

A simulation-based inference of the Milky Way merger history

Andrea Sante^{1,★}, Andreea S. Font¹, Daisuke Kawata^{2,3}, T. Lucas Makinen^{4,5}
and Robert J. J. Grand¹

¹*Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK*

²*Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK*

³*National Astronomical Observatory of Japan, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan*

⁴*Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK*

⁵*Imperial Centre for Inference and Cosmology (ICIC), Imperial College London, Prince Consort Road, London SW7 2AZ, UK*

Accepted 2026 June 12. Received 2026 June 7; in original form 2026 March 12

ABSTRACT

Accreted stars in the Milky Way (MW) preserve information about their galaxies of origin in their chemo-dynamical properties. In this study, we use the chemo-dynamical signatures in the merger debris to approximate the posterior distribution of the properties of disrupted satellites at their accretion. Adopting a simulation-based inference framework, we train an ensemble of normalizing flows using samples of merger debris from the Auriga simulations. Applying this methodology to a local sample of accreted stars in the MW, we infer the lookback times, stellar and halo masses, and halo mass merger ratios of several known accretion events: *Gaia*-Enceladus-Sausage (GES), Helmi streams, Heracles, I'toi, LMS-1/Wukong, Sagittarius, Sequoia, and Thamnos. Our predictions for accretion time and mass align with previous estimates and follow the expected stellar mass–metallicity relation across redshifts. The total accreted stellar mass from these events is predicted to be $2.2^{+1.1}_{-0.6} \times 10^9 M_{\odot}$, with the GES and Sagittarius dwarfs being the largest contributors. The total accreted stellar mass from all disrupted progenitors is $1.4^{+1.1}_{-0.4} \times 10^9 M_{\odot}$, which is consistent with previous mass measurements of this component. We provide a prediction for the evolution of the MW halo mass until the accretion of Sagittarius ($z \approx 1$): We find that the halo mass growth from the time of its first merger ($z \approx 5$) to $z \approx 2$ significantly exceeds the total mass of the known progenitors accreted during that interval, suggesting the presence of unidentified substructures. Our estimate of the Galaxy halo mass after the Sagittarius merger, but prior to the accretion of the Magellanic Clouds, is $5.9^{+1.4}_{-1.1} \times 10^{11} M_{\odot}$.

Key words: software: machine learning – software: simulations – Galaxy: halo – Galaxy: kinematics and dynamics – Galaxy: stellar content.

1 INTRODUCTION

Within the hierarchical structure formation scenario, the Milky Way (MW) assembled its mass by accreting smaller satellite galaxies over time (S. D. M. White & M. J. Rees 1978). During these cosmic collisions, stars are tidally disrupted from the progenitor systems in which they formed, creating the stellar halo of the Galaxy (L. Searle & R. Zinn 1978; K. V. Johnston, L. Hernquist & M. Bolte 1996). Hence, this extended and diffuse component serves as a repository of fossil records of the events that describe the assembly history of the Galaxy (A. Helmi 2020; A. J. Deason & V. Belokurov 2024; A. Bonaca & A. M. Price-Whelan 2025).

The most apparent remnant of these merging events is the tidal tail of the Sagittarius dwarf galaxy (R. A. Ibata, G. Gilmore & M. J. Irwin 1994), which appears as a stream of spatially coherent stars wrapping around a vast region in the Galaxy. Stellar streams are usually associated with recent accretion events, whose debris

have not yet dynamically mixed with the stars that formed *in situ*, i.e. within the proto-MW. Signatures of more ancient accretion events can still be recovered as phase-mixed accreted stars share similar orbital properties in the *integrals of motion* space (A. Helmi & S. D. M. White 1999), which is defined by quantities expected to remain constant over time in a static gravitational potential. Using this principle, A. Helmi et al. (1999) discovered the first evidence of a phase-mixed accretion event – the Helmi streams – which appeared as a substructure in the velocity space of stars in the solar neighbourhood.

Because of the high precision required to separate stellar streams and accreted substructures from the dense background of the MW field stars, it was not until the second data release of the European Space Agency *Gaia* mission (Gaia Collaboration 2018) that it became possible to constrain the properties of the progenitors of the Galaxy. By providing accurate astrometry and proper motions for about a billion stars, fossil remnants of a plethora of merging events were uncovered from the stars in the inner halo (for reviews, see A. J. Deason & V. Belokurov 2024; A. Bonaca & A. M. Price-Whelan 2025). The most significant

* E-mail: A.Sante@2022.ljmu.ac.uk

detection was the identification of a large population of stars in highly radial and eccentric orbits, which represents the remnant of the *Gaia*-Enceladus-Sausage (GES) galaxy (V. Belokurov et al. 2018; A. Helmi et al. 2018), a massive dwarf galaxy believed to have been accreted around 8–10 billion years ago.

While the debris of the GES accretion event are expected to dominate the accreted stars population in the stellar halo (R. P. Naidu et al. 2020), several smaller phase-mixed substructures have also been identified by complementing the *Gaia* catalogue with radial velocity and chemical abundance measurements from high-resolution spectroscopic surveys, such as the Apache Point Observatory Galactic Evolution Experiment (APOGEE; S. R. Majewski et al. 2017), the Galactic Archaeology with HERMES (S. Buder et al. 2021), the Hectochelle in the Halo at High Resolution Survey (C. Conroy et al. 2019), and the Large sky Area Multi-Object fibre Spectroscopic Telescope (X.-Q. Cui et al. 2012). The application of clustering algorithms in this chemo-dynamical space has led to a proliferation of identified substructures, which include the retrograde Sequoia progenitor (G. C. Myeong et al. 2019), the low-energy retrograde structure Thamnos (H. H. Koppelman et al. 2019b), and the prograde LMS-1/Wukong stream (Z. Yuan et al. 2020). A further division of the retrograde halo into multiple populations with different metallicity has been proposed by R. P. Naidu et al. (2020), which includes the presence of the Arjuna and I'toi substructures, which are, respectively, more metal rich and metal poor than the Sequoia debris.

Eventually, the search for the progenitors of the MW has extended into the inner regions of the Galaxy, where the debris of the most ancient mergers are expected to orbit. By analysing the properties of globular clusters and field stars, several groups have postulated the existence of one or more massive, very early mergers (J. M. D. Kruijssen et al. 2020; D. Horta et al. 2021; K. Malhan & H.-W. Rix 2024). These events would have occurred more than 10–12 billion years ago and contributed to the formation of the proto-Galaxy. However, these debris are extremely phase-mixed and chemically similar to the primordial *in-situ* population (V. Belokurov & A. Kravtsov 2023) making their definitive identification an ongoing challenge.

Similarly, as the census of substructures in the stellar halo grows, so does the complexity of their interpretation. For instance, D. Horta et al. (2023) analysed the detailed chemical abundance patterns of stars in these substructures and suggested that some of the retrograde systems, including Arjuna, I'toi, and some definitions of Sequoia, are chemically indistinguishable from GES. This raises the possibility that, rather than being separate accretion events, these may be different dynamical components or fragments of a single, massive, and complex merger. Moreover, recent works have shown that stellar debris from a single progenitor can occupy a vast region of the chemo-dynamical space (I. Jean-Baptiste et al. 2017; A. Mori et al. 2024) and clustering algorithms might return spurious or heavily contaminated substructures if not optimized (G. F. Thomas et al. 2025; A. Sante et al. 2026).

While efforts are made to refine the identification of the individual progenitors of the MW, it is crucial to develop techniques for the characterization of the merger events in order to piece together the MW assembly. Traditionally, the lookback time and mass of the progenitors at infall are estimated by fitting the properties of the respective debris with analytical models. For instance, the infall time of a disrupted galaxy can be inferred from the distribution of the stellar ages of its debris (C. Gallart

et al. 2019; J. Montalbán et al. 2021; T. Woody et al. 2025) as the star formation history is affected by the merger via the tidal stripping of cold gas. The stellar mass of progenitors can instead be determined by fitting the density profile of their debris (J. T. Mackereth & J. Bovy 2020; J. J. Han et al. 2022; J. M. M. Lane, J. Bovy & J. T. Mackereth 2023), or comparing the observed abundance of different chemical elements with the outputs of chemical evolution models (F. Vincenzo et al. 2019; S. Hasselquist et al. 2021). Alternatively, the properties of the progenitor galaxies at accretion can be estimated by comparing the observed chemo-dynamical distribution of their debris with those expected from specific mergers in cosmological simulations.

While these methods provide interpretable estimates of the properties of progenitors, they are often fitted for single substructures following analytical assumptions (i.e. the isochrone model library for colour-magnitude diagram (CMD) fitting, the shape of the distribution function of the debris for the density profile fitting, or the choice of the chemical evolution model) that can be decoupled from the cosmological context in which the assembly occurs. Current cosmological simulations of MW-like galaxies provide a more comprehensive framework for the characterization of mergers, as they self-consistently incorporate a wide range of coupled physical processes on both interstellar and cosmological scales.

One of the first studies to use the information encoded within cosmological simulations to reconstruct the assembly history of the MW was conducted by J. M. D. Kruijssen et al. (2020), who applied an artificial neural network to effectively map globular cluster properties (age and metallicity) in the E-MOSAICS suite (J. M. D. Kruijssen et al. 2019) on to the orbital parameters of the dwarf galaxies from which they originated. This methodology, combined with results from clustering algorithms, was then applied to the globular cluster population in the MW to recover single-point estimates for the time of infall and stellar mass of the Galaxy progenitors. By treating the reconstruction of the merger tree as a regression task, potential degeneracies in the merger parameters were not accounted for in these predictions, which only reflected the uncertainty due to the configuration of the neural network.

In contrast, physical degeneracies can be explicitly accounted for by casting predictions of merger properties as a probabilistic inference problem using the mathematical framework of simulation-based inference (SBI; K. Cranmer, J. Brehmer & G. Louppe 2020; M. Ho et al. 2024). SBI adopts generative machine learning techniques to directly approximate the likelihood or posterior distribution of intractable or unknown models from simulations. The main idea is to perform a Bayesian analysis by obtaining an approximation of the posterior distribution as informed by some observed realizations of the simulated model. Samples can then be drawn from this distribution effectively determining the most likely set of model parameters associated with the observations.

One of the first applications of SBI to the reconstruction of the properties of disrupted mergers is the GalactiKit methodology (A. Sante et al. 2025). By using the information contained in the merger histories from the 39 cosmological magnetohydrodynamical (MHD) zoom-in simulations of MW-like galaxies in the Auriga suite (R. J. J. Grand et al. 2017, 2024), a model was developed to infer the lookback time, stellar mass, halo mass, and halo mass merger ratio at infall of disrupted dwarf galaxies directly from the present-day chemo-dynamical distribution of their debris.

The aim of this paper is to apply the GalactiKit methodology to the accreted stars in the MW in order to provide a data-driven, self-consistent reconstruction of the assembly history of the Galaxy as depicted by the current census of its disrupted progenitors. For this analysis, we follow the data construction and division of stars in their respective accreted progenitor of origin defined in D. Horta et al. (2023). More details on the observational data used in this analysis are provided in Section 2. While the Auriga simulations reproduce morphological, chemical, and dynamical trends observed in the stellar haloes of MW-mass galaxies (A. Monachesi et al. 2019), there are still numerical offsets between the observed chemo-dynamical properties of stars and the ones of the particles in the simulations. In order to account for these discrepancies, we introduce a number of changes to the GalactiKit methodology, which are described in detail in Section 3. Finally, the chronology and properties of the merger events characterizing the formation history of the MW as inferred through GalactiKit are reported in Section 4. Section 5 presents our conclusions.

2 DATA

The chemo-dynamical properties for the stars considered in this analysis are obtained from a combination of the SDSS-IV/APOGEE (DR17; S. R. Majewski et al. 2017; Abdurro'uf et al. 2022) and *Gaia* (EDR3; Gaia Collaboration 2021) surveys. The celestial coordinates and proper motions provided by *Gaia* are complemented by the radial velocities and chemical abundance measurements from APOGEE. The distance to each star is taken from the *astroNN* catalogue produced by H. W. Leung & J. Bovy (2019), as it accounts for parallax zero-point biases in *Gaia*-EDR3 using photometric distance estimates from an artificial neural network model trained on APOGEE spectra. The 6D phase-space coordinates are converted to a Galactocentric cylindrical coordinate system adopting the following solar parameters: a distance to the Galactic centre of $R_0 = 8.178$ kpc (GRAVITY Collaboration 2019), a vertical height above the Galactic mid-plane of $Z_0 = 0.02$ kpc (M. Bennett & J. Bovy 2019), and a velocity vector $[U_\odot, V_\odot, W_\odot] = [-11.1, 248.0, 8.5]$ km s⁻¹, which accounts for the motion of the Sun relative to the Local Standard of Rest and Sgr A* (R. Schönrich, J. Binney & W. Dehnen 2010; M. J. Reid & A. Brunthaler 2020). The total energy of the stars is calculated assuming the P. J. McMillan (2017) model of the gravitational potential of the MW.

To ensure that the properties of the stars considered in the analysis are reliably measured and the samples of accreted debris of the disrupted progenitors are chemically distinct from each other and consistent with the dwarf satellite population of the MW, we define the accreted substructures in the Galaxy following the procedure described in D. Horta et al. (2023). First, the following selection criteria are applied to the combined APOGEE-*Gaia* catalogue to identify the stars with the most reliable chemical abundance measurements:

- (i) $3500 < T_{\text{eff}} < 5500$ K and $\log g < 3.6$,
- (ii) APOGEE spectral S/N > 70 ,
- (iii) APOGEE STARFLAG = 0,
- (iv) $d_{\odot, \text{err}}/d_\odot < 0.2$,

where T_{eff} and $\log g$ are stellar atmospheric parameters determined through the APOGEE Stellar Parameter and Chemical Abundances Pipeline (ASPCAP; A. E. García Pérez et al. 2016); APOGEE STARFLAG = 0 ensures that stars with issues in the




Figure 1. Distribution of the orbital properties (E and L_z) of the accreted stars in the solar neighbourhood of the MW. The accreted stars are grouped by the proposed progenitor galaxy in which they formed, according to the criteria of D. Horta et al. (2023), and are plotted on top of the overall distribution of stars in the APOGEE sample from which they were selected.

ASPCAP fit are discarded, while $d_{\odot, \text{err}}$ and $d_\odot$ are the relative uncertainty and value of the *astroNN*-derived distances to the Sun. Stars in the APOGEE Value Added Catalogue of globular cluster stars (R. P. Schiavon et al. 2024) and those associated with the Large and Small Magellanic Clouds samples from S. Hasselquist et al. (2021) are also excluded from the analysis. The resulting parent sample comprises 198 503 stars.

Then, the debris from the disrupted progenitors of the MW are identified following the chemo-dynamical selection cuts described in table 1 of D. Horta et al. (2023). However, we exclude from the analysis the stars from the Aleph and Nyx substructures as they are chemically consistent with the low- α and high- α disc populations, respectively. We also do not include stars from the Arjuna substructure, given their potential overlap with the high-energy retrograde debris from the GES merger (R. P. Naidu et al. 2021; D. Horta et al. 2023; T. Woody et al. 2025). Moreover, we do not consider the Icarus and Pontus substructures because they are comprised of a statistically insignificant number of stars in the parent sample (< 5). The accreted substructures (and their respective number of stars in their sample) considered in this analysis are as follows: GES (2380), the Helmi streams (104), Heracles (263), I'toi (24), LMS-1/Wukong (89), Sagittarius (116), and Sequoia and Thamnos (234). The Sequoia sample is defined in three different ways following the selection criteria adopted in previous works by H. H. Koppelman et al. (2019a), G. C. Myeong et al. (2019), and R. P. Naidu et al. (2020). We follow the notation of D. Horta et al. (2023) and refer to the samples selected through these selection criteria as *K19* (41), *M19* (165), and *N20* (77), respectively.

The chemo-dynamical properties of the stars in the accreted substructures considered in this study are reported in Figs 1 and 2, which show the distribution of debris in the integrals-of-motion and α -iron planes, respectively. In each plot, stars belonging to the same substructures share the same colour and symbol.




Figure 2. Chemical abundance ratios ($[\text{Mg}/\text{Fe}]$ and $[\text{Fe}/\text{H}]$) distribution for the same samples of stars plotted in Fig. 1.

3 INFERRING MERGER PROPERTIES WITH SIMULATION-BASED INFERENCE

The primary objective of this work is to reconstruct the MW merger history by estimating the properties of the past merger events. We approach this as a Bayesian inference problem, aiming to compute the posterior probability distribution for a set of progenitor parameters, θ , of an unknown model, $\mathcal{M}$, which determines the present-day chemo-dynamical properties, $\mathbf{x}$, of its stellar debris, i.e. $\mathcal{M} : \theta \rightarrow \mathbf{x}$.

An analytical definition of the posterior, $p(\theta|\mathbf{x}) \propto p(\theta)p(\mathbf{x}|\theta)$, is intractable as it would require an explicit formulation of the likelihood probability, $p(\mathbf{x}|\theta)$, which accurately encapsulates the complex coupled physical processes governing galaxy formation. These include the hierarchical assembly of galaxies within a cosmological framework, the distinct chemical evolution histories of the host and satellite, and the subsequent phase-mixing of accreted material within the host's evolving gravitational potential.

However, modern cosmological simulations, such as the Auriga suite (R. J. J. Grand et al. 2017, 2024), effectively forward-model these processes, providing a rich data set of realistic merger events. As these simulations connect progenitor properties at the time of accretion to the corresponding stellar debris properties at redshift $z = 0$, they are an implicit representation of the complex physical model, $\mathcal{M}$, that we want to approximate.

The Auriga simulations are a suite of 39 MHD zoom-in simulations of MW mass, relatively isolated, haloes selected from the L100N1504 EAGLE dark matter only simulation (J. Schaye et al. 2015). The zoom-in simulations are run with the moving-mesh code AREPO (V. Springel 2010) and include a comprehensive galaxy formation physics model, which is described in detail

in R. J. J. Grand et al. (2017). Star particles are treated as single stellar populations following a Chabrier (2003) initial mass function. Star formation is regulated by a pressurized two-phase model for the interstellar medium as gas cells become denser than 0.13 particles per cc. The chemical evolution is tracked by accounting for enrichment from asymptotic giant branch stars and both Type Ia and Type II supernovae. Feedback processes are also included as galactic-scale gaseous outflows powered by supernovae winds and active galactic nuclei feedback in both quasar and radio modes. Additionally, the model includes primordial and metal-line gas cooling with self-shielding corrections, as well as a unidirectional, homogeneous magnetic field of 10^{-14} G.

The combination of these physical prescriptions results in a population of simulated galaxies that successfully reproduce a wide range of observational scaling relations found in the stellar haloes of nearby MW analogues (A. Monachesi et al. 2019).

We extract the information in the Auriga simulation with SBI (K. Cranmer et al. 2020; M. Ho et al. 2024), a class of statistical methods that uses simulations to perform Bayesian inference when the likelihood is unknown or computationally prohibitive. In particular, we base our approach on GalactiKit, an SBI framework developed by A. Sante et al. (2026) specifically to reconstruct the merger properties. The core idea of GalactiKit is to use an ensemble of Masked Autoregressive Flows (MAFs; G. Papamakarios, T. Pavlakou & I. Murray 2017), a type of generative artificial neural network capable of learning complex high-dimensional probability distributions, to approximate the posterior distribution of the merger parameters directly from the joint distribution of the merger–debris pairs contained in the Auriga simulations.

Following the GalactiKit framework, we define the vector of merger parameters θ , composed of the quantities of interest to be inferred, i.e. the lookback infall time τ , stellar mass $\log(M_*/M_\odot)$, halo mass $\log(M/M_\odot)$, and the halo mass merger ratio $\log(\text{MMR})$ at infall. The data vector $\mathbf{x}$, which is used to inform the prediction of the posterior distribution, is defined as the chemo-dynamical properties of 100 star particles randomly selected from the present-day ($z = 0$) debris of the merger. For each star particle, we consider the total energy $[\log(E)]$, total specific angular momentum $[\log(L)]$, iron-to-hydrogen abundance $[\text{Fe}/\text{H}]$, and α -to-iron abundance $[\alpha/\text{Fe}]$. From each merger event in the Auriga simulations, we draw multiple 100-star samples to augment the data set, with a maximum of 10 samples per event to ensure that individual mergers with vast amounts of debris do not dominate the training process. In total, this sampling strategy yields 20 433 unique merger–debris pairs, 80 per cent of which are randomly selected to train the posterior estimator and the rest used to validate its performance.

The original GalactiKit framework was designed for application to simulated data. Therefore, the methodology needs to be adapted to account for potential discrepancies between the simulations and observations before being applied to real observations of the MW. These discrepancies can arise from a wide range of sources, such as the systematic differences between the subgrid physics models in Auriga and the true physical processes of the MW (e.g. chemical yield tables, feedback models) or the uncertainties inherent in the measurement of stellar properties.

3.1 A calibration model for generating realistic training data

Because the real deviations between the stellar properties in the simulations and observations, $\Delta\mathbf{x}$, are unknown, we assume that they can be described by an underlying stochastic model, $\mathcal{C} : \theta_c \rightarrow \Delta\mathbf{x}$.

By perturbing the same debris sample with offset values drawn from $\mathcal{C}$, i.e. $\hat{\mathbf{x}} = \mathbf{x} + \Delta\mathbf{x}$, a new training data set of merger–debris pairs, $(\theta, \hat{\mathbf{x}})$, can be created. As each debris configuration, $\mathbf{x}$, is represented by multiple realizations with different calibration values, $\hat{\mathbf{x}}$, a posterior density estimator trained on this new data set learns to estimate the posterior distribution, $p(\theta|\hat{\mathbf{x}})$, of the merger parameters marginalizing over the differences between simulations and observations.

A schematic of the relations between the merger–debris model, $\mathcal{M}$, and the calibration model, $\mathcal{C}$, is shown in Fig 3. The actual $z = 0$ chemo-dynamical distribution of debris $\mathbf{x}$ is parametrized by the properties of the associated disrupted progenitor galaxy, θ , while the observed distribution, $\hat{\mathbf{x}}$, is perturbed by the stochastic noise component, $\Delta\mathbf{x}$. This is modelled as a multivariate Gaussian distribution with mean and standard deviation vectors, $\mathbf{M}$ and Σ . For each $\Delta\mathbf{x}$, the values of the i th components in $\mathbf{M}$ are drawn by a normal distribution, $\mathcal{N}_i(\mu_{M_i}, \sigma_{M_i}^2)$. The parameters $(\mu_{M_i}, \sigma_{M_i}^2)$ are chosen to reflect the degree of discrepancy between the simulated and observed debris properties. For instance, the gravitational potential and Solar motion are assumed to be reasonably well modelled. Thus, the parameters of the calibration distribution for these properties are drawn from a Gaussian centred at zero, with a standard deviation of 0.10 to account for minor model inaccuracies and observational error, i.e. $\mu_{\log(E)} = \mu_{\log(L)} = 0$ and $\sigma_{\log(E)}^2 = \sigma_{\log(L)}^2 = 0.10$.

A more significant difference between the simulations and observations is noticed for the chemical abundance ratio




Figure 3. Relations between data, $\mathbf{x}$, and parameters, θ for the model, $\mathcal{M}$, describing the relation between merger and debris properties, and the calibration model, $\mathcal{C}$, describing the difference between simulated and observed stellar properties. The parameters referring to each model are indicated by dashed boxes.

distributions. Comparing the overall $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ distributions for the star particles in each of the 39 Auriga simulations to the ones for the stars in the MW, we find the former are systematically less α -enhanced and more iron-rich than the latter. This could be related to several factors related to the subgrid physics implemented in the Auriga simulations, such as inaccuracies in the chemical yields tables or secondary processes unaccounted for by the stellar evolution or feedback models used in the simulations. To approximately quantify the effect of these discrepancies, we calculate the difference between the peak in the chemical abundance ratio distributions of each Auriga simulation and the MW, $[\text{Fe}/\text{H}]_{\text{MW}} - [\text{Fe}/\text{H}]_{\text{Auriga},i}$, where i indicates one of the 39 simulations considered in the analysis. The mean and variance values for the calibration model are then derived from the median and variance of the set of these differences, i.e. $\mu_{[\text{Fe}/\text{H}]} = \text{med}([\text{Fe}/\text{H}]_{\text{MW}} - [\text{Fe}/\text{H}]_{\text{Auriga},i}) = -0.20$, $\mu_{[\text{Mg}/\text{Fe}]} = \text{med}([\text{Mg}/\text{Fe}]_{\text{MW}} - [\text{Mg}/\text{Fe}]_{\text{Auriga},i}) = 0.42$, $\sigma_{[\text{Fe}/\text{H}]}^2 = \text{var}([\text{Fe}/\text{H}]_{\text{MW}} - [\text{Fe}/\text{H}]_{\text{Auriga},i}) = 0.08^2$, and $\sigma_{[\text{Mg}/\text{Fe}]}^2 = \text{var}([\text{Mg}/\text{Fe}]_{\text{MW}} - [\text{Mg}/\text{Fe}]_{\text{Auriga},i}) = 0.02^2$.

To ensure the calibration model is realistic, we assume that there is an unknown underlying correlation across the dimensions of the calibration model, hence avoiding to drawing the offset values from each component independently. This is achieved by sampling the correlation matrix of the model, $\mathbf{C}$, from a Lewandowski–Kurowicka–Joe (2009) distribution with a concentration parameter $\eta = 1$. The covariance matrix of the calibration model is then calculated as $\Sigma = \mathbf{S} \cdot \mathbf{C} \cdot \mathbf{S}$, where $\mathbf{S}$ is a diagonal matrix that contains the standard deviation of each stellar property. The values of the non-zero elements in $\mathbf{S}$ are drawn from an exponential distribution, which describes the internal spread in the noise assigned to each star in the debris sample associated with a specific merger event. This is assumed to be marginal compared to the values of the mean vector, $\mathbf{M}$.

A comparison between the chemo-dynamical properties of the accreted star particles in the training data set as taken from the Auriga simulations and perturbed by the stochastically generated noise (blue) and the ones of the stars in the MW parent sample (red) is shown in Fig. 4. The stellar distributions are plotted as




Figure 4. Chemo-dynamical properties of the accreted star particles from the Auriga simulations used to train the posterior estimator (blue) and of the stars in the MW parent sample obtained from the APOGEE and *Gaia* surveys (red). The distribution of the MW stars belonging to the accreted substructures considered is also shown (pink). The distributions are normalized by the total number of samples.

contours showing the 0.118, 0.393, 0.675, and 0.864 confidence levels. A selection of accreted stars (pink) is derived from the MW parent sample by considering all the stars in the accreted substructure samples.

3.2 Data aggregation with Fishnets

Along with the introduction of a stochastic noise model, a data aggregation step is also added to the original method developed in A. Sante et al. (2025) before the data are passed to the MAFs ensemble for the posterior estimation. By compressing the debris data, $\hat{\mathbf{x}}$, into a lower dimensional representation, $\hat{\mathbf{x}}_{\text{agg}}$, the inference becomes more robust to the specific values of the properties of the individual stars in the sample. For this purpose, we employ the Fishnets model (T. L. Makinen, J. Alsing & B. D. Wandelt 2023), which is designed to produce robust, information-optimal embeddings from heterogeneous, noisy sets of data. The Fishnets framework is based on the additive nature of the Fisher information matrix and the log-likelihood score for independent data samples. The model operates via a pair of twin neural networks that learn to approximate the per-star score vector $[\mathbf{t}_i = \nabla_{\theta} \ln p(\mathbf{x}_i|\theta)]$ and the Fisher information matrix ($\mathbf{F}_i = -\mathbb{E}_{\theta}[\nabla \nabla^T \ln p(\mathbf{x}_i|\theta)]$) for each star i in a debris set.

For a set of N stars comprising a single merger’s debris, the total score and Fisher matrix are obtained by simple summation:

$$\mathbf{t}_{\text{total}} = \sum_{i=1}^N \mathbf{t}_i \quad ; \quad \mathbf{F}_{\text{total}} = \sum_{i=1}^N \mathbf{F}_i. \quad (1)$$

These aggregated quantities, which can be combined to form a maximum likelihood estimate of the parameters ($\hat{\mathbf{x}}_{\text{agg}} = \hat{\theta} \propto \mathbf{F}_{\text{total}}^{-1} \mathbf{t}_{\text{total}}$), serve as an information-optimal summary statistic of




Figure 5. Outline of the pipeline used for the modelling of the posterior distribution of the properties of merger events, θ , from the observed $z = 0$ chemo-dynamical properties of their debris, $\hat{\mathbf{x}}$.

the entire set of stellar debris. This summary is then used as an input to the posterior estimator.

The Fisher matrix acts as an optimal weighting scheme for each star. This makes the aggregation resilient to changes in the underlying data distribution, such as those introduced by the stochastic calibration model, ensuring that the inference is not biased by these variations and that the constraining power of the data is maximally preserved.

3.3 Posterior density estimation and validation

An end-to-end overview of the pipeline adopted to approximate the posterior probability of the properties of the disrupted progenitors of the MW is shown in Fig. 5. At first, merger–debris pairs, representing the ground-truth merger parameters, $\theta = (\tau/\text{Gyr}, \log(M_*/M_{\odot}), \log(M/M_{\odot}), \text{MMR})$, and the chemo-dynamical data of the resulting stellar debris, $\mathbf{x} = (\log(E), \log(L), [\text{Fe}/\text{H}], [\text{Mg}/\text{Fe}])$, are extracted from the Auriga simulations. To simulate realistic observational conditions, we generate 10 000 unique realizations, $\Delta \mathbf{x}$, from the calibration model. Data vectors are generated by randomly adding calibration realizations to the vectors of stellar debris.

A compressed summary representation $\hat{\mathbf{x}}_{\text{agg}}$ is obtained using a Fishnets model with four layers of 128 neurons each. To ensure robustness, a different, randomly selected calibration realization is applied to the debris data for each training epoch. This forces the model to learn the underlying features of the merger debris invariant to the specific calibration instance. The debris data are standardized using the median and inter-quartile range of the combined MW stellar sample. The Fishnets model is trained for 10 000 epochs with a batch size of 256 and a learning rate of 0.0001 using the Adam optimization algorithm (D. P. Kingma & J. Ba 2017).

Once the Fishnets model is trained, we use it to generate the data set of compressed data, $\hat{\mathbf{x}}_{\text{agg}}$, which is used for the posterior density estimation. This data set comprises 10 examples of each merger–debris pair in the Fishnet training set,




Figure 6. Predicted versus true parameters for the validation set. The top panels show the median of the predicted posterior (y-axis) against the true value (x-axis). Error bars represent the 16th–84th percentile range. The bottom panels show the normalized residuals, with the green shaded region indicating the ± 1 deviation range.




Figure 7. Validation of posterior calibration using the TARP method. The plot shows the empirical coverage probability (y-axis) versus the predicted credibility level (x-axis). The dashed black line represents a perfect estimator.

where the stellar properties are perturbed each time by a different calibration realization. The posterior distribution of the merger parameters given the aggregated chemo-dynamical information of the debris, $p(\theta|\mathbf{x}_{\text{agg}})$, is approximated by an ensemble of three MAFs. Given the reduced dimensionality of the input data after the Fishnets compression, we simplified the MAF architecture to 10 layers of 100 neurons each. Each MAF is trained for 5000 epochs, using the same optimizer, batch size, and learning rate as the Fishnets model. For each merger property, the associated posterior is derived assuming a uniform prior distribution spanning all the values expected for a merging dwarf galaxy, i.e. $p(\tau/\text{Gyr}) = \mathcal{U}(0, 14)$, $p(\log(M_*/M_\odot)) = \mathcal{U}(6, 11)$, $p(\log(M/M_\odot)) = \mathcal{U}(8, 12)$, and $p(\text{MMR}) = \mathcal{U}(-3, 0)$. The implementation and training of the density estimator models is done using the Learning the Universe Implicit Likelihood Inference framework (M. Ho et al. 2024) with the `sbi` backend.

The goodness of the inference of the trained model is estimated by comparing the predicted posterior distributions to the true merger parameters for all examples in the validation set. Each validation example is perturbed by a calibration realization not used in the training of the models, and then passed to the Fishnets model to obtain an aggregated representation. This is in turn used to estimate the posterior distribution of the merger properties with the MAFs ensemble. We sample the predicted posterior distribution 1000 times per validation example to compute the median and the 16th–84th percentile range.

Fig. 6 shows a direct comparison between the true merger property values and the median of our predictions for the un-

seen Auriga mergers. The error bars indicate the predicted 16th–84th percentile range. The bottom panels display the normalized residuals, showing the difference between the predicted median and the true value, scaled by the 16th–84th percentile range. Overall, the predictions align well with the true values across the entire parameter range, with the residuals scattered around zero, indicating that there is no significant systematic bias.

Moreover, the statistical accuracy of the posterior estimation is assessed with the Tests of Accuracy with Random Points (TARP; P. Lemos et al. 2023) method. This test compares the expected coverage probability to the actual credibility levels of the estimated distribution. For an accurate posterior estimation, the n per cent credibility region should contain the true merger parameters n per cent of the time. The results of the TARP test, shown in Fig. 7, reveal that the model is generally well calibrated, but with some parameter-specific behaviours. While the calibration for infall time and stellar mass is good, with the empirical coverage probability following closely the ideal one-to-one relationship, the one for the total progenitor mass and the merger-mass ratio show more significant deviations. In both cases, the curves lie slightly above the diagonal for low credibility levels and dip slightly below for high ones. This indicates that the model is mildly under-confident in its narrower credible intervals (e.g. the 20 per cent credible region contains the true value more than 20 per cent of the time), meaning that the predicted uncertainties for the parameters are, on average, larger than they need to be. While this is preferable to overconfidence, it suggests that the




Figure 8. Comparison of the lookback infall time (top) and stellar mass (bottom) of the MW progenitors as predicted in this work (bands) and through other studies in the literature (scatter points and bars). Scatter points represent single-value estimates, whereas bars indicate estimates quoting a confidence interval. The estimates from the literature for the Sequoia are all represented with the same marker and colour corresponding to the *N20* sample, but should be compared to the predictions from the *M19* and *K19* samples, too.

model could be more precise in its predictions for this specific property.

4 THE TIME OF ACCRETION AND STELLAR MASS OF THE MW DISRUPTED PROGENITORS

The chemo-dynamical properties of the accreted stars from the MW progenitors are used as conditions for the posterior estimation model to infer the lookback infall time, stellar mass, halo mass, and halo mass merger ratio of accreted galaxies at infall. Given that the model can be conditioned only by the properties of a hundred stars, we implement a resampling strategy to ensure a robust and unbiased estimate of the posterior distribution of the selection of the debris sample. For each accreted substructure, we randomly select $n = \lceil N/100 \rceil \times 10$ samples of 100 stars each, where N is the total number of accreted stars in that specific progenitor sample. Each of these n star distributions is then used to obtain an approximation of the posterior distribution of the merger properties. The resulting n independently generated posteriors are sampled 1000 times. For each progenitor, the summary statistics describing the properties of the galaxy at infall are calculated from the aggregated $n \times 1000$ posterior samples.

A comparison between the estimated lookback infall time and stellar mass of the MW progenitors as predicted by the GalactiKit methodology to the corresponding estimates in the literature is shown in the upper and lower panels of Fig. 8, respectively. The estimates from the inferred posterior distributions are reported as

coloured bands, whose upper and lower limits represent the 84th and 16th percentiles of the distributions, respectively, while the continuous lines crossing the bands indicate the median values of the merger property. The lookback infall time and stellar mass predictions reported by other studies are also shown for both single value (scatter points) and confidence range (error bars) estimates. The colour and symbol of the bands, scatter points and error bars for each progenitor are consistent with what are used in Figs 1 and 2. References to the studies used for the comparison as well as the specific values of the lookback infall time and stellar mass of the progenitors shown in the figure are reported in Table 1. Note that these works differ from each other not only in terms of the methodologies adopted, but also in the definition of the infall time and stellar mass and in the way the samples of accreted stars from each progenitor are selected. To the best of our knowledge, our study is the first to provide self-consistent estimates for both the lookback time and stellar mass at infall for such a large number of MW disrupted progenitors.

In general, the GalactiKit predictions for the infall times are in agreement with the estimates extrapolated through other methodologies for the most studied MW progenitors (GES, Helmi, and Sagittarius galaxies). A tension is present in the literature for the infall time of the Helmi streams, where *N*-body simulations (A. A. Kopley et al. 2007; H. H. Koppelman et al. 2019a) and star formation history from CMD fitting (T. Ruiz-Lara et al. 2022) set the accretion around 8 Gyr ago, while estimates from the associated globular cluster population (J. M.

Table 1. Lookback infall time and stellar mass predictions of the MW progenitors. From left to right, the columns indicate the name of the progenitor accreted substructure (i), the lookback infall time and stellar mass predictions provided by GalactiKit (ii, iv), and studies from the literature (iii, v).

	Lookback infall time (Gyr)		$\log(M_*/M_\odot)$	
GES	$9.5^{+0.8}_{-1.4}$	~ 10 (a) $10.2^{+0.2}_{-0.1}$ (b) $9.1^{+0.7}_{-0.7}$ (c) $7-9$ (d) ~ 10 (e) 10.5 (f) 10.5 (g) 10.2 ± 0.2 (h)	$9.0^{+0.3}_{-0.3}$ 	8.8 (q) $8.5-9$ (r) 9.7 (s) $8.5^{+0.1}_{-0.2}$ (t) $8.85-9.85$ (u) $8.43^{+0.15}_{-0.16}$ (c) 9.5 (d) 8.7 (f) $8.64^{+0.40}_{-0.44}$ (v) $8.8-8.9$ (w) $8.16^{+0.21}_{-0.17}$ (x)
Helmi streams	$10.1^{+0.7}_{-0.9}$	$6-9$ (i) $5-8$ (j) $10.1^{+0.7}_{-1.2}$ (c) 7.9 (f) ~ 8 (k) 9.4 ± 1.9 (h)	$8.4^{+0.2}_{-0.2}$ 	8.3 (y) ~ 8 (j) $7.96^{+0.19}_{-0.18}$ (c) $8.26^{+0.38}_{-0.42}$ (u) 8 (f)
Heracles	$10.4^{+0.6}_{-0.7}$	$10.5-11.6$ (l)	$8.2^{+0.2}_{-0.2}$	~ 8.7 (l) $8.86^{+0.01}_{-0.01}$ (z)
I'toi	$12.2^{+0.5}_{-0.6}$	10.4 ± 3.2 (h)	$6.7^{+0.4}_{-0.3}$	6.3 (f)
LMS-1/Wukong	$12.6^{+0.3}_{-0.4}$	~ 8 (m) 8.3 (f)	$7.3^{+0.4}_{-0.3}$	$6-7$ (m) 7.1 (f)
Sagittarius	$7.2^{+1.1}_{-1.2}$	12.9 ± 1.3 (h) $5-7$ (n) $6.8^{+1.1}_{-1.1}$ (c) 5.5 (k) $3-4$ (d) 5.9 (f)	$8.8^{+0.2}_{-0.2}$	~ 8 (aa) 8.8 (f) $8.44^{+0.22}_{-0.21}$ (c) 9.3 (d) $7.69^{+0.43}_{-0.47}$ (u)
Sequoia	$11.3^{+0.6}_{-0.7}$ (K19) $10.4^{+0.7}_{-0.9}$ (M19) $11.4^{+0.6}_{-0.7}$ (N20)	~ 9 (o) $9.4^{+0.4}_{-0.5}$ (c) $8-11$ (p) 11.6 ± 2.4 (h)	$7.4^{+0.5}_{-0.4}$ (K19) $8.5^{+0.3}_{-0.3}$ (M19) $7.3^{+0.5}_{-0.4}$ (N20)	~ 7.7 (o) $7.9^{+0.1}_{-0.1}$ (c) $7.74^{+0.42}_{-0.46}$ (u) 7.2 (f)
Thamnos	$11.8^{+0.5}_{-0.6}$	> 13 (p) 13.4 ± 0.3 (h)	$8.1^{+0.3}_{-0.4}$	< 6.7 (cc) 6.7 (f)

Note. References: (a) C. Gallart et al. (2019), (b) A. Bonaca et al. (2020), (c) J. M. D. Kruijssen et al. (2020), (d) S. Hasselquist et al. (2021), (e) J. Montalbán et al. (2021), (f) R. P. Naidu et al. (2022), (g) Y. K. González-Koda et al. (2025), (h) T. Woody et al. (2025), (i) A. A. Kepley et al. (2007), j/K19 H. H. Koppelman et al. (2019a), (k) T. Ruiz-Lara et al. (2022), (l) D. Horta et al. (2021), (m) K. Malhan et al. (2021), (n) T. J. L. de Boer, V. Belokurov & S. Koposov (2015), o/M19 G. C. Myeong et al. (2019), (p) E. Dodd et al. (2025), (q) A. Helmi et al. (2018), (r) J. T. Mackereth et al. (2019), (s) F. Vincenzo et al. (2019), (t) J. T. Mackereth & J. Bovy (2020), (u) D. K. Feuillet et al. (2020), (v) T. M. Callingham et al. (2022), (w) J. J. Han et al. (2022), (x) J. M. M. Lane et al. (2023), (y) A. Helmi & S. D. M. White (1999), (z) D. Horta & R. P. Schiavon (2025), (aa) M. Niederste-Ostholt, V. Belokurov & N. W. Evans (2012), (bb) A. W. McConnachie (2012), and (cc) H. H. Koppelman et al. (2019b).

D. Kruijssen et al. 2020) and the ages distribution of field stars as determined through isochrone fitting (T. Woody et al. 2025) suggest an earlier accretion, around 10 Gyr ago. The GalactiKit prediction favours the latter scenario (with a predicted lookback infall time of $10.1^{+0.7}_{-0.9}$ Gyr; see Table 1). However, we note that the predicted distribution of τ is sufficiently wide to accommodate

other scenarios; for example, an initial infall at ~ 10 Gyr and a second and final one at 8–9 Gyr.

Similarly, the GalactiKit predictions for the GES, Helmi, and Sagittarius galaxies' stellar masses are in agreement with results from the majority of previous studies. There is a significant variation among the stellar mass estimates reported for

Sagittarius, which span the range of $10^{7.7}–10^{9.3} M_{\odot}$. GalactiKit predicts a stellar mass in the range $10^{8.6–9.0} M_{\odot}$ for the Sagittarius galaxy at accretion, which is similar to the estimate from the mass–metallicity relation (MZR) of the globular cluster population associated with this system (J. M. D. Kruijssen et al. 2020).

The nature of the Sequoia substructure (G. C. Myeong et al. 2019; H. H. Koppelman et al. 2019b) is a matter of debate in the literature, with some studies assigning a standalone origin as the remnant of a disrupted satellite galaxy (G. C. Myeong et al. 2019; T. Matsuno, W. Aoki & T. Suda 2019; G. García-Bethencourt et al. 2023; E. Ceccarelli et al. 2024), while others argue for a connection with the GES merger as a group of stars in high-energy, retrograde orbits (H. H. Koppelman et al. 2019b; H. H. Koppelman, R. O. Y. Bos & A. Helmi 2020; J. A. S. Amarante et al. 2022) or a satellite of this system (R. P. Naidu et al. 2021). In our study, the predictions of the infall time and stellar mass of the Sequoia dwarf galaxy are calculated considering the debris samples obtained from three different selection criteria. If the selection criteria were equivalent, i.e. different methods to sample debris from the same substructure, the merger properties estimated by GalactiKit conditioned on each debris sample would be consistent with each other. This follows from the way GalactiKit was implemented as the training set contains multiple independent, identically distributed samples drawn from the debris distribution of each merger event in the Auriga simulations. However, we note a similarity between the predictions from the K19 and N20 debris samples, in which a progenitor of $\sim 10^7 M_{\odot}$ was accreted ~ 11.4 Gyr ago. The predictions obtained from the M19 sample suggest a more massive progenitor ($\sim 10^{8.5} M_{\odot}$) and a slightly more recent time of infall (~ 10.4 Gyr ago). However, the latter is consistent with the infall time inferred from the K19 and N20 samples within the errors. The estimates from the K19 and N20 samples seem to point towards Sequoia being a separate MW progenitor than GES, while the higher mass estimate from the M19 selection could also be consistent with a scenario where the debris is associated with GES. Interestingly, D. Horta et al. (2023) also find a dichotomy in the origin of Sequoia, as the chemical abundance patterns of the stars in the M19 sample are similar to those of GES, while the K19 and GES debris appear chemically inconsistent. This can be seen in Fig. 2, which shows that the M19 sample is characterized by a higher [Fe/H] and lower [Mg/Fe] population than the K19 and N20 samples, which could explain the higher stellar mass and slightly later accretion time predictions for the Sequoia substructure in this case.

Stronger disagreements are present between the GalactiKit predictions and those from previous studies for the other substructures (Heracles, I'toi, LMS-1/Wukong, and Thamnos), which might be related to an incorrect physical description of low-mass systems ($M_{*} < 10^7 M_{\odot}$) in the Auriga simulations due to particle resolution limitations. While this might explain the discrepancies between the GalactiKit predictions and the values from literature for the progenitors of the I'toi, LMS-1/Wukong, and Thamnos substructures, the same is not true for the Heracles galaxy, whose stellar mass is more precisely estimated through density fitting as $10^{8.8} M_{\odot}$ (D. Horta & R. P. Schiavon 2025). Another possible explanation is that the debris samples used to inform the GalactiKit model are not representative of the actual accreted dwarf galaxies. For instance, E. Ceccarelli et al. (2025) conducted a high-resolution chemical abundance analysis on some of the stars associated with the Thamnos substructure




Figure 9. The mass–metallicity relation for the disrupted progenitors of the MW. Each progenitor is represented as a scatter point consistently with the scheme used in Fig. 8. The point is located at the median of the observed [Fe/H] debris distribution and the median of the inferred posterior distribution of the progenitor stellar mass. The extent of the errorbars shows the 16th–84th percentile range in each dimension. In the background, the distribution of the Auriga mergers is shown in mass–metallicity bins colour-coded by their average redshift of infall. The continuous black line represents the median metallicity in each mass bin, while the dashed line shows the MZR inferred for the MW progenitors by R. P. Naidu et al. (2020).

finding a significant contamination from stars belonging to the MW *in-situ* population and, to a lesser degree, the GES galaxy.

4.1 The mass–metallicity relation of the MW progenitors

Alongside the direct comparison with the infall time and stellar mass values from the literature, an implicit test of the validity of the GalactiKit predictions can be performed by investigating whether the predicted merger properties follow known physical relations. For this purpose, we compare the inferred progenitor stellar mass with the average measured metallicity of the stellar debris. Empirically, a positive relation is observed between these two quantities for galaxies in the mass range $10^{3.5–12} M_{\odot}$ (E. N. Kirby et al. 2013). This stellar MZR can be explained in terms of the amount of gas retained by galaxies through their gravitational pull.

Fig. 9 shows the inferred MZR for the disrupted progenitors of the MW, represented as scatter points using the same scheme as in Fig. 8. The MZR in the Auriga simulations is also shown by binning the merger galaxies in both mass and metallicity. Random samples from the calibration model were added to the simulation data to allow a comparison with the observed metallicities. The colour of each bin indicates the median redshift at which the galaxies crossed the virial radius of the host. The continuous black line represents the median metallicity values in the mass bins. Another prediction for the MZR of the disrupted MW progenitors, inferred by R. P. Naidu et al. (2022) using comparisons with *N*-body simulations, is also shown as a dashed black line. In general, the GalactiKit stellar mass predictions are consistent with this previously predicted MZR. However, because the progenitors were accreted at different redshifts and because the MZR is expected to evolve with time, we generally expect the




Figure 10. Mass–metallicity relation for the disrupted satellites of the MW as in Fig. 9. The dashed lines represent the empirical MZR calculated with equation (2) for $z = 0$ (top line), $z = 1$, $z = 2$, $z = 3$, and $z = 4$ (top to bottom), respectively. The MZR at $z = 0$ is the same as the E. N. Kirby et al. (2013) relation.

progenitors to follow different MZR. The time evolution of MZR is seen in both simulations (X. Ma et al. 2016; P. Torrey et al. 2019; A. M. Garcia et al. 2024; S. E. Grillozzi, A. S. Font & M. E. De Rossi 2024) and observations (R. L. Sanders et al. 2021; G. Roberts-Borsani et al. 2024), a behaviour which is mainly due to a lower metal abundance of galaxies at higher redshifts. Thus, while the slope of the MZR tends to stay constant, the normalization factor of the relation scales over time. This can also be noticed by looking at the coloured bins in Fig. 9, which show that in the Auriga simulations galaxies with higher metallicity in a given mass bin have been accreted at earlier times. Therefore, galaxies from late merger events, such as Sagittarius, are expected to follow an MZR with higher normalization than galaxies accreted at earlier epochs (such as I’itoi and LMS-1/Wukong).

To investigate whether MW progenitors accreted at different times are consistent with different observed MZRs, we compute the redshift evolution of this relation by considering the MZR at $z = 0$ (E. N. Kirby et al. 2013) and modelling the change in the normalization factor as done in R. L. Sanders et al. (2021), i.e. using a gradient of $d[\text{Fe}/\text{H}]/dz \approx -0.11$. This evolution of the metallicity content comes from a study of the gas-phase metallicity of galaxies in redshift $z = 0 - 3.3$ from the MOSDEF survey. Although some of the MW accreted progenitors are predicted to be accreted at $z > 3.3$, recent *JWST* observations (K. Nakajima et al. 2023; A. Sarkar et al. 2025) show that the MZR does not evolve significantly past this redshift. Therefore, we assume a constant ratio between the average stellar and gas-phase metallicity for galaxies of similar mass across redshifts and calculate the observed MZR as follows:

$$\begin{aligned} [\text{Fe}/\text{H}]_{z=z'} &= -0.11 \times z' + [\text{Fe}/\text{H}]_{z=0} \\ &= -0.11 \times z' + (-1.69) + 0.30 \times (\log(M_*/M_\odot) - 6). \end{aligned} \quad (2)$$

Fig. 10 shows the predicted MZR for MW progenitors against the same empirical relation calculated using equation (2) at $z = 0$ (top line), $z = 1$, $z = 2$, $z = 3$, and $z = 4$ (bottom line). The

empirical lines are coloured by the redshift value at which they were estimated and, similarly, the predicted scatter points are coloured by the redshift of infall predicted for each merger. It appears that most of the MW accreted progenitors follow the MZR evaluated at the closest redshift to their time of infall. A discrepancy can be noticed between the different sample selections for the Sequoia galaxy. Although the *K19* (left triangle) and *N20* (diamond) stellar mass predictions are similar to each other, only the former lies on an MZR consistent with its predicted infall time. Similarly, the *M19* (right triangle) stellar mass prediction, which is considerably larger than the one from the *K19* and *N20* samples, lies on an MZR inconsistent with its redshift of infall. Therefore, for the rest of the analysis, we consider the *K19* sample as our preferred representation for the debris of the Sequoia galaxy because it also has a lower risk of contamination from GES stars compared to the more metal-rich *M19* sample.

4.2 The growth of the accreted stellar mass and halo mass of the MW

The predicted lookback infall time and stellar mass of the MW progenitors allow us to trace back the evolution of the accreted mass content of the Galaxy across cosmic time. Because GalactiKit returns the probability distributions of these quantities, the cumulative distribution of the MW accreted stellar mass as a function of redshift, $\text{CDF}_{\text{MW}}(M_{*,\text{ACC}}, z)$, can be estimated numerically from the ensemble of the single $\text{CDF}_{\text{MW},i}(M_{*,\text{ACC}}, z)$. A single (*i*th) prediction of the MW assembly history can be computed by drawing a sample from the posterior distribution of each merger event. The cumulative MW accreted stellar mass distribution, $\text{CDF}_{\text{MW},i}(M_{*,\text{ACC}}, z)$, can then be calculated by adding the stellar mass estimates of each event ordered by their redshift of infall, which is obtained from the lookback infall time assuming the same cosmology as in the Auriga simulations. We repeat this procedure for $i = 1000$ times obtaining a distribution of $\text{CDF}_{\text{MW},i}(M_{*,\text{ACC}}, z)$. For this analysis, we exclude Thamnos, whose mass estimate is significantly overpredicted by the model compared to the estimates from the literature.

Fig. 11 shows the median (black line) and 16th-to-84th percentile range (grey area) of the predicted cumulative accreted mass distribution of the MW, $\text{CDF}_{\text{MW}}(M_{*,\text{ACC}}, z)$. For reference, the median redshifts of infall of the considered merger events are shown as vertical dashed lines. The Galaxy is predicted to have obtained the majority of its accreted stellar mass between $z = 1$ and $z = 2$, following the mergers with the GES, Helmi streams, and Heracles dwarf galaxies. A significant mass contribution is also provided by the latest merger with the Sagittarius dwarf galaxy, which in this calculation is considered in its totality, i.e. debris and surviving satellite.

More quantitatively, the fraction of accreted stellar mass contributed by each merger event is estimated by calculating the ratio of the stellar mass of a progenitor and the total MW accreted stellar mass for each *i*-th assembly history. The median and 16th-to-84th percentile range of the distribution of the 1000 mass fractions of each merger event is reported in Table 2. GES is predicted to be the most significant contributor of accreted stars in the MW (~ 0.43), which may have contributed to the disc formation phases (N. Funakoshi et al. 2025). Only the Helmi streams (~ 0.10) and Heracles (~ 0.07) among the fully phase-mixed substructures contribute more than 0.01. The total stellar mass of Sagittarius is also a significant fraction (~ 0.31) of the mass supposedly accreted by the MW. However, because this is




Figure 11. Cumulative accreted stellar mass distribution for the MW through cosmic time with (top) and without (bottom) the Sagittarius merger. This calculation only accounts for 1 Sequoia sample (*K19*). The plot also reports measurements of the stellar mass of the MW halo, excluding the Magellanic Clouds, obtained through isochrone fitting (A. J. Deason, V. Belokurov & J. L. Sanders 2019) and density modelling (J. T. Mackereth & J. Bovy 2020) of RGB stars.

Table 2. Fraction of progenitor stellar mass to MW total accreted stellar mass. From left to right, the columns indicate the median (50th), 84th, and 16th percentile of the mass fractions MW progenitors estimated from 1000 sampled realizations of the Galaxy assembly history.

	50th	84th	16th
GES	0.425	0.612	0.252
Helmi streams	0.099	0.180	0.054
Heracles	0.071	0.124	0.037
I'toi	0.002	0.006	0.001
LMS-1/Wukong	0.009	0.024	0.004
Sagittarius	0.309	0.474	0.188
Sequoia	0.013	0.042	0.005

an ongoing merger, we cannot predict how much of this mass is still bound to the surviving satellite as GalactiKit predicts the mass of the system at the time it first interacted with the MW by crossing its virial radius.

However, under the assumption that the majority of the accreted stellar content is within the stellar halo, an investigation of

the amount of the accreted stellar mass retained by the MW after its mergers can be conducted by comparing the predicted total accreted stellar mass to measurements of the stellar halo mass. For this purpose, we report in Fig. 11 two estimates of the $z = 0$ total stellar halo mass obtained from number counting (red band; A. J. Deason et al. 2019) and density modelling (blue band; J. T. Mackereth & J. Bovy 2020) of red giant branch (RGB) stars in the MW stellar halo. Neither of the estimates includes the contribution of the Magellanic Clouds. The present-day MW accreted stellar mass predicted by GalactiKit is $2.2^{+1.1}_{-0.6} \times 10^9 M_\odot$. This estimate is consistent within the uncertainty range with the stellar halo mass measured by A. J. Deason et al. (2019; $1.4 \pm 0.4 \times 10^9 M_\odot$), while being larger than the estimate from J. T. Mackereth & J. Bovy (2020; $1.3^{+0.3}_{-0.2} \times 10^9 M_\odot$). However, the latter study does not account for the contribution of the Sagittarius dwarf galaxy (J. T. Mackereth & J. Bovy 2020). Ignoring this contribution, we found the total MW accreted stellar mass to be $1.4^{+1.1}_{-0.4} \times 10^9 M_\odot$ (bottom panel of Fig. 11), which is in agreement with both studies. Therefore, the total stellar mass accreted from the progenitors of the substructures identified in the stellar halo is consistent with measurements of the mass of this component. This suggests that, in line with other findings (e.g. R. P. Naidu et al. 2020), the stellar halo is dominated by accreted stars.

Alongside providing predictions for the stellar mass and time of infall, GalactiKit infers also the halo mass of the accreted progenitor, $\log(M/M_\odot)$, and halo mass merger ratio, $\text{MMR} = \log(M/M_\odot) - \log(M_{\text{MW}}/M_\odot)$, of each progenitor. This information can be combined to infer the MW halo mass, $\log(M_{\text{MW}})$, at the merger time, i.e.

$$\log(M_{\text{MW}})_{z=z_{\text{infall}}} = \log(M_{\text{prog}})_{z=z_{\text{infall}}} - \text{MMR}_{z=z_{\text{infall}}}. \quad (3)$$

Therefore, the growth of the MW dark matter halo across cosmic time can be reconstructed from its assembly history. Given the underlying assumption that each merger is an independent event and that the progenitor properties are independent and identically distributed samples, a distribution of the MW halo mass across redshifts is calculated, using equation (3), by pooling the samples from all the progenitor posteriors. The median and 16th-to-84th percentile range of the MW halo mass estimates in different redshift bins is shown in Fig. 12 as a black line and a grey shaded area. For context, we also show the density contours of the 2D distribution of the halo mass and redshift of infall inferred for the progenitors considered in the calculation. Similarly to the previous analysis, we exclude samples from the Thamnos substructure.

At the earliest traced time ($z \approx 5$), the Galaxy is predicted to live within a dark matter halo of mass $1.0^{+0.4}_{-0.3} \times 10^{11} M_\odot$. After the first merger events (LMS-1/Wukong, I'toi, and Sequoia), the mass of the halo increases to $2.6^{+0.8}_{-0.8} \times 10^{11} M_\odot$ at $z \approx 2$. This is unlikely due only to the accretion of those galaxies, whose dark matter halo masses are predicted to be $< 10^{10} M_\odot$. A significant growth is observed after the accretion of the Heracles, Helmi streams, and GES dwarf galaxies, as the Galaxy halo reaches a mass of $5.5^{+1.4}_{-1.1} \times 10^{11} M_\odot$ before starting its interaction with Sagittarius at $z \approx 1$. It is interesting to see that GES contributed to a large mass increase around $z \approx 1.5$, which may have led to the thick-to-thin disc formation transition due to the cold-to-hot gas accretion phase change (R. J. J. Grand et al. 2018; N. Funakoshi et al. 2025). Because the MW halo mass estimates refer to the mass of this component prior to an accretion event, the total dark matter mass of the MW after the accretion of the Sagittarius dwarf galaxy cannot be estimated directly from the pooled




Figure 12. Distribution of the MW dark matter halo mass estimates as a function of redshift, inferred from the posterior samples of the progenitor halo masses and corresponding halo-mass ratios with the MW at infall. The black line indicates the median MW halo mass at a given redshift, while the grey shaded area shows the 16th–84th percentile range. For context, the 2σ coloured contours show the joint posterior distribution of the halo mass and infall redshift for each progenitor.

measurements (grey line in Fig. 12). However, the total MW halo mass after the merger can still be inferred by adding the Sagittarius dark matter halo mass to the estimate of the Galactic halo at the time of infall of the satellite. Thus, the ‘post-merger’ estimate of the MW dark matter halo is $5.9^{+1.4}_{-1.1} \times 10^{11} M_{\odot}$, which appears to be already in agreement with the lower end of the expected range from dynamical measurements of the total MW mass (e.g. J. Wang, F. Hammer & Y. Yang 2022). However, this estimate does not account for the evolution of the Galaxy after $z \approx 1$, which includes significant accretion events such as the merger with the Magellanic Clouds.

5 CONCLUSIONS

In this study, we presented simulation-based predictions of the lookback infall time, stellar mass, halo mass, and halo mass merger ratio of the MW accreted progenitors. Applying the GalactiKit SBI framework to Auriga simulations, we developed a model of the posterior distribution of the properties of mergers at infall conditioned on the chemo-dynamical distribution of debris at $z = 0$. However, we modified the original methodology of A. Sante et al. (2025) to explicitly account for the discrepancy between stellar properties in simulations and observations (Section 3.1), hence ensuring predictions are robust against this uncertainty. Moreover, an additional data aggregation step (Section 3.2) was

added to maximize the amount of information extracted from the stellar properties. We applied the new version of GalactiKit on a sample of unseen merger–debris examples from the Auriga simulations, finding that the posterior predictions are accurate and consistent with the expected distributions (Figs 6 and 7). We then proceeded to estimate the properties of the MW mergers conditioned on the chemo-dynamics of their $z = 0$ debris. Our main findings are as follows:

- (i) a general agreement exists between the GalactiKit predictions and values from the literature for the lookback infall time and stellar mass of the MW progenitors (Fig. 8). A notable exception is the Thamnos substructure, whose stellar mass is considerably larger than the value reported by H. H. Koppelman et al. (2019b). However, this can be caused by significant contamination by *in-situ* stars in the merger debris sample (E. Ceccarelli et al. 2025).
- (ii) the infall time and stellar mass predictions obtained from samples of the Sequoia progenitor obtained from different selection criteria are inconsistent with each other. The *M19* sample predicts a merger event similar in time and mass to the accretion of GES, while the *K19* and *N20* samples point towards an earlier and less massive merger. This might suggest that the three selection criteria used here may not fully represent this system or that there is significant contamination from GES stars in the *M19* sample.

(iii) the predicted progenitor stellar masses are positively related to the debris median metallicities, with mergers following different relations depending on their redshift of infall (Fig. 10).

(iv) the predicted total stellar mass accreted by the MW is $2.2^{+1.1}_{-0.6} \times 10^9 M_{\odot}$ (Fig. 11), with GES (~ 0.43) and Sagittarius (~ 0.31) accounting for the majority of it. Excluding the mass contributions from Sagittarius, we predict a total accreted stellar mass of $1.4^{+1.1}_{-0.4} \times 10^9 M_{\odot}$ from the substructures in the stellar halo. This is in good agreement with the mass measurements of this component of the Galaxy.

(v) the inferred growth of the MW dark matter halo mass is not entirely explained by the current census of mergers, suggesting the potential presence of undetected accretion events or by a significant accretion of dark matter-dominated systems, especially between $z \approx 5$ and $z \approx 2$. The prediction for the MW halo mass after the accretion of Sagittarius is $5.9^{+1.4}_{-1.1} \times 10^{11} M_{\odot}$.

We caution that our results rely strongly on the selection of the debris samples defining each accreted progenitor and that GalactiKit assumes that each sample conditioning the merger property posterior is drawn independently from the complete chemo-dynamical distribution of the debris. Hence, the impact of the contamination from the *in-situ* population or from stars accreted from other mergers is not accounted for. This is a scope for future investigations. Similarly, potential biases caused by the survey selection functions affecting the observable debris chemo-dynamical distribution remain to be studied. In addition, the choice of particle resolution in simulations has an effect on the realism of small galaxies and, consequently, on the predictions of their properties at infall. A comparison between the infall time and stellar mass predictions for these systems obtained from posterior estimators trained at different resolution levels is needed to quantify the importance of resolution effects on SBI predictions of merger events.

Nevertheless, we find that the SBI methodology provides remarkable insights into the Galaxy formation history and that the accuracy of these results is on a par with those from the equivalent corpus of studies in the literature. Modern and upcoming surveys provide a unique opportunity to advance these machine learning techniques, unlocking the potential to reconstruct the formation of the Galaxy and the Local Group with unprecedented fidelity.

ACKNOWLEDGEMENTS

We would like to thank the reviewer for their helpful suggestions. AS acknowledges a Science Technologies Facilities Council (STFC) PhD studentship at the LIV.INNO Centre for Doctoral Training 'Innovation in Data Intensive Science'. This study used the Prospero high-performance computing facility at Liverpool John Moores University. AS thanks Danny Horta for insightful discussions on the MW formation history. AF acknowledges support from UKRI (ST/W006766/1). RJG acknowledges support from an STFC Ernest Rutherford Fellowship (ST/W003643/1). DK acknowledges MWGaiaDN, a Horizon Europe Marie Skłodowska-Curie Actions Doctoral Network funded under grant agreement no. 101072454 and also funded by UK Research and Innovation (EP/X031756/1). TLM acknowledges support from the InfoSys Physics-AI Lab at the University of Cambridge.

DATA AVAILABILITY

The Auriga simulations can be downloaded through the Globus file transfer service (<https://globus.org/>). Detailed instructions on how to access the data are provided in <https://wwwmpa.mpa-garching.mpg.de/auriga/data.html>.

The codes for reproducing the results of the analysis are available at <https://github.com/ariasant/MWsbj>.

REFERENCES

- Abdurro'uf et al., 2022, *ApJS*, 259, 35
Amarante J. A. S., Debattista V. P., Beraldo e Silva L., Laporte C. F. P., Deg N., 2022, *ApJ*, 937, 12
Belokurov V., Kravtsov A., 2023, *MNRAS*, 525, 4456
Belokurov V., Erkal D., Evans N. W., Koposov S. E., Deason A. J., 2018, *MNRAS*, 478, 611
Bennett M., Bovy J., 2019, *MNRAS*, 482, 1417
Bonaca A., Price-Whelan A. M., 2025, *New Astron. Rev.*, 100, 101713
Bonaca A. et al., 2020, *ApJ*, 897, L18
Buder S. et al., 2021, *MNRAS*, 506, 150
Callingham T. M., Cautun M., Deason A. J., Frenk C. S., Grand R. J. J., Marinacci F., 2022, *MNRAS*, 513, 4107
Ceccarelli E. et al., 2024, *A&A*, 684, A37
Ceccarelli E., Massari D., Palla M., Mucciarelli A., Bellazzini M., Helmi A., 2025, *A&A*, 704, A180
Chabrier G., 2003, *PASP*, 115, 763
Conroy C. et al., 2019, *ApJ*, 883, 107
Cranmer K., Brehmer J., Louppe G., 2020, *Proc. Natl. Acad. Sci. USA*, 117, 30055
Cui X.-Q. et al., 2012, *Res. Astron. Astrophys.*, 12, 1197
de Boer T. J. L., Belokurov V., Koposov S., 2015, *MNRAS*, 451, 3489
Deason A. J., Belokurov V., 2024, *New Astron. Rev.*, 99, 101706
Deason A. J., Belokurov V., Sanders J. L., 2019, *MNRAS*, 490, 3426
Dodd E., Ruiz-Lara T., Helmi A., Gallart C., Callingham T. M., Cassisi S., Fernández-Alvar E., Surot F., 2025, *A&A*, 698, A277
Feuillet D. K., Feltzing S., Sahlholdt C. L., Casagrande L., 2020, *MNRAS*, 497, 109
Funakoshi N., Kawata D., Sanders J. L., Ciucă I., Grand R. J. J., Zhāng H., 2025, *MNRAS*, 543, 2275
Gaia Collaboration, 2018, *A&A*, 616, A1
Gaia Collaboration, 2021, *A&A*, 649, A1
Gallart C., Bernard E. J., Brook C. B., Ruiz-Lara T., Cassisi S., Hill V., Monelli M., 2019, *Nat. Astron.*, 3, 932
García-Bethencourt G., Brook C. B., Grand R. J. J., Kawata D., 2023, *MNRAS*, 526, 1190
García Pérez A. E. et al., 2016, *AJ*, 151, 144
García A. M. et al., 2024, *MNRAS*, 529, 3342
González-Koda Y. K. et al., 2025, *A&A*, 704, A259
Grand R. J. J. et al., 2017, *MNRAS*, 467, 179
Grand R. J. J. et al., 2018, *MNRAS*, 474, 3629
Grand R. J. J., Fragkoudi F., Gómez F. A., Jenkins A., Marinacci F., Pakmor R., Springel V., 2024, *MNRAS*, 532, 1814
GRAVITY Collaboration, 2019, *A&A*, 625, L10
Grimozzi S. E., Font A. S., De Rossi M. E., 2024, *MNRAS*, 530, 95
Han J. J. et al., 2022, *AJ*, 164, 249
Hasselquist S. et al., 2021, *ApJ*, 923, 172
Helmi A., 2020, *ARA&A*, 58, 205
Helmi A., White S. D. M., 1999, *MNRAS*, 307, 495
Helmi A., White S. D. M., de Zeeuw P. T., Zhao H., 1999, *Nature*, 402, 53
Helmi A., Babusiaux C., Koppelman H. H., Massari D., Veljanoski J., Brown A. G. A., 2018, *Nature*, 563, 85
Ho M. et al., 2024, *Open J. Astrophys.*, 7, 54
Horta D., Schiavon R. P., 2025, *MNRAS*, 537, 3730
Horta D. et al., 2021, *MNRAS*, 500, 1385
Horta D. et al., 2023, *MNRAS*, 520, 5671
Ibata R. A., Gilmore G., Irwin M. J., 1994, *Nature*, 370, 194

- Jean-Baptiste I., Di Matteo P., Haywood M., Gómez A., Montuori M., Combes F., Semelin B., 2017, *A&A*, 604, A106
- Johnston K. V., Hernquist L., Bolte M., 1996, *ApJ*, 465, 278
- Kepley A. A. et al., 2007, *AJ*, 134, 1579
- Kingma D. P., Ba J., 2017, preprint ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980))
- Kirby E. N., Cohen J. G., Guhathakurta P., Cheng L., Bullock J. S., Gallazzi A., 2013, *ApJ*, 779, 102
- Koppelman H. H., Helmi A., Massari D., Roelenga S., Bastian U., 2019a, *A&A*, 625, A5
- Koppelman H. H., Helmi A., Massari D., Price-Whelan A. M., Starkenburg T. K., 2019b, *A&A*, 631, L9
- Koppelman H. H., Bos R. O. Y., Helmi A., 2020, *A&A*, 642, L18
- Kruijssen J. M. D., Pfeffer J. L., Crain R. A., Bastian N., 2019, *MNRAS*, 486, 3134
- Kruijssen J. M. D. et al., 2020, *MNRAS*, 498, 2472
- Lane J. M. M., Bovy J., Mackereth J. T., 2023, *MNRAS*, 526, 1209
- Lemos P., Coogan A., Hezaveh Y., Perreault-Levasseur L., 2023, in Krause A., Brunskill E., Cho K., Engelhardt B., Sabato S., Scarlett J., eds, Proc. Machine Learning Research, Vol. 202. PMLR, Honolulu, Hawaii, p. 19256
- Leung H. W., Bovy J., 2019, *MNRAS*, 489, 2079
- Lewandowski D., Kurowicka D., Joe H., 2009, *J. Multivar. Anal.*, 100, 1989
- Ma X., Hopkins P. F., Faucher-Giguère C.-A., Zolman N., Muratov A. L., Kereš D., Quataert E., 2016, *MNRAS*, 456, 2140
- Mackereth J. T., Bovy J., 2020, *MNRAS*, 492, 3631
- Mackereth J. T. et al., 2019, *MNRAS*, 482, 3426
- Majewski S. R. et al., 2017, *AJ*, 154, 94
- Makinen T. L., Alsing J., Wandelt B. D., 2023, preprint ([arXiv:2310.03812](https://arxiv.org/abs/2310.03812))
- Malhan K., Rix H.-W., 2024, *ApJ*, 964, 104
- Malhan K., Yuan Z., Ibata R. A., Arentsen A., Bellazzini M., Martin N. F., 2021, *ApJ*, 920, 51
- Matsuno T., Aoki W., Suda T., 2019, *ApJ*, 874, L35
- McConnachie A. W., 2012, *AJ*, 144, 4
- McMillan P. J., 2017, *MNRAS*, 465, 76
- Monachesi A. et al., 2019, *MNRAS*, 485, 2589
- Montalbán J. et al., 2021, *Nat. Astron.*, 5, 640
- Mori A., Di Matteo P., Salvadori S., Khoperskov S., Pagnini G., Haywood M., 2024, *A&A*, 690, A136
- Myeong G. C., Vasiliev E., Iorio G., Evans N. W., Belokurov V., 2019, *MNRAS*, 488, 1235
- Naidu R. P., Conroy C., Bonaca A., Johnson B. D., Ting Y.-S., Caldwell N., Zaritsky D., Cargile P. A., 2020, *ApJ*, 901, 48
- Naidu R. P. et al., 2021, *ApJ*, 923, 92
- Naidu R. P. et al., 2022, preprint ([arXiv:2204.09057](https://arxiv.org/abs/2204.09057))
- Nakajima K., Ouchi M., Isobe Y., Harikane Y., Zhang Y., Ono Y., Umeda H., Oguri M., 2023, *ApJS*, 269, 33
- Niederste-Ostholt M., Belokurov V., Evans N. W., 2012, *MNRAS*, 422, 207
- Papamakarios G., Pavlakou T., Murray I., 2017, in Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., eds, NeurIPS Proc. Vol. 30, Advances in Neural Information Processing Systems 30 (NIPS 2017). Curran Associates, Inc.
- Reid M. J., Brunthaler A., 2020, *ApJ*, 892, 39
- Roberts-Borsani G. et al., 2024, *ApJ*, 976, 193
- Ruiz-Lara T., Helmi A., Gallart C., Surot F., Cassisi S., 2022, *A&A*, 668, L10
- Sanders R. L. et al., 2021, *ApJ*, 914, 19
- Sante A., Kawata D., Font A. S., Grand R. J. J., 2025, *MNRAS*, 542, 1776
- Sante A., Font A. S., Mistry D., Ortega-Martorell S., Olier I., 2026, *MNRAS*, 547, stag503
- Sarkar A. et al., 2025, *ApJ*, 978, 136
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schiavon R. P. et al., 2024, *MNRAS*, 528, 1393
- Schönrich R., Binney J., Dehnen W., 2010, *MNRAS*, 403, 1829
- Searle L., Zinn R., 1978, *ApJ*, 225, 357
- Springel V., 2010, *MNRAS*, 401, 791
- Thomas G. F., Battaglia G., Grand R. J. J., Aguiar Álvarez A., 2025, *A&A*, 704, A40
- Torrey P. et al., 2019, *MNRAS*, 484, 5587
- Vincenzo F., Spitoni E., Calura F., Matteucci F., Silva Aguirre V., Miglio A., Cescutti G., 2019, *MNRAS*, 487, L47
- Wang J., Hammer F., Yang Y., 2022, *MNRAS*, 510, 2242
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Woody T. et al., 2025, *ApJ*, 978, 152
- Yuan Z. et al., 2020, *ApJ*, 891, 39

APPENDIX A: MERGER POSTERIOR DISTRIBUTIONS OF MW ACCRETED SUBSTRUCTURES

The posterior distributions of the lookback infall time, stellar and halo mass, and halo mass merger ratio of the MW accreted satellites were estimated from the present-day chemo-dynamics of their debris using an ensemble of Masked Autoregressive Flows following the GalactiKit methodology. Throughout the paper, we used samples from these distributions to compute summary statistics, obtaining a quantitative description of the properties of each merger event and, overall, of the MW assembly history.

In this section, the posterior sample distributions for the accreted substructures are reported in individual figures: GES (Fig. A1), Helmi streams (Fig. A2), Heracles (Fig. A3), I'toi (Fig. A4), LMS-1/Wukong (Fig. A5), Sagittarius (Fig. A6), and Thamnos (Fig. A10). For Sequoia, we provide three distributions conditioned on debris selected with different criteria: *K19* (Fig. A7), *M19* (Fig. A8), and *N20* (Fig. A9). Each distribution is presented as a corner plot of the inferred galaxy merger properties, with one-dimensional density distributions showing the median and the 16th–84th percentile range. Samples are coloured consistently with the substructure colour scheme adopted in the analysis.




Figure A1. bibkey1.




Figure A2. bibkey2.




Figure A3. bibkey3.




Figure A4. bibkey4.




Figure A5. bibkey5.




Figure A6. bibkey6.




Figure A7. bibkey7.




Figure A8. bibkey8.




Figure A9. bibkey9.




Figure A10. bibkey10.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.