column:* all-sky integrated dust extinction maps of the same region, but for sightlines towards the location of the Sun. *First row:* result of the reconstruction discussed in this paper, *second row:* reconstruction performed by Lallement et al. (2018), *last row:* reconstruction by Green et al. (2018).

Fig. 9. Natural logarithmic version of Fig. 8.

a detailed break down of methodological differences to other papers. There are three notable differences between our method and other methods that we would like to stress. On the one hand,

the data set used here is one of the largest of its kind used so far. Additionally, we use a large amount of data while still taking 3D correlations into account. Furthermore, we reconstruct the

Fig. 10. Uncertainty of the reconstruction of this paper derived from posterior samples (*first row*) and of the reconstruction of [Green et al. \(2018\)](#) (*second row*), both in the sky projection. The uncertainties are in the same units as the corresponding maps in Fig. 8, while the logarithmic uncertainties are dimensionless. *First column*: variance for the dust extinction and *second column*: variance of the logarithmic projected dust density on a natural log-scale, which can be interpreted as a relative error.

spatial correlation power spectrum. The motivation and impact of this was already briefly discussed in Sect. 3.3

We compare our dust map to other maps. Comparisons to 2D dust maps are only possible on a qualitative level, since it is not clear whether structures visible in the 2D maps that are not present in the 3D map are simply further away or are too noisy in the data for the algorithm to pick them up. On a qualitative level it is possible to see several morphological similarities between our reconstruction in Fig. 8b and the *Planck* dust map ([Planck Collaboration I 2019](#)) in Fig. 1. These figures also show that many dust structures that are not inside the Galactic plane are local features.

The two 3D dust maps mentioned in Sect. 1 permit a more thorough analysis. Figure 8 shows a compilation of projected dust densities for our reconstruction as well as the reconstruction of [Lallement et al. \(2018\)](#) and [Green et al. \(2018\)](#)⁶, restricted to the same volume as the reconstruction discussed in this paper. A logarithmic version of this figure is provided in Fig. 9.

⁶ It should be noted that there is a new version ([Green et al. 2019](#)) that appeared during the revision of this paper.

While our map seems to agree on large scales with the other maps, there seems to be a pronounced tension in the predictions of the position of some dust clouds compared to the reconstruction of [Lallement et al. \(2018\)](#). Compared to the map of [Green et al. \(2018\)](#) our map recovers the small scales significantly better and suffers far less from radial smearing. It should be noted that [Green et al. \(2018\)](#) mapped a significantly larger part of our galaxy, and that the region that overlaps with our map was declared as relatively unreliable by the authors themselves. The differences are probably due to the different nature of the data sets used. *Gaia* DR2 which was used in our reconstruction contains a much greater number of data points than the data set used for the other reconstructions. These data points, taken from *Gaia* DR2, have a very small parallax error. Additionally, our reconstruction takes the full 3D correlation structure into account.

Our reconstruction and that of [Green et al. \(2018\)](#) allow uncertainties to be quantified using samples. A plot of uncertainties of the dust density reconstructions projected into the Galactic plane can be seen in Fig. 11. Uncertainties of the dust density reconstructions in the sky projection can be seen in Fig. 10.

Fig. 11. Posterior uncertainty of the reconstruction of this paper derived from samples (*first row*) and of the reconstruction of [Green et al. \(2018\)](#) (*second row*) in the plane projection. The uncertainties are in the same units as the corresponding maps in Fig. 8, while the logarithmic uncertainties are dimensionless. *First column*: variance for the dust extinction and *second column*: variance of the logarithmic projected dust density on natural log-scale which can be interpreted as a relative error.

To quantify the dynamic range of the reconstruction and as a prediction on the variability of the logarithmic dust density we calculated histograms of dust density which show the number of voxels that can be attributed to each density. These histograms can be seen in Fig. 12a. One can see that the histogram of our reconstruction extends slightly more towards high dust densities and substantially towards low dust densities. This is possibly because our reconstruction is more sharply resolved, thus regions of high dust density get captured better and bleed less into the regions where dust is absent.

We characterize the agreement between pairs of those reconstructions using the heat maps of their voxel-wise value pairs. These heat maps can be seen in Fig. 12. For two perfectly agreeing reconstructions, the heat map shows a line with a slope of 1. Again it can be seen that the dust density in our reconstruction varies significantly more than in the two other reconstructions. While all maps agree more or less for high dust densities, our dust map exhibits a vastly greater volume with low dust density.

The reconstruction of [Lallement et al. \(2018\)](#) is performed using Gaussian process regression, as is ours. Therefore, one can

compute the prior Gaussian process correlation power spectrum used in their reconstruction. Figure 7 shows both ours and their inferred power spectrum. These two power spectra agree more or less for the larger modes (low k), where the data are very constraining.

One can empirically compute power spectra of the dust density using a Fourier transformation. A comparison plot with all the three mentioned reconstructions can be found in Fig. 15. This shows a white noise floor in the reconstruction of [Green et al. \(2018\)](#), which can visually also be seen as small scale structures in the plane projections shown in Figs. 8 and 9.

7. Conclusions

We were able to obtain highly resolved map of the local dust density using only *Gaia* data. This map agrees on large scales with previously published maps of [Lallement et al. \(2018\)](#) and [Green et al. \(2018\)](#), but shows significant differences on small scales. These differences might to a large degree stem from the different

Table 1. Comparison of the different dust inference methods with the one performed in this paper.

	This paper	Sale & Magorrian (2018)	Rezaei Kh et al. (2018b)	Lallement et al. (2018)	Green et al. (2018)
Parallax uncertainty	Smoothing only	Marginalization by sampling	Neglected	Neglected	Proper uncertainty handling
Max distance	$300 \sqrt{3}$ pc	5 kpc	6 kpc	$\approx 2 \sqrt{2}$ kpc	3 kpc
Max voxel resolution	2.3 pc	Not applicable	About 200 pc	5 pc	16.4 pc/0.063 pc
Number of data points	3.7 million	6349	21 000	71 357	806 million
Power spectrum inference	Yes	No	No	No	No
Correlations	3D	3D	2D map only	3D	1D correlations only
Positiveness	Yes	Only of reddening	No	Yes	Yes
Statistical method	Variational Bayes	Expectation propagation	Analytic	Maximum posterior	Hamiltonian Monte Carlo
Data sets	<i>Gaia</i> DR2	Synthetic <i>Gaia</i> data	APOGEE	<i>Gaia</i> DR1 + APOGEE + 2MASS	Pan-STARRS + 2MASS

Notes. The first row indicates how the parallax uncertainty of the stars was treated. Hereby smoothing refers to weighting a voxel in the line of sight by the survival function of the star radial distance, as is described in Eq. (7). The distance of the furthest point in the reconstruction is given in the second row. The dimensions of the smallest voxel are given in the third row. For the reconstruction of Sale & Magorrian (2018) the concept of voxel resolution is not readily applicable; Sale & Magorrian (2018) use 140 inducing points spanning a region for which one could evaluate the posterior mean at any point. The resolution for Green et al. (2018) contains two values because the resolution is different in radial and angular directions. The fourth row provides the number of data points used. The fifth row indicates whether or not the power spectrum is inferred. The sixth row states which kind of correlations are assumed for the reconstruction. Whether or not positivity of dust density is enforced can be read in the seventh row. The second-to-last row states the method with which the posterior summary statistics was calculated from the non-normalized log posterior. In the last row the data sets used for the reconstruction are listed.

Fig. 12. Normalized histograms of dust densities (*panel a*) and heat maps of voxel-wise correlations (*panels b–d*). The solid line corresponds to our reconstruction, while the dashed line is the reconstruction of Lallement et al. (2018) and the dash-dotted line is the reconstruction of Green et al. (2018). The other three plots are heat maps of voxel-wise correlations between reconstructed logarithmic dust densities, where the colors show bin counts. The black line in the heat maps is the identity function, corresponding to perfect correlation.

Fig. 13. Reconstructed dust density in different projections. The rows show integrated dust extinction for sightlines parallel to the z -, x -, and y -axis, respectively. In the first row, the Galactic center is located towards the left of the plot, in the other two rows the Galactic north is located towards the top of the plot. The cube is in Galactic coordinates, thus the x -axis is oriented towards the Galactic center and the z -axis is perpendicular to the Galactic plane. *First column:* integrated G -band extinction in e -folds of extinction, *second column:* logarithmic version of the first column.

Fig. 14. Reconstructed dust pseudo-density in a slice of the Galactic plane. *First plot:* differential dust extinction in the plane containing the Sun. *Second plot of first row:* logarithmic version of the first plot. *Second row:* corresponding uncertainty maps. The unit of the dust is *G*-band extinction in *e*-folds per parsec. The coordinates are Galactic Cartesian coordinates.

Fig. 15. Empirical spatial correlation power spectra of the reconstructed mean dust density in parsecs. The black line was computed from our reconstruction, the dark-gray line was computed from the reconstruction of Lallement et al. (2018), and the light-gray line was computed from the reconstruction of Green et al. (2018). For the reconstruction of Green et al. (2018), unspecified voxel values on sightlines that lacked data were replaced with a zero.

data used. Our map shows many structures that are visible in the *Planck* dust map (Planck Collaboration I 2019).

In comparison to previous maps, we were able to improve on 3D resolution while still being mostly consistent on the large scales. A comparison to 2D maps like the *Planck* dust map seems to confirm the features present in our map.

We find that the logarithmic density of dust exhibits a power-law power spectrum with a 3D spectral index of 3.1, corresponding to a 1D index of 1.1. This is a significantly harder spectrum than that expected for a passive tracer in Kolmogorov turbulence, which would be a 1D index of 5/3. The harder spectrum is probably caused by the sharp edges of the local bubble and other ionization or dust evaporation fronts.

In contrast to other dust reconstructions, we predict very low dust densities inside the local bubble. This discrepancy is possibly an artifact of our reconstruction as there are known dust clouds in our vicinity; for example the northern high-latitude shells (Puspitarini & Lallement 2012) and the local Leo cold cloud (Peek et al. 2011). The Leo cold cloud is however considerably smaller than a voxel of our simulation. The possibility that *Gaia* extinction estimates are biased for small distances can also not be excluded.

We hope that by providing accurate reconstructions of the nearby dust clouds, further studies of dust morphology will be possible, as well as the construction of more accurate extinction models for photon observations in a large range of frequency bands.

Acknowledgements. We acknowledge helpful commentaries and suggestion from our anonymous referee as well as fruitful discussions with S. Hutschenreuter, J. Knollmüller, P. Arras and others from the information field theory group at the MPI for astrophysics, Garching. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

References

- Albareti, F. D., Prieto, C. A., Almeida, A., et al. 2017, *ApJS*, **233**, 25
- Amari, S.-i. 1997, in *Advances in Neural Information Processing Systems* (Cambridge: MIT Press), 127
- Andrae, R., Fouesneau, M., Creevey, O., et al. 2018, *A&A*, **616**, A8
- Arenou, F., Grenon, M., & Gomez, A. 1992, *A&A*, **258**, 104
- Arras, P., Baltac, M., Ensslin, T. A., et al. 2019a, *Astrophysics Source Code Library* [record ascl:1903.008]
- Arras, P., Frank, P., Leike, R., Westermann, R., & Enßlin, T. A. 2019b, *A&A*, **627**, A134
- Brown, A. G., Vallenari, A., Prusti, T., et al. 2016, *A&A*, **595**, A2
- Burstein, D., & Heiles, C. 1978, *ApJ*, **225**, 40
- Chen, B., Huang, Y., Yuan, H., et al. 2018, *MNRAS*, **483**, 4277
- Chiang, Y.-K., & Ménard, B. 2019, *ApJ*, **870**, 120
- Crámér, H. 1946, *Mathematical Methods of Statistics* (Princeton: Princeton University Press)
- Doob, J. L. 1953, *Stochastic Processes* (New York: John Wiley & Sons), 101
- Enßlin, T. A. 2019, *Ann. Phys.*, **531**, 1800127
- Enßlin, T. A., & Frommert, M. 2011, *Phys. Rev. D*, **83**, 105014
- Enßlin, T. A., & Weig, C. 2010, *Phys. Rev. E*, **82**, 051112
- Gaia Collaboration (Brown, A. G. A., et al.) 2018, *A&A*, **616**, A1
- Gaia Collaboration et al. 2016, *Summary description of Gaia DR1*
- Girichidis, P., Seifried, D., Naab, T., et al. 2018, *MNRAS*, **480**, 3511
- Gontcharov, G. 2012, *Astron. Lett.*, **38**, 87
- Gontcharov, G. A. 2017, *Astron. Lett.*, **43**, 472
- Green, G. M., Schlafly, E. F., Finkbeiner, D., et al. 2018, *MNRAS*, **478**, 651
- Green, G. M., Schlafly, E. F., Zucker, C., Speagle, J. S., & Finkbeiner, D. P. 2019, ArXiv e-prints [arXiv:1905.02734]
- Haid, S., Walch, S., Seifried, D., et al. 2018, *MNRAS*, **482**, 4062
- Jackson, J. 1972, *MNRAS*, **156**, 1P
- Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, *Proc. SPIE*, **4836**, 154
- Knollmüller, J., & Enßlin, T. A. 2018, ArXiv e-prints [arXiv:1812.04403]
- Knollmüller, J., & Enßlin, T. A. 2019, ArXiv e-prints [arXiv:1901.11033]
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., & Blei, D. M. 2017, *J. Mach. Learn. Res.*, **18**, 430
- Kullback, S., & Leibler, R. 1951, *Ann. Math. Stat.*, **22**, 79
- Lallement, R., Capitanio, L., Ruiz-Dern, L., et al. 2018, *A&A*, **616**, A132
- Luri, X., Brown, A. G. A., Sarro, L. M., et al. 2018, *A&A*, **616**, A9
- Nasrabadi, N. M. 2007, *J. Electron. Imag.*, **16**, 049901
- Peek, J., Heiles, C., Peek, K. M., Meyer, D. M., & Lauroesch, J. 2011, *ApJ*, **735**, 129
- Puspitarini, L., & Lallement, R. 2012, *A&A*, **545**, A21
- Rao, C. R. 1947, in *Mathematical Proceedings of the Cambridge Philosophical Society* (Cambridge: Cambridge University Press) 43, 280
- Rezaei Kh, S., Bailer-Jones, C. A., Fouesneau, M., & Hanson, R. 2017, *Proc. IAU*, **12**, 189
- Rezaei Kh, S., Bailer-Jones, C., Schlafly, E., & Fouesneau, M. 2018a, *A&A*, **616**, A44
- Rezaei Kh, S., Bailer-Jones, C. A., Hogg, D. W., & Schultheis, M. 2018b, *A&A*, **618**, A168
- Sale, S., Drew, J., Barentsen, G., et al. 2014, *MNRAS*, **443**, 2907
- Sale, S., & Magorrian, J. 2018, *MNRAS*, **481**, 494
- Selig, M. 2013, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering* (Berlin: Springer)
- Skrutskie, M., Cutri, R., Stiening, R., et al. 2006, *AJ*, **131**, 1163
- Steininger, T., Dixit, J., Frank, P., et al. 2017, ArXiv e-prints [arXiv:1708.01073]
- Planck Collaboration I. 2019, *A&A*, accepted [arXiv:1807.06205]
- Tully, R. B., & Fisher, J. R. 1978, *IAU Symp.*, **79**, 31
- Vergely, J.-L., Egret, D., Freire Ferrero, R., Valette, B., & Koeppen, J. 1997, *HIPPARCOS-Venice'97*, **402**, 603
- Vergely, J.-L., Ferrero, R. F., Siebert, A., & Valette, B. 2001, *A&A*, **366**, 1016
- Walch, S., Girichidis, P., Naab, T., et al. 2015, *MNRAS*, **454**, 238
- Wright, E. L., Eisenhardt, P. R., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868