

Local dark matter density from *Gaia* DR3 K-dwarfs using Gaussian processes

Laurin Söding *,  Ruben L. Bartel and Philipp Mertsch

Institute for Theoretical Particle Physics and Cosmology, RWTH Aachen University, Department of Physics, Sommerfeldstr 16, D-52074 Aachen, Germany

Accepted 2025 August 14. Received 2025 August 12; in original form 2025 June 4

ABSTRACT

The local dark matter density provides constraints on dark matter models and is of importance for experiments hoping to detect dark matter particles in the laboratory. The advent of extensive survey data calls for more complex physical modelling and more sophisticated statistical analysis, particularly to account for correlated uncertainties. In this paper, we perform a vertical Jeans analysis, including a local approximation of the tilt term, using a sample of 200 000 K-dwarf stars from the *Gaia* DR3 catalogue. After combination with the Survey-of-Surveys catalogue, 160 888 of those have radial velocity measurements. We use Gaussian processes as priors for the covariance matrix of radial and vertical velocities. Joint inference of the posterior distribution of the local dark matter density and the velocity moments is performed using geometric variational inference. We find a local dark matter density of $\rho_{\text{dm}} = 0.0117 \pm 0.0035 M_{\odot} \text{pc}^{-3} = 0.44 \pm 0.13 \text{ GeV cm}^{-3}$ at the Sun's position, which is in agreement with most other recent analyses. By comparing a (z -dependent) Gaussian process prior with a (z -independent) scalar prior for the tilt term, we quantify its impact on estimates of the local dark matter density and argue that careful modelling is required to mitigate systematic biases.

Key words: methods: data analysis – methods: statistical – Galaxy: kinematics and dynamics – solar neighbourhood – dark matter.

1 INTRODUCTION

Most of the mass in the Milky Way is believed to be in the form of dark matter. While a non-gravitational detection is still lacking, the gravitational effects on the dynamics and evolution of galaxies are sizeable and can be quantified with ever-increasing accuracy.

In recent years, the *Gaia* mission (Gaia Collaboration 2023) has presented unprecedented data on the positions and velocities of stars in the nearby region of the Milky Way. This wealth of information has opened up new avenues for studying the structure of the Milky Way and the distribution of dark matter within it, as assumptions that had to be made in the past can now be informed by data.

The increase in data quantity has been accompanied by an increase in the complexity of the models used to analyse and interpret it. Simple fits of extremely symmetric parametric models are no longer sufficient to capture the complexity of the data (Wang et al. 2017; Antoja et al. 2018; Rehemtulla, Valluri & Vasiliev 2022; Sivertsson et al. 2022; Li et al. 2025). To properly estimate statistical uncertainties, Bayesian inference has emerged as the method of choice. However, tracing the posterior distribution, for example, using Markov Chain Monte Carlo (MCMC) methods, is quickly becoming prohibitively expensive as the number of parameters increases.

Another way of performing Bayesian inference is variational inference, that is, approximating the posterior distribution with a simpler parametric distribution. In recent years, new methods based on

information field theory have been developed to perform variational inference in high-dimensional parameter spaces (Knollmüller & Enßlin 2019; Frank, Leike & Enßlin 2021). With these methods, the prior space is not restricted to a few scalar parameters, but can easily be extended to thousands of latent parameters describing, e.g. physical fields on a fine grid.

Extracting the local dark matter density from the data is a non-trivial task and there are many methodological approaches for this. First attempts to model the local matter density using stellar kinematics date back at least a century (e.g. Kapteyn 1922; Oort 1932; Bahcall 1984). As the main observable quantity at our disposal are stellar positions and velocities, most methods focus on the kinematics of some stellar population. For example, there are recent methods based on an analysis of circular velocities (e.g. Salucci et al. 2010; Nesti & Salucci 2013; Pato, Iocco & Bertone 2015; Huang et al. 2016; McMillan 2017; Benito, Cuoco & Iocco 2019; de Salas et al. 2019; Karukes et al. 2019; Lin & Li 2019; Ablimit et al. 2020; Cautun et al. 2020; Sofue 2020; Benito, Iocco & Cuoco 2021; Zhou et al. 2023; Ou et al. 2024), fitting a parametric distribution function (e.g. Bienaymé et al. 2014; Piffl et al. 2014; Binney & Piffl 2015; Cole & Binney 2017), normalizing flows (Lim et al. 2025), halo star mass models (e.g. Kafle et al. 2014; Wegg, Gerhard & Bieth 2019; Hattori, Valluri & Vasiliev 2021), Jeans anisotropic modelling (e.g. Nitschai, Cappellari & Neumayer 2020; Nitschai et al. 2021), modelling of phase-space-spirals (e.g. Widmark et al. 2021; Guo et al. 2022), very local analyses (e.g. Holmberg & Flynn 2000; Schutz et al. 2018; Buch, Leung & Fan 2019), and vertical Jeans analyses (e.g. Garbari et al. 2012; McKee, Parravano & Hollenbach 2015; Xia et al.

* E-mail: soeding@physik.rwth-aachen.de

2016; Hagen & Helmi 2018; Sivertsson et al. 2018; Guo et al. 2020; Salomon et al. 2020; Wardana et al. 2020; Cheng et al. 2024). For a more comprehensive overview and list of previous works, we refer to a review by de Salas & Widmark (2021). While most analyses agree on a local dark matter density of the order of $0.01 M_{\odot} \text{pc}^{-3}$ at the Sun’s position, they differ in their methodological approach and the level of rigour in their treatment of uncertainties.

In this analysis, we use a vertical Jeans analysis to estimate the local dark matter density. We use stellar data from mainly the *Gaia* DR3 catalogue (Gaia Collaboration 2023) and extend them with data from the Survey-of-Surveys (SoS) catalogue (Tsantaki et al. 2022). Doing so, we focus on main-sequence stars, which are abundant enough to yield good statistics in velocity and density space. Setting us apart from previous analyses, we use a non-parametric model for the velocity moments in terms of Gaussian processes and describe correlations between vertical and radial motion of stars as a function of height above and below the disc. This allows us to model the velocity moments in a flexible way without having to assume a specific functional form while naturally propagating various uncertainties. As a result of our Bayesian inference, we obtain samples tracing the posterior probability distribution not only of the local dark matter density but also of the velocity moments. We quantify the effect of the so-called tilt term in the Jeans equation and find that it is subdominant for our analysis.

The Jeans equation connects the density profile of a stellar population and their velocity moments to the gravitational potential of the Milky Way (Binney & Tremaine 2008). The gravitational potential is then connected to the total mass density of the Milky Way by means of the Poisson equation. The total mass density can then be separated into a visible and a dark matter component.

The paper is structured as follows: In Section 2, we present the data used in this analysis, including the disc mass model and the selection of our tracer population. In Section 3, we present the methods used including the Poisson–Jeans system of equations, the variational inference method, and our forward model connecting the data to the parameters. In Section 4, we present the results of our analysis, including the local dark matter density and the velocity moments, before we conclude in Section 5.

2 DATA

In recent times, due to the *Gaia* mission, data sets containing 6D stellar measurements, that is, 3D positions and 3D velocities, have grown significantly. For a Jeans analysis, such data are crucial as they constrain the velocity statistics, which are related to the gravitational potential of the Milky Way. If one knew the gravitational potential, constraining the dark matter component would be a problem of component separation. We will make use of the *Gaia* DR3 catalogue (Gaia Collaboration 2023) extensively. We extend the *Gaia* catalogue with data from the SoS catalogue (Tsantaki et al. 2022) to obtain a larger sample of stars with 3D velocity information. For the reference, we will fix some general numbers that are used for coordinate transformations, noting that their exact values are not important for the results of this analysis. We assume a height of the Sun above the Galactic disc of $z_{\odot} = 14.5 \text{ pc}$ (Skowron et al. 2019), a Solar radius of $R_{\odot} = 8.15 \text{ kpc}$ (Reid et al. 2019), a circular velocity of the local standard of rest of $v_{\odot} = 232 \text{ km s}^{-1}$ (McMillan 2017), and a peculiar motion of the Sun of $(U_{\odot}, V_{\odot}, W_{\odot}) = (11.1, 12.24, 7.25) \text{ km s}^{-1}$ (Schönrich, Binney & Dehnen 2010).

In the following text, we will shortly present the disc mass model used in Section 2.1. Then, we will discuss the data selection process for our tracer population and how we extract the density profile in

Table 1. Parameters of the visible components j of the disc mass model: local volume mass density in the galactic mid-plane $\rho_{s,0}^j$ and vertical velocity dispersion in the galactic mid-plane $(v_{s,0}^j)^2$. All gaseous components (indicated with *) are assumed to have uncertainties of 50 percent on their densities and the stellar components have 20 percent. Model taken from Flynn et al. (2006).

Component	$\rho_{s,0}^j [M_{\odot} \text{pc}^{-3}]$	$(v_{s,0}^j)^{1/2} [\text{km s}^{-1}]$
H ₂ *	0.021	4.0 ± 1.0
HI(1)*	0.016	7.0 ± 1.0
HI(2)*	0.012	9.0 ± 1.0
Warm gas*	0.0009	40.0 ± 1.0
Giants	0.0006	20.0 ± 2.0
$M_V < 2.5$	0.0031	7.5 ± 2.0
$2.5 < M_V < 3.0$	0.0015	10.5 ± 2.0
$3.0 < M_V < 4.0$	0.0020	14.0 ± 2.0
$4.0 < M_V < 5.0$	0.0022	18.0 ± 2.0
$5.0 < M_V < 8.0$	0.007	18.5 ± 2.0
$M_V > 8.0$	0.0135	18.5 ± 2.0
White dwarfs	0.006	20.0 ± 5.0
Brown dwarfs	0.002	20.0 ± 5.0
Thick disc	0.0035	37.0 ± 5.0
Stellar halo	0.0001	100.0 ± 10.0

Section 2.2. We finish this section by discussing our data fusion with the SoS catalogue in Section 2.4.

2.1 Disc mass model

For separating the dark matter density from other gravitating masses, we need to model those other components. We use the same model as Garbari et al. (2012), that is, we use estimates of the mid-plane densities from Flynn et al. (2006) with uncertainties as proposed by Garbari et al. (2012). For convenience, we list those values in Table 1. For the z -dependence of the density components, we assume an isothermal profile, that is,

$$\rho_s(z) = \sum_j \rho_{s,0}^j \exp\left(-\frac{\Phi(z)}{(v_{s,0}^j)^2}\right), \quad (1)$$

where ρ_s is the total density of visible mass, $\rho_{s,0}^j$ is the mid-plane density of component j , Φ is the gravitational Potential of the Milky Way, and $(v_{s,0}^j)^2$ is the effective velocity dispersion of component j .

Several alternative Galactic models have been developed in the literature, including, for example, the Besançon model (Robin et al. 2003, 2022) and the Galactic model of McMillan (2017). However, the accuracy of any such model remains challenging to evaluate for the following two reasons: first, these models assume a parametric, often highly symmetric Galaxy, whereas – as will become apparent below – the Milky Way exhibits significant asymmetries and inhomogeneities. Secondly, the observational data available to constrain model parameters are incomplete, as certain stellar populations are difficult to observe directly and gaseous components remain undetected or poorly characterized (e.g. the value of the conversion factor from CO to H₂). We do not expect to gain much by using a more complex model in this analysis. While there are undoubtedly improvements in modern models, they do not sufficiently address the most pressing issue of describing a non-ideal system. Therefore, we will use the model from Flynn et al. (2006) as it allows for a direct comparison of our results with previous works. We will however incorporate some newer analysis of the stellar and

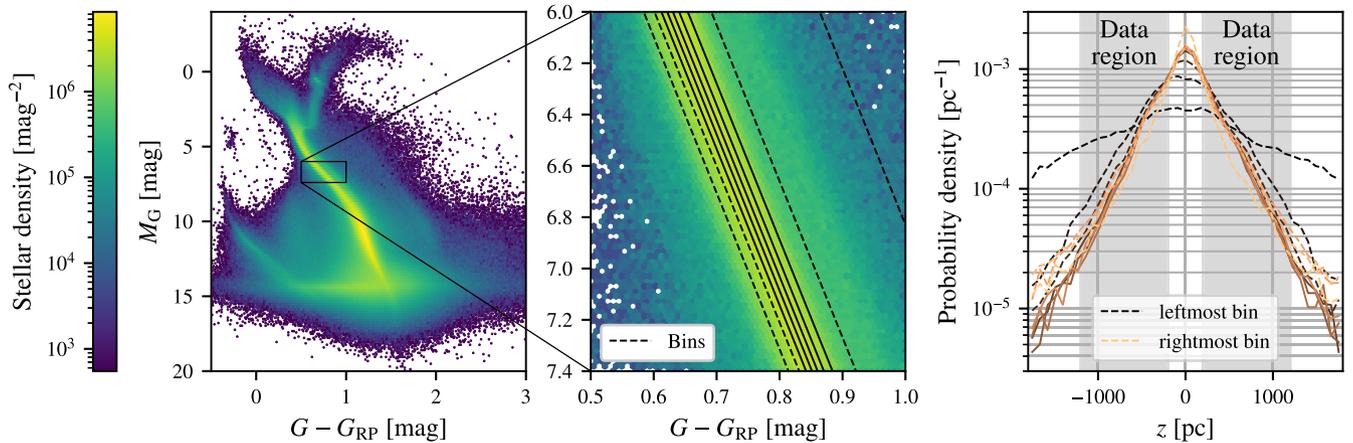


Figure 1. *Left panel:* colour–magnitude diagram of the initially selected stars in the *Gaia* DR3 catalogue within our cylindrical volume. *Centre panel:* zoom-in on the considered region within the main sequence. The lines indicate the edges of our equal-number bins. Every bin, except for the rightmost one, contains 50 000 stars. After applying our cuts, only stars in-between the solid lines were kept. *Right panel:* the vertical density profile of the stars in the selected bins. The dashed lines indicate discarded bins. All lines are normalized so that the area under the curve is equal to 1.

gaseous components in the form of surface mass densities within 1.1 kpc of the Galactic mid-plane by McKee et al. (2015). They estimate $\Sigma_{g,1.1} = 12.6 \pm 1.6 M_{\odot} \text{pc}^{-2}$ for the gaseous component and $\Sigma_{*,1.1} = 31.2 \pm 3.0 M_{\odot} \text{pc}^{-2}$ for the stellar component.

2.2 Stellar tracer population

The stellar population to be analysed necessarily has to fulfil the following two criteria: first, it has to be long-lived so that it is in dynamical equilibrium with the gravitational potential of the disc. Secondly, it has to be bright enough so that volume-complete measurements are available up to sufficiently large distances above and below the Sun’s position in the Galactic disc. To this end, we make use of the *Gaia* DR3 catalogue (Gaia Collaboration 2023), selecting stars within a cylinder of radius 200 pc and height $|z| \leq 1800$ pc as computed using the mean estimated parallax, corrected according to Lindegren et al. (2021). Using the mean parallax here omits some stars whose true parallax would lie inside the region of interest, whereas its estimated mean does not. As we aim to achieve sufficient volume completeness, we will address this later by performing additional cuts. The instruments of *Gaia* are limited in apparent magnitude both towards very bright and very faint stars. In Fabricius et al. (2021), the completeness of *Gaia* data was investigated by comparison to other surveys and models. For our purposes, completeness can be assumed approximately for apparent magnitudes in *G* band between $G \geq 12$ mag and $G \leq 17$ mag. Especially the faint limit on completeness is ill-defined and strongly depends on position and crowding. Since magnitudes as faint as $G \approx 21$ mag can be detected, the upper limit above should be understood as a conservative estimate. An apparent magnitude lower (upper) limit, together with a minimum (maximum) distance we want to consider, yields the lower (upper) limit on the absolute magnitude. Choosing these limits too tight (lenient) leads to a loss of statistics due to restrictions in distance (magnitude) space.

As a compromise, we restrict our analysis to stars with an absolute magnitude $M_G \in [6.0, 7.4]$, which, for main sequence stars, falls into the K-dwarf category. Using the apparent magnitude limits from above, this corresponds approximately to a distance range of $z \in [200, 1000]$ pc. We look at another quantity as an indicator of completeness, that is, the quality of parallax measurements. The

fraction of stars with a parallax error exceeding 20 per cent increases sharply from almost 0 per cent up to $|z| \approx 1200$ pc to 70 per cent at $|z| \approx 1800$ pc, which could indicate a loss of data quality in this regime. We choose this value as our upper limit and note that the following results in this paper are insensitive to small changes in this value. This leaves us with 464 172 stars, of which only 4123 have no colour (G_{RP}) information (mostly at small distances where they are negligible in number).

Next, we clean our sample by removing objects that do not lie on the main sequence. Above the main sequence, we expect to find, e.g. unresolved binaries, which would typically appear brighter; and stars with an unusually high metallicity. Below the main sequence, we expect mainly very old low-metallicity sub-dwarfs. Therefore, we define a magnitude-dependent cut in colour–magnitude space parallel to our chosen segment of the main sequence with slope

$$\alpha = \frac{M_2 - M_1}{c_2 - c_1}, \quad (2)$$

where $M_1 = 6$ mag, $M_2 = 7.4$ mag, $c_1 = 0.67$ mag, and $c_2 = 0.9$ mag. We deem our population sufficiently homogeneous, if all selected stars follow the same number statistics in z independent of their position within the main sequence. For that purpose, we perform an equal-numbered binning of stars parallel to the main sequence (bin size 50 000 stars), discarding those bins below and above the main sequence that do not share the same statistics modulo sampling noise (see Fig. 1). We note that the cut above the main sequence has only very little effect on the derived results later, while the cut from below is crucial to obtain a homogeneous sample. This can already clearly be seen in the strongly deviating density profile of the stars in these bins. The final cuts chosen can be parametrized as

$$2.02 \text{ mag} \leq M_G - \alpha(G - G_{\text{RP}}) \leq 2.27 \text{ mag}, \quad (3)$$

where G and G_{RP} denote the apparent magnitudes in the *Gaia* broad and red photometric bands, respectively, and M_G denotes the absolute magnitude in the *G* band. After discarding these bins, we are left with 200 000 stars.

2.3 Density profile

From the selected sample, we compute the density profile in the following way: The number n_* of stars in a given volume as a

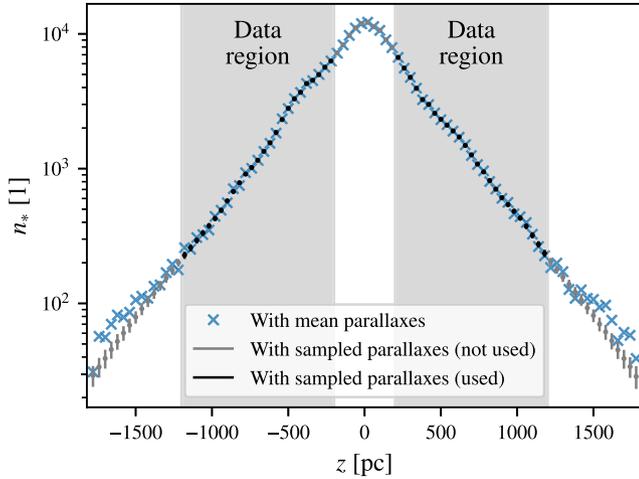


Figure 2. Star counts n_* as a function of height z above the Galactic disc. The crosses assume the mean parallax of each star, while the points show the mean and standard deviation (in log-space) of 1000 realizations sampled from the parallax errors. The bin width is 40 pc.

function of height z is a Poisson process. This demands a binning of the stars in z -space. This binning is subject to uncertainties because the distances, and therefore the heights, are affected by parallax uncertainties. Every star in our sample possesses a measurement of the mean and standard deviation of the parallax p^i , where i indexes the stars. We incorporate this uncertainty by sampling from the parallax distribution of all our stars to create 1000 realizations of the density profile. From these realizations, we compute the mean N_*^k and standard deviation $\sigma_{n_*^k}$ in every bin, where k indexes the z -bins. Since densities are strictly positive, we use a lognormal distribution for describing the uncertainty. On the one hand, this Gaussian approximation requires a sufficiently large number of stars in each bin. On the other hand, binning inherently destroys information and larger bins lead to a loss of resolution. We find a sufficient balance between these two effects using a bin size of 40 pc. The results of this can be seen in Fig. 2. Two effects are important to note: first, using the mean parallax only, the distance is biased towards larger values because parallax and distance are non-linearly related. Secondly, the height restriction of the cylinder leads to a loss of stars in the tails of the distribution because by sampling, stars can effectively be moved out of the cylinder while no new stars can enter. These two effects lead to the divergence of the samples from the raw density profiles at the edges of the cylinder. We find this effect to be negligible for $|z| \leq 1200$ pc. This is because over 98 per cent of these stars have a parallax error smaller than 20 per cent, so that the mean parallax is a reasonably good approximation of the true parallax.

2.4 Stellar velocities

The *Gaia* DR3 catalogue (Gaia Collaboration 2023) contains measurements of the proper motion on the plane of sky for all 200 000 stars in our selection, but radial velocities are only available for 154 217 stars, mostly at smaller distances and often with large uncertainties. To extend this towards larger distances, we combine our sample with data from APOGEE (Majewski et al. 2017), GALAH (De Silva et al. 2015; Buder et al. 2018), GES (Gilmore et al. 2012; Randich, Gilmore & *Gaia*-ESO Consortium 2013), LAMOST (Zhao et al. 2012), and RAVE (Steinmetz et al. 2006, 2020) by means of the SoS catalogue (Tsantaki et al. 2022). To this end, we cross-match

our *Gaia* DR3 sample to *Gaia* DR2 source IDs as are provided by the SoS catalogue. For this, we use the `dr2_neighbourhood` table accompanying the *Gaia* DR3 catalogue. Our DR3 stars contain a few duplicate cross-matches, where there are multiple DR2 neighbours within the search radius. We only keep the entry with the smallest angular distance, that is, the closest neighbour. Having cross-matched a one-to-one correspondence, we overwrite *Gaia* radial velocities with the SoS values where available. This leads to a total of 160 888 stars with radial velocity measurements. With only *Gaia* DR3 data, radial velocity data extend only up to approximately 600 pc above and below the disc. With the SoS catalogue, this is extended up to all considered heights ($z \leq 1800$ pc) in positive z direction. In negative z direction, there is only a minor improvement of statistics and height range.

3 METHODS

3.1 Poisson–Jeans system of equations

There are several methods with varied assumptions connecting the statistical properties of a stellar population with the gravitational potential governing their motion (Binney & Tremaine 2008). We will make use of a particularly simple way, the so-called ‘minimal assumption method’ presented in Garbari, Read & Lake (2011) and Garbari et al. (2012) (consult these papers and references therein for details). We will briefly present the main points in the following text. The collisionless Boltzmann equation describes the time evolution of the phase-space density of some tracer population in a gravitational potential Φ . Expressing it in terms of the first three moments of the distribution function, we obtain the so-called Jeans equations. Expressing these in cylindrical coordinates and assuming an axisymmetric system in a steady state, we obtain a system of two equations. Only the vertical equation is of interest to us, as it connects the total density distribution (by means of the gravitational potential) to the observable velocity moments and density fall-off of the tracer population:

$$\frac{\partial v}{\partial z} + v \left(\frac{\partial \Phi}{\partial z} + \frac{\partial \overline{v_z^2}}{\partial z} + \frac{1}{Rv} \frac{\partial (Rv \overline{v_z v_R})}{\partial R} \right) = 0. \quad (4)$$

Here, v is the density fall-off of some stellar population, $\overline{v_z^2}$ and $\overline{v_z v_R}$ are velocity moments, z is the height above the Galactic disc, Φ is the gravitational potential, and R is the cylindrical radius. The last expression in the parentheses is called the ‘tilt term’ and is often neglected (see e.g. Sivertsson et al. 2018 or Guo et al. 2020 for notable exceptions). We attempt to improve on this by approximating it as

$$\frac{1}{Rv} \frac{\partial (v R \overline{v_z v_R})}{\partial R} \approx \frac{\overline{v_z v_R}}{R} + \frac{v_z v_R}{R_0} + \frac{v_z v_R}{R_T} \quad (5)$$

assuming that v and $\overline{v_z v_R}$ change on some characteristic length-scales R_0 and R_T . One would naively expect these scales to be comparable to the disc length-scale, which is of the order of 2–3 kpc in magnitude. We compare our expectations to the data by splitting our cylinder in two halves, one closer to the Galactic centre, and one further away, using the mean parallaxes and velocities for estimating these length-scales. We find a sufficient agreement for R_0 being approximately constant over the considered range in z and consistent with the disc length-scale. For R_T however, an estimate is not reliably possible as regions with good velocity statistics (small $|z|$) have negligible $|\overline{v_z v_R}|$, making the length-scale essentially meaningless, and regions with large $|z|$ (and expected large $|\overline{v_z v_R}|$) have too few stars to yield a useful estimate. For a more detailed discussion about the sign and

magnitude of these corrections, see Sivertsson et al. (2018) or Binney et al. (2014).

Sivertsson et al. (2022) analyse the impact of the tilt term using two simulated Milky-Way-like galaxies out of equilibrium and find that careful modelling of the tilt term is necessary for obtaining an accurate estimate of the local dark matter density. They further find that, surprisingly, time-dependent effects are subdominant even if the disc experienced a recent perturbation.

With this approximation, the differential equation can be integrated to yield

$$\frac{v(z)}{v(z_0)} = \frac{\bar{v}_z^2(z_0)}{\bar{v}_z^2(z)} \exp\left(-\int_{z_0}^z \frac{1}{\bar{v}_z^2} \frac{\partial \Phi}{\partial z'} + \frac{\bar{v}_z \bar{v}_R}{\bar{v}_z^2} \left(\frac{1}{R} + \frac{1}{R_0} + \frac{1}{R_T}\right) dz'\right), \quad (6)$$

with some arbitrary reference height z_0 , which we will take to be $z_0 = 0$ in the following text.

Every tracer population evolving in the gravitational potential Φ has to satisfy this equation. The gravitational potential is further constrained by the Poisson equation, here in cylindrical coordinates:

$$\frac{\partial^2 \Phi}{\partial z^2} = 4\pi G (\rho_s(z) + \rho_{\text{dm}}^{\text{eff}}), \quad (7)$$

where we have split the total matter density ρ into a visible part ρ_s and a dark matter part $\rho_{\text{dm}}^{\text{eff}}$ with

$$\rho_{\text{dm}}^{\text{eff}} = \rho_{\text{dm}} - \frac{1}{4\pi G R} \frac{\partial v_c^2(R)}{\partial R} \quad (8)$$

being the effective dark matter density. It contains the local dark matter density (that we assume to be independent of height z within the considered volume), and another term given by the local slope of the rotation curve $v_c = (R\partial\Phi/\partial R)^{1/2}$ but independent of z . This term can be calculated from the Oort constants A and B (Binney & Tremaine 2008) as

$$\frac{1}{4\pi G R} \frac{\partial v_c^2(R)}{\partial R} = \frac{1}{2\pi G} (B^2 - A^2). \quad (9)$$

using the values $A = 14.79 \pm 0.11 \text{ km s}^{-1} \text{ kpc}^{-1}$ and $B = -13.73 \pm 0.11 \text{ km s}^{-1} \text{ kpc}^{-1}$ as computed by Akhmetov et al. (2024).

3.2 Variational inference

Our main interest lies in one single scalar quantity, the local dark matter density ρ_{dm} . It appears in the Poisson equation (7) as one constituent of the mass distribution, which dictates the gravitational potential. This is connected via equation (6) to the vertical density fall-off v of some tracer population and its velocity moments \bar{v}_z^2 and $\bar{v}_z \bar{v}_R$. The latter two are only indirectly observable and have to be inferred from the data. To make matters worse, all measurements and input parameters (such as the disc mass model) come with significant uncertainties, and the data are both noisy and (in the case of the stellar velocities) sparse. This prevents us from using simple histograms to estimate the density fall-off and velocity moments. The quantities \bar{v}_z^2 and $\bar{v}_z \bar{v}_R$ are generally functions of height z that, as properties of a smooth stellar distribution, can be assumed to be somewhat smooth but are not strongly constrained by data beyond a height of a couple of hundred parsecs. We will use Gaussian processes to model these functions and capture their uncertainties while exploiting the concept of local correlations. This allows us to properly incorporate measurement uncertainties, which is a big improvement compared to fitting a function to binned data. This description inevitably leads to a very large number of latent parameters requiring a sophisticated inference method as sampling, using, e.g. MCMC algorithms, is becoming prohibitively costly.

For properly incorporating measurement uncertainties of sparse noisy data we employ a probabilistic approach. In order to deal with the large number of latent parameters, we use a novel variational inference method called ‘Geometric Variational Inference’ (geoVI) (Frank et al. 2021), a generalization of ‘Metric Gaussian Variational Inference’ (MGVI) (Knollmüller & Enßlin 2019). It approximates the posterior probability distribution by a multivariate normal distribution in a transformed coordinate system induced by the Fisher information metric. It has been shown to be a powerful tool for Bayesian inference in high-dimensional parameter spaces and has recently been applied to a variety of problems in astrophysics (e.g. Roth et al. 2023; Hutschenreuter et al. 2024; Tsouros et al. 2024; Edenhofer et al. 2024b; Söding et al. 2025). The algorithm works in an iterative manner by switching between sampling around an expansion point according to the local curvature of the posterior, and updating said expansion point by minimizing the Kullback–Leibler-divergence of the approximated posterior and the true posterior until a self-consistent solution is obtained. The resulting samples trace the true posterior and implicitly contain the correlations between all latent parameters. The algorithm is implemented in the Python package NIFTY8 (Selig et al. 2013; Arras et al. 2019; Steininger et al. 2019; Edenhofer et al. 2024a).

To perform Bayesian inference, we have to set up a prior, a likelihood, and a forward model.

We incorporate the measurement uncertainties on the parallaxes p^i , the proper motions v_μ^i and v_δ^i , and the radial velocities v_r^i of all stars (indexed by i) by sampling from the respective distributions in the prior. In the forward model, we then compute the vertical velocity v_z^i , the cylindrical radial velocity v_R^i , and the vertical position z^i of each star. These will later provide information on the velocity moments \bar{v}_z^2 and $\bar{v}_z \bar{v}_R$. Note that even if the errors in parallax-proper-motion space are uncorrelated, the errors in the cylindrical velocities will be correlated. This is due to projection effects, but also, more importantly, because the parallax to distance transformation is non-linear.

Here, we add a new assumption, namely that the statistics of the velocities of our sample stars in cylindrical coordinates can locally be well described by a multivariate Gaussian distribution. Since we do not incorporate angular velocities, this reduces to a two-dimensional Gaussian distribution in the (v_z, v_R) -plane. The steady-state and axisymmetry assumptions require that the mean velocities \bar{v}_z and \bar{v}_R are zero. The covariance matrix of the velocity distribution is best modelled in its eigenspace. The two eigenvalues Σ_1 and Σ_2 of the covariance matrix Σ are equal to \bar{v}_z^2 and \bar{v}_R^2 if the matrix is diagonal. Inspection of the velocity data shows that the velocity moments as a function of height z can be well modelled by a linear function of $|z|$, plus deviations. We model the deviations using a lognormal Gaussian process by means of the NIFTY8 correlated field model (Arras et al. 2022). This provides the necessary flexibility for the inference to capture the level of smoothness without enforcing a specific functional form or preferred correlation length by hand. Our model for the eigenvalues then takes the form

$$\Sigma_1(z) = (b_1 + m_1|z|) \exp(g_{\Sigma_1}(z)) \quad (10)$$

$$\Sigma_2(z) = (b_2 + m_2|z|) \exp(g_{\Sigma_2}(z)), \quad (11)$$

where g_{Σ_1} and g_{Σ_2} are the Gaussian processes, b_1 and b_2 are the intercepts, and m_1 and m_2 are the slopes of the linear relations.

We expect – and will allow for – deviations from a diagonal covariance matrix by introducing an angle θ rotating the covariance matrix. Our prior consists of two parts: first, a tilt originating from rotations relative to the Galactic centre. This is expected to be the

majority contribution, though may not fit exactly. To give the angle more freedom, we add a second Gaussian process g_θ that is added to a uniformly distributed (arbitrary) starting angle θ_0 :

$$\theta(z) = -\arctan\left(\frac{z}{R_\odot}\right) + \theta_0 + (g_\theta(z) - g_\theta(0)). \quad (12)$$

The covariance matrix is then assembled as

$$\Sigma(z) = R(\theta(z)) \begin{pmatrix} \Sigma_1(z) & 0 \\ 0 & \Sigma_2(z) \end{pmatrix} R(\theta(z))^T, \quad (13)$$

with a 2D rotation matrix $R(\theta)$. The components describe the velocity moments, that is, $\Sigma_{00} = \overline{v_z^2}$, $\Sigma_{01} = \Sigma_{10} = \overline{v_z v_R}$, and $\Sigma_{11} = \overline{v_R^2}$. We demand the stellar velocities to be distributed according to this covariance matrix, which means that the residuals of the velocities are distributed according to a multivariate Gaussian distribution with zero mean and covariance matrix Σ . This constitutes our first likelihood:

$$\text{LH}_1 = \mathcal{G}\left(\begin{pmatrix} v_z \\ v_R \end{pmatrix} \middle| \vec{0}, \Sigma(z)\right). \quad (14)$$

Now that we have a model for the velocity moments, we can turn to the vertical density fall-off $\nu(z)$. For computing it, using equation (6), we need to know the first derivative of the gravitational potential Φ . We can compute this by integrating the Poisson equation (7) with respect to z using the boundary condition

$$\Phi(0) = \frac{\partial \Phi}{\partial z}(0) = 0. \quad (15)$$

The right-hand side of the Poisson equation is given by the sum of the visible mass density ρ_s (see equation 1) and the effective dark matter density $\rho_{\text{dm}}^{\text{eff}}$. We sample the constituents of the former from the values given in Table 1, the dark matter density from a uniform prior between $0 \text{ M}_\odot \text{ pc}^{-3}$ and $0.1 \text{ M}_\odot \text{ pc}^{-3}$, and the Oort's constants from a normal prior according to their uncertainties (cf. equations 8 and 9).

Next, we need the scale lengths R_0 and R_T for the tilt term (equation 5). For the number density scale length R_0 , we choose a Gaussian prior centred around -2.5 kpc with a standard deviation of 0.5 kpc . For the velocity tilt scale length R_T , there is no conclusive data available that confirm the magnitude, constancy, or even the sign of this term. Indications of disequilibria in the velocity distribution that would lead to a contribution that is potentially large in magnitude and varying in height $|z|$ have been known for a long time (e.g. Dehnen & Binney 1998). We therefore model this term as a Gaussian process as well, choosing a broad prior centred around zero. Another option would be to choose a scalar Gaussian prior, assuming that the value is constant in z . To estimate the effect of such a prior choice, we will also discuss results derived using a scalar prior.

With these ingredients, equation (6) can be evaluated to yield the vertical density fall-off $\nu(z)$. To account for the uncertainties in our data due to parallax errors, we multiply this density fall-off by the above-derived (see Section 2.3) star count error $\sigma_{n_*}^k$ in log-space to obtain the density profile $n_*(z)$ in every bin k as

$$n_*^k = n_{*,0} \nu(z^k) \exp(\sigma_{n_*}^k). \quad (16)$$

Here, $n_{*,0}$ is the density at the reference height $z_0 = 0 \text{ pc}$, which is a priori unknown. We draw it from a broad lognormal prior around the actually measured density $N_{*,0}$ (which we cannot take directly because data in this region might be volume-incomplete). This can then be compared to the measured star counts N_*^k using a Poissonian likelihood:

$$\text{LH}_2 = \mathcal{P}(N_*^k | n_*^k). \quad (17)$$

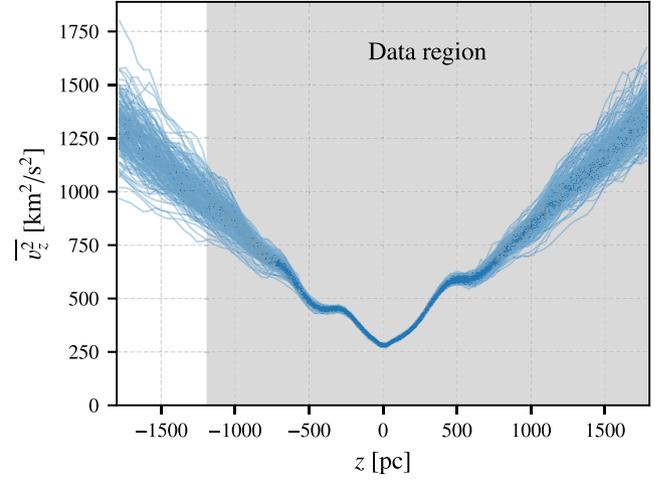


Figure 3. Vertical velocity dispersion $\overline{v_z^2}$ as a function of height z above the Galactic disc. All samples are shown. The shaded region indicates the availability of stellar velocity data.

This procedure captures both the stochastic Poissonian uncertainties in the star counts and the parallax-induced uncertainties in the density fall-off.

Finally, we include the constraints on the surface mass density within 1.1 kpc of stellar and gaseous components of the disc mass model from McKee et al. (2015) as two additional scalar likelihoods (cf. Section 2.1):

$$\text{LH}_3 = \mathcal{G}(\Sigma_{g,1.1} | 12.6 \pm 1.6 \text{ M}_\odot \text{ pc}^{-2}) \quad (18)$$

$$\text{LH}_4 = \mathcal{G}(\Sigma_{*,1.1} | 31.2 \pm 3.0 \text{ M}_\odot \text{ pc}^{-2}). \quad (19)$$

These constraints help in alleviating some degeneracies between baryonic and dark matter components in the disc mass model.

4 RESULTS

We ran the variational inference algorithm using 100 antithetical (200 total) samples for 12 iterations. We report a mean reduced χ^2 (per degree of freedom, averaged over all samples) of 1.0 for the fit indicating that all data points included are statistically consistent with the model. This includes the Poissonian star counts and the multivariate Gaussian velocity data.

The main difference setting our work apart from others is our use of Gaussian processes to model a priori unknown but smooth functions and the variational inference algorithm to sample the posterior distribution. We will first discuss the inferred velocity moments and the fit to the density fall-off, and then the estimated local dark matter density and its dependence on modelling of the tilt term.

There are three velocity moments that we can compute from the inferred covariance matrix: the vertical velocity dispersion $\overline{v_z^2}$, the radial velocity dispersion $\overline{v_R^2}$, and the cross term $\overline{v_z v_R}$. The inferred vertical velocity dispersion is presented in Fig. 3. It shows the approximately linear relation with $|z|$ that we expected. Within the data region, especially towards small $|z|$, the function is tightly constrained by the data. Towards larger $|z|$, where data become sparse, the uncertainties increase significantly as indicated by the spread of samples. We observe an asymmetry in the vertical velocity dispersion above and below the disc (see e.g. the bump at around

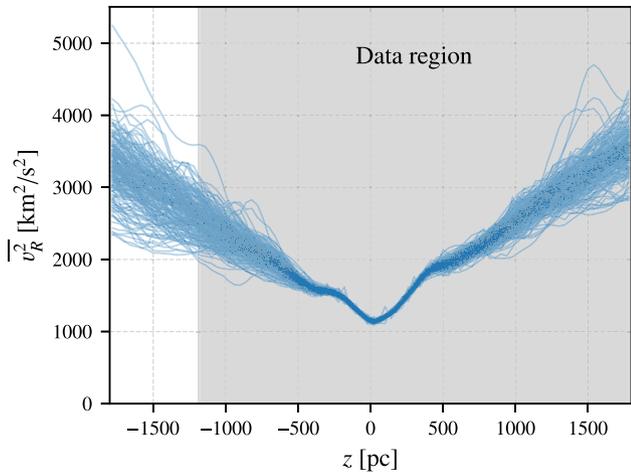


Figure 4. Like Fig. 3, but for the radial velocity dispersion $\overline{v_R^2}$.

$z = \pm 400$ pc), even in regions well constrained by data. This hints towards one of our assumptions being violated, e.g. the assumption of axisymmetry or steady state. In the density fall-off data (see Fig. 2), a similar asymmetry is observed at the same heights, but opposite in sign. This is somewhat expected from the solution of the Jeans equation (equation 6), which relates the density fall-off ν and the vertical velocity dispersion $\overline{v_z^2}$ approximately antilinearly. As the gravitational potential Φ is enforced to be smooth and symmetric, the asymmetries in the other terms have to, and do, cancel each other (within the uncertainty margins) to match the observed star-count and velocity data.

The radial velocity dispersion $\overline{v_R^2}$ is shown in Fig. 4. It is generally larger than the vertical velocity dispersion and also shows an approximately linear relation with $|z|$. A similar asymmetrical behaviour as for the vertical velocity dispersion is observed, but this function is generally less constrained by the data. This may be due to the fact that only the stellar velocity data directly contribute here, while in case of the vertical velocity dispersion, also the star counts contribute indirectly (by means of equation 6). Especially at negative z , where there are no velocity data, many samples deviate significantly from the mean, reflecting the lack of velocity data in this region.

The cross term $\overline{v_z v_R}$ is shown in Fig. 5. It follows a comparatively smoother behaviour than the other two velocity moments and is on average about an order of magnitude smaller. It exhibits, again, a wide spread at large distances and is close to zero in the mid-plane. The gradient at $z = 0$ pc is positive, which indicates that on both sides of the disc, stars that are moving vertically outwards are also preferentially moving radially outwards. A simple tilt in the average stellar orbits leads to the expectation of $\overline{v_z v_R}$ being antisymmetric in z , which we confirm within the uncertainty margins. At least in regions well constrained by data, so at small $|z|$, the cross term, and also the tilt angle, increases, indicating that the correlation of velocity components becomes increasingly important for analyses of stars towards larger distances from the disc. However, velocity data are sparse in these regions, leading to large uncertainties. Within our error estimate, the inferred tilt is consistent with a simple rotation of the stellar orbits around the Galactic centre.

The covariance matrix of the stellar velocities is shown in Fig. 6 for various z . This emphasizes the tilt of the covariance matrix in the considered two-dimensional velocity space. It should be noted that the spread of the samples, which reflects the uncertainty of the

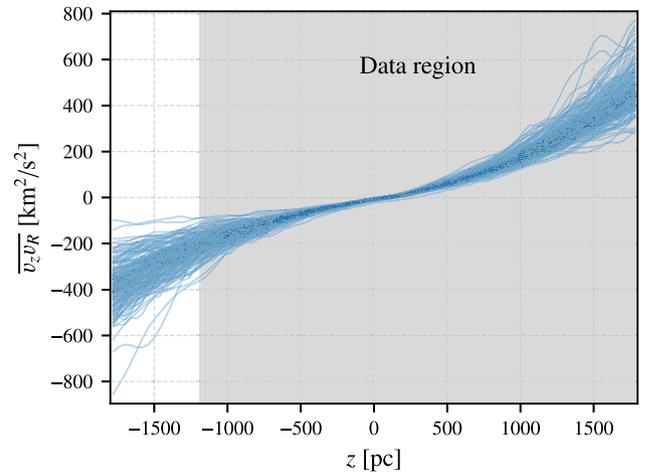


Figure 5. Like Fig. 3, but for the velocity covariance $\overline{v_z v_R}$.

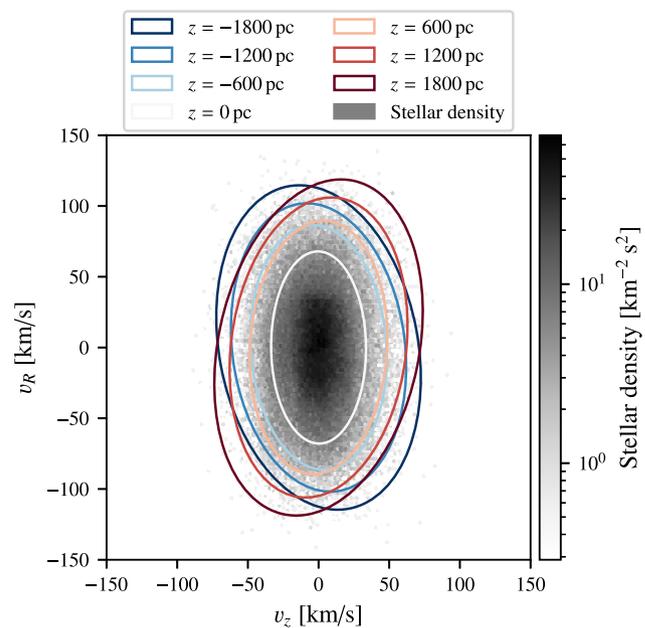


Figure 6. Density map of stellar velocities in the (v_z, v_R) -plane in our data set. The sample-mean of the inferred covariance matrix is indicated by the 2σ ellipse contours for various heights z .

covariance matrix estimation, is very large for the $z = \pm 1800$ pc samples. At these distances, it is only constrained by extrapolation from smaller $|z|$ and our prior structure for the Gaussian processes and angles (cf. equations 10 and 12). However, the rotation away from a diagonal covariance matrix is clearly visible, statistically significant, and consistent with findings of previous analyses such as Bündenbender, van de Ven & Watkins (2015).

The inferred density fall-off $\nu(z)$ is shown in Fig. 7. The model is in agreement with the data, indicated by a good fit of the star counts in the considered region. At distances farther than $|z| = 1200$ pc, so outside the range of star count data considered, our samples show a large spread, reflecting the lack of constraining data. See the discussion in Section 2.3 for an explanation of why data outside this region were excluded. Similar to the velocity moments, we can also see a clear asymmetry in the density fall-off above and below the

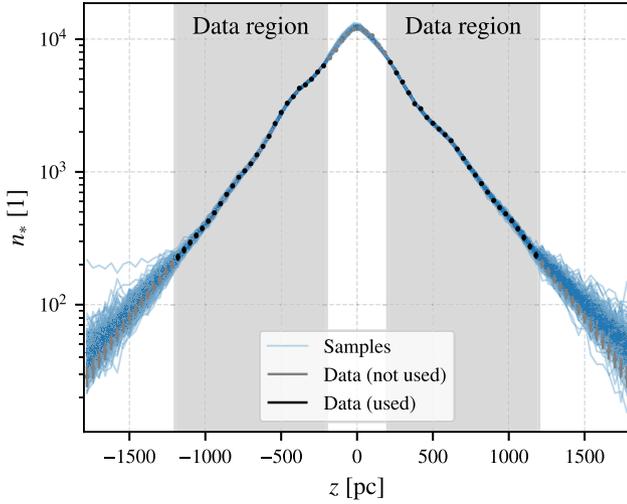


Figure 7. Star counts n_* as a function of height z above the Galactic disc. All samples are shown. The shaded region indicates the availability of star count data.

disc at approximately $z = \pm 400$ pc. Interestingly, these asymmetries fit together very well, being able to explain the asymmetric data even in a symmetric gravitational potential. It should be emphasized that these are asymmetries in the tracer population only, whereas the disc mass model is still fully symmetric about the mid-plane.

We find a local dark matter density of

$$\rho_{\text{dm}} = 0.0117 \pm 0.0035 M_{\odot} \text{pc}^{-3} = 0.44 \pm 0.13 \text{ GeV cm}^{-3}, \quad (20)$$

which already contains the correction term using Oort’s constants. This result uses the Gaussian process prior for R_{T}^{-1} in the tilt term (equation 5). This result is consistent with values typically found in the literature (compare with, e.g. de Salas 2020) for local analyses, but significantly lower than the (mean) value derived by Garbari et al. (2012) using a similar class of stars. As a consistency check, we also ran our analysis using only data above or below the disc, finding dark matter densities of $\rho_{\text{dm}}^{\text{north}} = 0.0149 \pm 0.0045 M_{\odot} \text{pc}^{-3}$ and $\rho_{\text{dm}}^{\text{south}} = 0.0149 \pm 0.0037 M_{\odot} \text{pc}^{-3}$, which agree with both each other and with the value derived using the full data set. The total baryonic surface mass density is prior-dominated, with a value of $\Sigma_{\text{b}} = 43.75 \pm 3.04 M_{\odot} \text{pc}^{-2}$.

In Fig. 8, we compare our result to other recent Jeans analyses of the local dark matter density. Apart from consistency with the literature, it also illustrates that the estimated dark matter density is anticorrelated with both the baryonic surface mass density Σ_{b} and the inverse characteristic length-scale R_{T}^{-1} of the tilt term. Deviations between different Jeans analyses can stem from different choices of the tracer population, the particular data used, and the approximations made in the Jeans equation. In the following text, we want to examine the influence of our modelling of the tilt term on the inferred dark matter density.

In order to quantify the effect of our particular modelling of the tilt term, we also ran an analysis with a scalar prior for R_{T}^{-1} , assuming a constant behaviour in z . We show a comparison of the inferred dark matter densities as a histogram, R_{T}^{-1} as a function of z , and the relative contribution of the tilt term in Fig. 9.

Using the scalar prior, we find a local dark matter density of

$$\rho_{\text{dm}} = 0.0053 \pm 0.0028 M_{\odot} \text{pc}^{-3} = 0.20 \pm 0.11 \text{ GeV cm}^{-3}, \quad (21)$$

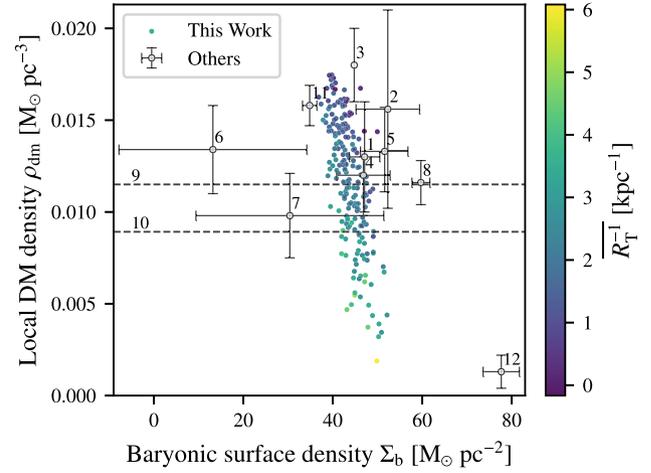


Figure 8. Our estimation of the local dark matter density in comparison to recent estimates using a vertical Jeans analysis. The colour coding indicates the mean value of R_{T}^{-1} of each posterior sample, averaged over the z -axis. The other works, shown in grey, are McKee et al. (2015) (1), Xia et al. (2016) (2), Hagen & Helmi (2018) (3), Sivertsson et al. (2018) (4), Guo et al. (2020) (5), Salomon et al. (2020) (6 north and 7 south), Wardana et al. (2020) (8), Nitschai et al. (2020); Nitschai et al. (2021) (9 and 10), and Cheng et al. (2024) (11 thin disc and 12 thick disc). We have omitted one recent result by López-Corredoira (2025), because the error bars are so large that readability would suffer. For some works, we had to assume a locally constant dark matter density for calculating their baryonic surface mass density.

which is significantly lower than the value found using the Gaussian process prior. Here, we find a difference to the values derived using only data above or below the disc, in particular $\rho_{\text{dm}}^{\text{north}} = 0.0135 \pm 0.0045 M_{\odot} \text{pc}^{-3}$ and $\rho_{\text{dm}}^{\text{south}} = 0.0155 \pm 0.0052 M_{\odot} \text{pc}^{-3}$. Note that even though the difference may seem large, the uncertainties are also large, and the values are still consistent with each other. We attribute this difference to the asymmetry in the data. Fitting both sides of the disc simultaneously demands more flexibility in the model, favouring larger values of R_{T}^{-1} and thus lower values of ρ_{dm} . The baryonic surface mass density is slightly larger than in the Gaussian process prior case, with a value of $\Sigma_{\text{b}} = 46.07 \pm 2.70 M_{\odot} \text{pc}^{-2}$.

The characteristic length-scale R_{T}^{-1} of the tilt term is shown in the middle row of Fig. 9. Somewhat surprisingly, it prefers positive values indicating that the correlated velocity moment $\bar{v}_z \bar{v}_R$ is increasing towards larger R . This could be caused either by an increasing tilt angle, or by larger velocity dispersions towards larger R . The scalar prior prefers values of about $3.82 \pm 0.72 \text{ kpc}^{-1}$, while the Gaussian process prior prefers values of about $2.11 \pm 0.98 \text{ kpc}^{-1}$ after averaging over z . The preferred values of order unity lead to a net positive contribution of the tilt term to the exponential fall-off in equation (6), meaning a steepening of the vertical density fall-off v . Such a general steepening can also be achieved by increasing the local dark matter density ρ_{dm} as this translates to the gravitational force $\frac{\partial \Phi}{\partial z}$ gaining an additional, linear term. This explains why the dark matter density is anticorrelated with the product of $\bar{v}_z \bar{v}_R$ and R_{T}^{-1} (cf. equation 6). Since the allowed range of $\bar{v}_z \bar{v}_R$ is dictated by data that are equal for both priors, the difference in the inferred dark matter density is mainly due to the different values of R_{T}^{-1} . We show the relative contribution of the tilt term in the Jeans equation in the bottom row of Fig. 9. The relative contribution is defined as the ratio of the tilt term to the other term in the Jeans equation (equation 6),

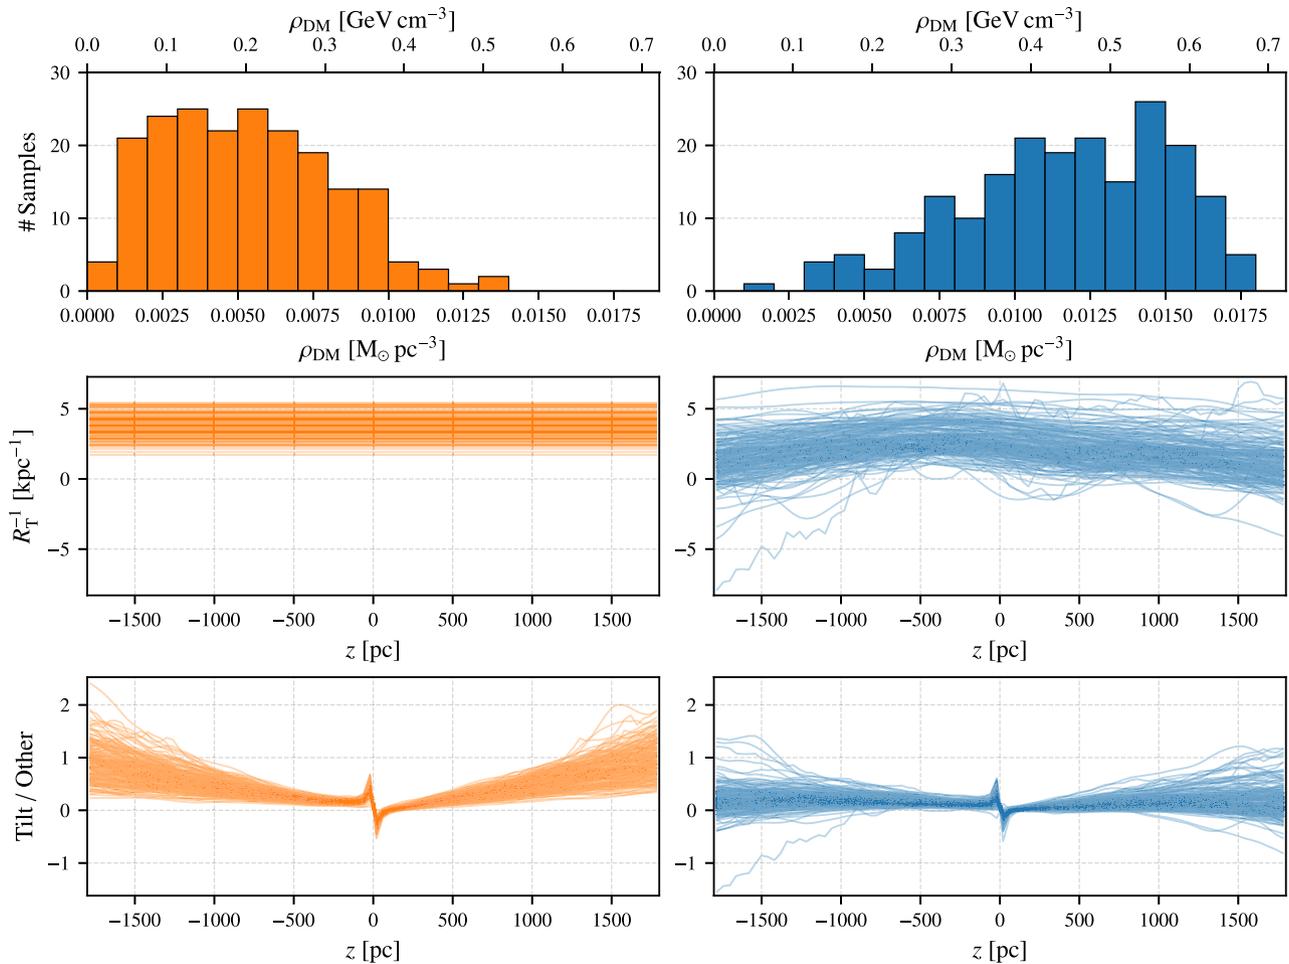


Figure 9. Comparison of the effect of modelling the tilt term on the inferred dark matter density. The left column shows the analysis using a scalar prior for R_T^{-1} , the right column the analysis using a Gaussian process prior. All samples are shown. *Top row:* histogram showing samples of the inferred posterior distribution of the local dark matter density ρ_{dm} . *Middle row:* the characteristic length-scale R_T^{-1} as a function of height z above the Galactic disc. *Bottom row:* the relative contribution of the tilt term in the Jeans equation as a function of height z above the Galactic disc. For details, refer to the text.

in particular

$$\frac{\text{Tilt}}{\text{Other}} = \frac{\overline{v_z v_R} \left(\frac{1}{R} + \frac{1}{R_0} + \frac{1}{R_T} \right)}{\frac{\partial \Phi}{\partial z}}. \quad (22)$$

As both the numerator and denominator change sign near $z = 0$ pc, the sharp peaks and change of sign there are not of interest.

For the Gaussian process prior, the tilt term contributes less than 30 per cent of the other term in the Jeans equation over the whole z region considered for most samples. Especially towards small $|z|$, the tilt term is small, as $\overline{v_z v_R}$ is close to zero there. There is a large spread around the mean, where, for some samples, the tilt term is the dominant contribution at large $|z|$.

For the scalar prior, the tilt term contributes much more on average. It has a net positive contribution for all samples and at all z , except for the very small $|z|$ region. Since these large values of R_T^{-1} are only obtained in the region around $z = -400$ pc for the Gaussian process prior, this region might be responsible for dictating the large value everywhere in case of the scalar prior. This enforces a significant contribution of the tilt term everywhere, increasing its importance as a consequence. The assumption of separability of the radial and vertical direction in the tilt term (equation 5) is not well justified and likely too restrictive. This makes the tilt term important for the Jeans

equation at all z considered, contributing more than 50 per cent of the other term in the Jeans equation for most samples at large $|z|$.

Generally, the tilt term is not well constrained, which is a consequence of sparse velocity data at large z , and the lack of explicit data and modelling. This result is consistent with the findings of Sivertsson et al. (2018) who find a strong tension of prior and posterior for this kind of approximation in one of their stellar sample populations (using G-dwarfs). They make use of a uniform prior for $\frac{1}{R_0} + \frac{1}{R_T}$ restricted to values smaller than those found here. Assuming for a moment that the value we found is physically correct, it is not surprising that they found a strong tension. Whether there really is a radial increase of $\overline{v_z v_R}$, as implied by the positive value of R_T^{-1} , is something only the data can tell. Unfortunately, the current data lack the statistical power to allow a definitive conclusion. It may well be that this approach to approximating the Jeans equations and parametrizing the tilt term is problematic, as it can compensate for broken assumptions (e.g. steady state or axisymmetry) by absorbing their effects into the tilt term. We however find an excellent data fit with our model, enabled mainly by the flexibility of the Gaussian processes.

We conclude that the tilt term is, at least in our approximation using a Gaussian process as prior, subdominant but not negligible, contributing less than approximately 50 per cent relative to the other

term in the Jeans equation at all heights z considered in this analysis and for most samples.

5 CONCLUSIONS

In this work, we present a new analysis of the local dark matter density in the Galactic disc using a sample of stars from the *Gaia* DR3 (Gaia Collaboration 2023) catalogue with available radial velocity measurements extended using the SoS (Tsantaki et al. 2022) catalogue. The analysis was performed using geometric variational inference (geoVI, Frank et al. 2021) implemented in the NIFTY8 package. The velocity moments are modelled as Gaussian processes in the eigenspace of the 2D covariance matrix of radial and vertical velocities. We also model the tilt term in the Jeans equation as a Gaussian process in z , avoiding the need for assuming separability of the radial and vertical components, or a restrictive functional form. We find a local dark matter density of $\rho_{\text{dm}} = 0.0117 \pm 0.0035 M_{\odot} \text{pc}^{-3}$, which is in agreement with previous similar estimates. Comparing this to an analysis using a scalar prior for the tilt term, we demonstrate that the modelling of the tilt term has a significant effect on the inferred local dark matter density and can lead to severe biases.

This work serves not only as an updated estimation of the local dark matter density (which has been performed many times before), but also as a proof of concept for two things: first, the application of Gaussian processes to model the velocity moments of a Galactic tracer population as a function of z , and secondly, the ability to infer non-diagonal elements of a multivariate Gaussian likelihood within this family of variational inference methods. The first point naturally leads to a steep increase in the number of latent parameters, requiring moving away from MCMC-based sampling methods of sampling the posterior. The application of variational inference methods also allows us to include measurement uncertainties in a much more natural way than usually done, which is not by Gaussian error propagation, but by sampling the uncertainties directly in the prior and then transforming them to the desired physical quantities within the forward model. This, too, increases the number of effective latent parameters by a large factor, further manifesting the need for an efficient inference method.

We confirm the smooth, approximately linear, but not featureless nature of the velocity moments as a function of height z . Furthermore, we find a significant correlation of radial and vertical stellar velocities (manifesting in a non-zero off-diagonal term in the covariance matrix), increasing with $|z|$ as well, which is consistent with previous findings (compare with e.g. Büdenbender et al. 2015). Finally, the effect of the tilt term is quantified using a local approximation and found to be subdominant but not negligible in the Jeans equation, contributing less than 50 per cent of the other term in the Jeans equation for all heights considered in this analysis and for most samples. Please note that most samples have a much smaller contribution, especially at small $|z|$.

For improving upon this in future analyses, there are multiple options: first, the mass model of the baryonic components can be improved by allowing for more freedom, especially an asymmetry in the mass distribution above and below the disc. The fact that there is an apparent asymmetry in the tracer population (here, and also in previous analyses using different families of tracers) implies the possibility of an asymmetry in the overall mass distribution, especially for large $|z|$. Secondly, the spatial dimensionality of the problem can be increased in various ways. The assumption of axisymmetry can be removed by adding the angular velocity as a third component to the velocity covariance. This change would likely go hand in hand with an increase of spatial dimensionality,

i.e. making the velocity moments and stellar density not only a function of height z , but also of the galactocentric radius R and azimuthal angle φ . This would give direct access to the required spatial derivatives. As a third, and possibly the most difficult, avenue for improvement, the time axis can be considered. Essentially all past analyses, apart from simulations, assume a steady-state system. This assumption is being increasingly questioned as evidence for disequilibria, excited, for example, by spiral arms, stellar streams, or the Large Magellanic Cloud, is mounting (e.g. Widrow et al. 2012; Hou & Han 2014; Antoja et al. 2018; Laporte et al. 2018; Vasiliev, Belokurov & Erkal 2021, among many others). All of these extensions to the model would lead to a significant increase in model complexity and computational requirements. Modelling using Gaussian processes also has some downsides, especially when simultaneously neglecting relevant parts of the model (e.g. the time dependence), since the freedom of the Gaussian processes can absorb effects that belong in a different unmodelled field. Thus, while we strongly recommend their use, we also recommend caution when interpreting the results.

We conclude by stressing that the methods used here are very versatile and can be applied to a wide range of problems, even beyond Galactic astrophysics. The combination of sophisticated inference methods and the ever-increasing amount of large high-quality data sets may soon lead to a new era of data-driven high-dimensional modelling in astrophysics.

ACKNOWLEDGEMENTS

We would like to thank João Alves for helpful comments regarding *Gaia* data. All figures in this publication have been created using MATPLOTLIB (Hunter 2007). Part of this work was supported by the German Deutsche Forschungsgemeinschaft, DFG project number 495252601. The authors gratefully acknowledge the computing time provided to them at the NHR Center NHR4CES at RWTH Aachen University (project number p0021785). This is funded by the Federal Ministry of Education and Research, and the state governments participating on the basis of the resolutions of the GWK for national high performance computing at universities (<https://www.nhr-verein.de/unsere-partner>). Part of this work is based on archival data, software, or online services provided by the Space Science Data Center – ASI. In particular, the GaiaPortal access tool was used for this research (<http://gaiaportal.ssd.c.asi.it>).

DATA AVAILABILITY

The *Gaia* DR3 data have been downloaded using the python package *astroquery* (Ginsburg et al. 2019). The Survey-of-Surveys (SoS) data were downloaded from the GaiaPortal website (<https://gaiaportal.ssd.c.asi.it/SoS/query/form>). The processed data sets and posterior samples will be shared on reasonable request to the corresponding author.

REFERENCES

- Ablimit I., Zhao G., Flynn C., Bird S. A., 2020, *ApJ*, 895, L12
- Akhmetov V. S., Bucciarelli B., Crosta M., Lattanzi M. G., Spagna A., Re Fiorentin P., Bannikova E. Y., 2024, *MNRAS*, 530, 710
- Antoja T. et al., 2018, *Nature*, 561, 360
- Arras P. et al., 2019, Astrophysics Source Code Library, record ascl:1903.008
- Arras P., Frank P., Haim P., Knollmüller J., Leike R., Reinecke M., Enßlin T., 2022, *Nat. Astron.*, 6, 259
- Bahcall J. N., 1984, *ApJ*, 276, 169

- Benito M., Cuoco A., Iocco F., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 033
- Benito M., Iocco F., Cuoco A., 2021, *Phys. Dark Universe*, 32, 100826
- Bienaymé O. et al., 2014, *A&A*, 571, A92
- Binney J., Piffl T., 2015, *MNRAS*, 454, 3653
- Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*. Princeton Univ. Press, Princeton, NJ
- Binney J. et al., 2014, *MNRAS*, 439, 1231
- Buch J., Leung J. S. C., Fan J., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 026
- Büdenbender A., van de Ven G., Watkins L. L., 2015, *MNRAS*, 452, 956
- Buder S. et al., 2018, *MNRAS*, 478, 4513
- Cautun M. et al., 2020, *MNRAS*, 494, 4291
- Cheng X., Anguiano B., Majewski S. R., Arras P., 2024, *MNRAS*, 527, 959
- Cole D. R., Binney J., 2017, *MNRAS*, 465, 798
- de Salas P. F., 2020, *J. Phys. Conf. Ser.*, 1468, 012020
- de Salas P. F., Widmark A., 2021, *Rep. Prog. Phys.*, 84, 104901
- de Salas P. F., Malhan K., Freese K., Hattori K., Valluri M., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 037
- De Silva G. M. et al., 2015, *MNRAS*, 449, 2604
- Dehnen W., Binney J. J., 1998, *MNRAS*, 298, 387
- Edenhofer G. et al., 2024a, *J. Open Source Softw.*, 9, 6593
- Edenhofer G., Zucker C., Frank P., Saydjari A. K., Speagle J. S., Finkbeiner D., Enßlin T. A., 2024b, *A&A*, 685, A82
- Fabricius C. et al., 2021, *A&A*, 649, A5
- Flynn C., Holmberg J., Portinari L., Fuchs B., Jahreiß H., 2006, *MNRAS*, 372, 1149
- Frank P., Leike R., Enßlin T. A., 2021, *Entropy*, 23, 853
- Gaia Collaboration, 2023, *A&A*, 674, A1
- Garbari S., Read J. I., Lake G., 2011, *MNRAS*, 416, 2318
- Garbari S., Liu C., Read J. I., Lake G., 2012, *MNRAS*, 425, 1445
- Gilmore G. et al., 2012, *The Messenger*, 147, 25
- Ginsburg A. et al., 2019, *AJ*, 157, 98
- Guo R., Liu C., Mao S., Xue X.-X., Long R. J., Zhang L., 2020, *MNRAS*, 495, 4828
- Guo R., Shen J., Li Z.-Y., Liu C., Mao S., 2022, *ApJ*, 936, 103
- Hagen J. H. J., Helmi A., 2018, *A&A*, 615, A99
- Hattori K., Valluri M., Vasiliev E., 2021, *MNRAS*, 508, 5468
- Holmberg J., Flynn C., 2000, *MNRAS*, 313, 209
- Hou L. G., Han J. L., 2014, *A&A*, 569, A125
- Huang Y. et al., 2016, *MNRAS*, 463, 2623
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Hutschenreuter S., Haverkorn M., Frank P., Raycheva N. C., Enßlin T. A., 2024, *A&A*, 690, A314
- Kafle P. R., Sharma S., Lewis G. F., Bland-Hawthorn J., 2014, *ApJ*, 794, 59
- Kapteyn J. C., 1922, *ApJ*, 55, 302
- Karukes E. V., Benito M., Iocco F., Trotta R., Geringer-Sameth A., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 046
- Knollmüller J., Enßlin T. A., 2019, preprint (arXiv:1901.11033)
- Laporte C. F. P., Johnston K. V., Gómez F. A., Garavito-Camargo N., Besla G., 2018, *MNRAS*, 481, 286
- Li Z., Han J., Wang W., Qian Y.-Z., Li Q., Jing Y., Li T. S., 2025, *MNRAS*, 538, 1442
- Lim S. H., Putney E., Buckley M. R., Shih D., 2025, *J. Cosmol. Astropart. Phys.*, 2025, 021
- Lin H.-N., Li X., 2019, *MNRAS*, 487, 5679
- Lindegren L. et al., 2021, *A&A*, 649, A4
- López-Corrodoira M., 2025, *ApJ*, 978, 45
- Majewski S. R. et al., 2017, *AJ*, 154, 94
- McKee C. F., Parravano A., Hollenbach D. J., 2015, *ApJ*, 814, 13
- McMillan P. J., 2017, *MNRAS*, 465, 76
- Nesti F., Salucci P., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 016
- Nitschai M. S., Cappellari M., Neumayer N., 2020, *MNRAS*, 494, 6001
- Nitschai M. S., Eilers A.-C., Neumayer N., Cappellari M., Rix H.-W., 2021, *ApJ*, 916, 112
- Oort J. H., 1932, *Bull. Astron. Inst. Netherlands*, 6, 249
- Ou X., Eilers A.-C., Necib L., Frebel A., 2024, *MNRAS*, 528, 693
- Pato M., Iocco F., Bertone G., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 001
- Piffl T. et al., 2014, *MNRAS*, 445, 3133
- Randich S., Gilmore G., Consortium Gaia-ESO, 2013, *The Messenger*, 154, 47
- Rehmentulla N., Valluri M., Vasiliev E., 2022, *MNRAS*, 511, 5536
- Reid M. J. et al., 2019, *ApJ*, 885, 131
- Robin A. C., Reylé C., Derrière S., Picaud S., 2003, *A&A*, 409, 523
- Robin A. C. et al., 2022, *A&A*, 667, A98
- Roth J., Arras P., Reinecke M., Perley R. A., Westermann R., Enßlin T. A., 2023, *A&A*, 678, A177
- Salomon J.-B., Bienaymé O., Reylé C., Robin A. C., Famaey B., 2020, *A&A*, 643, A75
- Salucci P., Nesti F., Gentile G., Frigerio Martins C., 2010, *A&A*, 523, A83
- Schönrich R., Binney J., Dehnen W., 2010, *MNRAS*, 403, 1829
- Schutz K., Lin T., Safdi B. R., Wu C.-L., 2018, *Phys. Rev. Lett.*, 121, 081101
- Selig M., Bell M. R., Junklewitz H., Oppermann N., Reinecke M., Greiner M., Pachajoa C., Enßlin T. A., 2013, *A&A*, 554, A26
- Sivertsson S., Silverwood H., Read J. I., Bertone G., Steger P., 2018, *MNRAS*, 478, 1677
- Sivertsson S. et al., 2022, *MNRAS*, 511, 1977
- Skowron D. M. et al., 2019, *Science*, 365, 478
- Söding L. et al., 2025, *A&A*, 693, A139
- Sofue Y., 2020, *Galaxies*, 8, 37
- Steininger T. et al., 2019, *Ann. Phys.*, 531, 1800290
- Steinmetz M. et al., 2006, *AJ*, 132, 1645
- Steinmetz M. et al., 2020, *AJ*, 160, 82
- Tsantaki M. et al., 2022, *A&A*, 659, A95
- Tsouros A., Edenhofer G., Enßlin T., Mastorakis M., Pavlidou V., 2024, *A&A*, 681, A111
- Vasiliev E., Belokurov V., Erkal D., 2021, *MNRAS*, 501, 2279
- Wang W., Han J., Cole S., Frenk C., Sawala T., 2017, *MNRAS*, 470, 2351
- Wardana M. D., Wulandari H., Sulistiyowati, Khatami A. H., 2020, in *European Physical Journal Web of Conferences*. EDP Sciences, France, p. 04002
- Wegg C., Gerhard O., Bieth M., 2019, *MNRAS*, 485, 3296
- Widmark A., Laporte C. F. P., de Salas P. F., Monari G., 2021, *A&A*, 653, A86
- Widrow L. M., Gardner S., Yanny B., Dodelson S., Chen H.-Y., 2012, *ApJ*, 750, L41
- Xia Q. et al., 2016, *MNRAS*, 458, 3839
- Zhao G., Zhao Y.-H., Chu Y.-Q., Jing Y.-P., Deng L.-C., 2012, *Res. Astron. Astrophys.*, 12, 723
- Zhou Y., Li X., Huang Y., Zhang H., 2023, *ApJ*, 946, 73

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