

GalactiKit: reconstructing mergers from $z = 0$ debris using simulation-based inference in Auriga

Andrea Sante ,  Daisuke Kawata ,  2,3 Andreea S. Font ,  1 and Robert J. J. Grand ,  1

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*

Accepted 2025 August 12. Received 2025 August 12; in original form 2025 February 21

ABSTRACT

We present GalactiKit, a data-driven methodology for estimating the lookback infall time, stellar mass, halo mass, and mass ratio of the disrupted progenitors of Milky Way-like galaxies at the time of infall. GalactiKit uses simulation-based inference to extract the information on galaxy formation processes encoded in the Auriga cosmological magnetohydrodynamic (MHD) simulations of Milky Way-mass haloes to create a model that relates the properties of mergers to those of the corresponding merger debris at $z = 0$. We investigate how well GalactiKit can reconstruct the merger properties given the dynamical, chemical, and the combined chemodynamical information of debris. For this purpose, three models were implemented considering the following properties of merger debris: (a) total energy and angular momentum, (b) iron-to-hydrogen and alpha-to-iron abundance ratios, and (c) a combination of all of these. We find that the kinematics of the debris can be used to trace the lookback time at which the progenitor was first accreted into the main halo. However, chemical information is necessary for inferring the stellar and halo masses of the progenitors. In both models (b) and (c), the stellar masses are predicted more accurately than the halo masses, which could be related to the scatter in the stellar mass–halo mass relation. Model (c) provides the most accurate predictions for the merger parameters, which suggests that combining chemical and dynamical data of debris can significantly improve the reconstruction of the Milky Way’s assembly history.

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

1 INTRODUCTION

In the standard cosmological model, galaxies form hierarchically by merging with smaller satellite systems over time (White & Rees 1978). A natural consequence of this ‘bottom-up’ formation scenario is the assembly of a stellar component in the outer region of galaxies, the stellar halo, dominated by tidally stripped stars once orbiting the cannibalized systems (Searle & Zinn 1978). Based on the idea that stellar populations preserve information on the history of the environment since they formed (Eggen, Lynden-Bell & Sandage 1962; Searle & Zinn 1978), stellar haloes have become objects of study as a bridge between the present and past of a galaxy’s evolution.

At the first stages of accretion, tidal debris tend to follow the orbit of the merging progenitor as spatially coherent structures, called streams. As the progenitor gradually sinks in the gravitational potential of the host, stellar streams lose their coherence blending with the stars formed *in situ* (Knebe et al. 2005). Because galaxies can be regarded, even when merging, as collisionless systems (Binney & Tremaine 2008), the probability density of the distribution function of stars around a point in phase-space is conserved over time (Liouville’s theorem). Exploiting this property, Helmi & White

(1999) devised a methodology to separate phase-mixed structures belonging to different progenitors as overdensities in a new physical space described by the *integrals of motion*, quantities which are constant in time.

A further method for inferring the assembly history of a galaxy from its present-day stellar content is by ‘chemical tagging’ (Freeman & Bland-Hawthorn 2002). Satellite galaxies have an independent chemical evolution history that imprints a unique signature on the chemical composition of forming stars; hence, the abundances of chemical elements can be used to ‘tag’ stars with a similar formation site and time.

While earlier ‘archaeological’ investigations focused on the stellar halo of the Milky Way (see reviews of Helmi 2020; Deason & Belokurov 2024; Bonaca & Price-Whelan 2025), where longer time-scales for phase-mixing enable the retention of information about the mergers, more recent studies have extended the search for debris into the disc and inner regions of the Galaxy, as more data have become available (Arentsen et al. 2020; Kawata et al. 2024). The first hints to the assembly history of the Galaxy were detected with the Sagittarius (Ibata, Gilmore & Irwin 1994) and Helmi (Helmi et al. 1999) streams, which led to the discovery of examples of late and early accretion events; however, it was only with the first and second data release of the European Space Agency’s *Gaia* mission (Gaia Collaboration 2016, 2018) that a more detailed picture of the past

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

of the Milky Way came into view. The measurements of accurate astrometry and photometry for over a billion stars significantly enhanced the discovery of streams and kinematic substructures associated to merging events. Moreover, the complementary data on radial velocities and chemical abundances for tens of millions of stars obtained with ground-based spectroscopic surveys – such as the Apache Point Observatory Galactic Evolution Experiment (APOGEE; Majewski et al. 2017), the Galactic Archaeology with HERMES (GALAH; Buder et al. 2021), and the Large sky Area Multi-Object fiber Spectroscopic Telescope (LAMOST; Cui et al. 2012), to mention a few – pushed further back our gaze into the first stages of the Milky Way evolution (Kruijssen et al. 2019b; Horta et al. 2021; Belokurov & Kravtsov 2022; Malhan & Rix 2024).

The general consensus is that the assembly history of the Milky Way was dominated by a radial, massive accretion event which occurred ~ 10 billion years ago, named *Gaia* Sausage/Enceladus (GS/E; Chiba & Beers 2000; Brook et al. 2003; Meza et al. 2005; Belokurov et al. 2018; Helmi et al. 2018). This merger is thought to have had a dramatic effect on the proto-Milky Way by dynamically heating the stars in the disc and bringing in a significant amount of gas which led to a burst in star formation (Bignone, Helmi & Tissera 2019; Grand et al. 2020; Ciucă et al. 2024). Although discoveries of a plethora of minor mergers have been claimed from the identification of stellar streams (Malhan, Ibata & Martin 2018) and chemodynamical studies in the inner halo (Naidu et al. 2020; Dodd et al. 2023; Horta et al. 2023), the GS/E is believed to be the last significant merger, making the assembly history of the Galaxy unusually quiet (Hammer et al. 2007). Currently, the chronology of the merger events describing the Milky Way assembly history is inferred through the ages of the associated globular cluster and stellar populations. The former are derived empirically from the age–metallicity relation (Massari, Koppelman & Helmi 2019; Kruijssen et al. 2019b), while the latter can be measured directly with asteroseismology techniques (Miglio et al. 2017; Montalbán et al. 2021), or inferred by comparing photometric and spectroscopic observations to stellar evolution models via isochrone (Queiroz et al. 2018; Sanders & Das 2018) and colour–magnitude diagram fitting (Gallart et al. 2024). Similarly, methods for estimating the mass of the disrupted progenitors of the Milky Way involve using empirical relations such as the mass–metallicity relation (Harmsen et al. 2017; Naidu et al. 2022), or fitting the density profile of debris (Mackereth & Bovy 2020; Lane, Bovy & Mackereth 2023).

In this study, we introduce GalactiKit: a data-driven methodology to infer the properties of the disrupted progenitors of the Milky Way. In particular, we focus on retrieving the infall time, stellar mass, halo mass, and the halo–mass ratio with the host of mergers at the time of infall. These are determined by performing a Bayesian inference analysis assuming that the progenitor properties are parameters of an undefined model that produces the chemodynamical distribution of the merger debris at $z = 0$. Despite the model being unknown, there are multiple examples of mergers and phase-space evolution of debris within cosmological simulations; hence, a mapping between the debris and mergers properties can be learned using simulation-based inference (SBI; Cranmer, Brehmer & Louppe 2020). Using the information contained in the simulations, SBI techniques are designed to perform Bayesian inference on models which are not analytically defined, or which have intractable likelihood distribution functions. Hence, SBI methods can be used to relate observations to the most likely set of parameters that generated them, independently of the degree of knowledge of the underlying model. Under the assumption that observations can be simulated for a given set of parameters, SBI methods rely on artificial neural networks to

approximate the posterior (or likelihood) distributions from the joint distribution of data and parameters. Once either of these quantities are defined, the parameters associated with a given observation can be determined probabilistically within the usual Bayesian formalism by sampling the posterior distribution.

Previous applications of SBI in Astrophysics include constraining cosmological parameters from galaxy cluster properties (Hahn et al. 2023; Hernández-Martínez et al. 2025) and weak lensing maps (Jeffrey, Alsing & Lanusse 2021), spectral energy distribution (SED) fitting of galaxies (Hahn & Melchior 2022) and estimating compact objects parameters from gravitational-wave data (Green, Simpson & Gair 2020). In the context of estimating the properties of merging galaxies, SBI has recently been used by Widmark & Johnston (2025) to develop a framework for inferring the mass density and orbital parameters of tidally perturbed dwarf galaxies based on the observed position and line-of-sight velocity field.

More closely related to the scope of this work, Viterbo & Buck (2024) have developed the Chemical Abundance Simulation Based Inference (CASBI) model, which uses the SBI formalism to predict the stellar mass and infall time of the disrupted components of the stellar halo of Milky Way-mass galaxies based on the oxygen and iron abundances of the debris. Following the idea proposed by Cunningham et al. (2022) that the chemical abundance ratio distributions (CARDS) of stellar haloes can be interpreted as the combination of the CARDS of the progenitor galaxies, CASBI was trained and tested on a data set of artificial haloes made up from a catalogue of dwarf galaxies from the Numerical Investigation of a Hundred Astronomical Objects (NIHAO) suite of cosmological simulations (Wang et al. 2015; Buck 2020). CASBI achieves remarkable accuracy in estimating the stellar masses of the dwarf galaxy components in artificial haloes. However, it is less clear what its performance is in estimating their times of infall.

Here, we consider a different approach by developing an SBI pipeline directly on the merger histories of galaxies within the Auriga suite of cosmological simulations (Grand et al. 2017, 2024). We also investigate the implications of using the dynamical and/or chemical properties of the debris to inform the inference.

An overview of the Auriga simulations, as well as a description of the procedure to extract and characterize the merger events in the suite, is given in Section 2. A general description of SBI and details about its implementation in the GalactiKit models are included in Section 3. The results of the analysis are presented in Section 4, while potential applications of the models are proposed in Section 5. We summarize our conclusions in Section 6.

2 THE AURIGA SIMULATIONS

The data we use to train Galactikit are taken from the Auriga suite of simulations (Grand et al. 2017, 2024). These are cosmological magnetohydrodynamical (MHD) zoom-in simulations of (relatively isolated) Milky Way-mass haloes run with the AREPO moving mesh code (Springel 2010).

The haloes were selected from the L100N1504 EAGLE dark matter only simulation (Schaye et al. 2015), which was run in a periodic cube of comoving side length of 100 cMpc, with a dark matter particle resolution of $1.15 \times 10^7 M_\odot$. The simulation adopts a Λ cold dark matter (Λ CDM) cosmology with model parameters taken from Planck Collaboration XVI (2014).

The initial conditions for the re-simulation of a given Auriga halo are prepared at $z = 127$. The high-resolution region is defined in the Lagrangian region around the halo, outside of which the resolution drops off with increasing distance. Baryons are added by splitting

each dark matter particle into a dark matter particle–gas cell pair with masses given by the cosmological abundance of baryons. The Auriga suite comprises haloes simulated primarily at two resolution levels, which are typically referred to as ‘level 4’ and ‘level 3’. In this work, we consider the ‘level 4’ simulations, which comprise 39 Milky Way analogues in the high ($30, 1 < M_{200}/[10^{12} M_{\odot}] < 2$) and low ($9, 0.5 < M_{200}/[10^{12} M_{\odot}] < 1$) end of the estimated virial mass range for the Milky Way (Wang, Hammer & Yang 2022). The typical particle mass resolution is $\sim 3 \times 10^5$ and $\sim 5 \times 10^4 M_{\odot}$ for dark and baryonic matter, respectively.

The Auriga physics model is described in Grand et al. (2017) and briefly summarized herein. Magnetic fields are seeded at the beginning of the simulations in a single direction and with a comoving strength of 10^{-14} G. Re-ionization is modelled with a spatially uniform, time-dependent ultraviolet background radiation field (Faucher-Giguère et al. 2009) that completes at $z = 6$. Primordial and metal-line cooling of gas with self-shielding corrections are also implemented. As gas collapses to densities larger than 0.11 particles per cubic centimetre it enters a sub-grid model for star formation: such gas is modelled as a two-phase interstellar medium, comprised of a cold and hot phase, with an effective equation of state (see Springel & Hernquist 2003, for a detailed description). Star particles are then created stochastically following a time-dependent, exponential probability function. Each particle represents a single stellar population, with a defined age and metallicity. A Chabrier (2003) initial mass function is assumed for the stellar evolution of the particle. The chemical enrichment of surrounding gas cells comes from asymptotic giant stars, Type II supernovae, and Type Ia supernovae. Once haloes surpass a mass of $5 \times 10^{10} M_{\odot}$, black holes are seeded with a mass of $5 \times 10^5 M_{\odot}$ at the gravitational potential minimum of the halo, and proceed to grow from gas accretion and merger processes following the prescription in Springel, Di Matteo & Hernquist (2005a). Energetic feedback is modelled phenomenologically as supernovae-driven galactic winds, as well as radiative and thermal energy injection from active galactic nuclei radio and quasar modes.

The Auriga galaxies have rotationally supported discs, flat rotation curves and are consistent with the expected stellar mass–halo mass, mass–metallicity and star formation rate relations (Grand et al. 2017). Moreover, the Auriga galaxies were directly compared to the observational counterparts from the GHOSTS (Galaxy Halos, Outer disks, Substructure, Thick disks, and Star clusters) survey (Radburn-Smith et al. 2011; Monachesi et al. 2016) by Monachesi et al. (2019), who showed the stellar haloes in Auriga reproduce the scatter in stellar mass, surface brightness, and metallicity profiles of Milky Way-mass galaxies. However, the Auriga galaxies were also found to have more massive stellar haloes due to an extended *in situ* component, which is discarded in this analysis.

2.1 Extracting mergers information

The $z = 0$ stellar composition of the Auriga galaxies is the result of *in situ* star formation and merger accretion. For each simulation, the evolutionary histories of all galaxies in the box are encoded in the merger trees, which were computed in post-processing following the LHaloTree algorithm (Springel et al. 2005b). At first, the SUBFIND group finding algorithm (Springel et al. 2001) is used to identify haloes and sub-haloes in all simulation outputs (snapshots). Then, each halo (progenitor) is matched to all the ones having common particles at the subsequent snapshot; a unique descendant is chosen by selecting the halo with the highest number of particles in common weighted by their binding energy. A merger occurs when two or more progenitors point to the same descendant.

From the merger trees, we can infer the stellar masses well and structural properties of galaxies at any given time. As part of the public data release of the Auriga Project (Grand et al. 2024), the accretion history of the star particles within R_{200} ¹ of the main haloes, as computed by the merger trees, is reported in the ‘Accreted particle lists’ catalogues.

There are two possible approaches for sorting the accreted particles in their systems of ‘origin’, hence defining the progenitors of the main halo: (i) using the `RootIndex` list, which associates each particle to the index in the merger tree of the progenitor halo which was bound to at formation; and (ii) with the `PeakMassIndex` list, which matches a particle to the merger tree index of the progenitor halo which was bound to at the time the progenitor reached its maximum stellar mass. We decide to consider the second definition because we are interested in retrieving the properties of a progenitor when it first interacts with the main halo, hereby considered as when it first crossed R_{200} , which often coincides or shortly follows the time the progenitor undergoes a quenching in star formation (Kawata & Mulchaey 2008; Simpson et al. 2018; Font et al. 2022).

We restricted our investigation to the most significant progenitors in the assembly history of each galaxy by selecting only those contributing with at least 100 star particles. This corresponds to a minimum stellar mass cut-off of $\sim 10^6 M_{\odot}$ for the satellites at the time of accretion, with most simulations having between 10 and 30 merger events meeting this criterion. In the whole suite, a total of 928 progenitors were identified, of which 792 are taken from the merger trees of the high-mass Auriga simulations ($1 < M_{200}/[10^{12} M_{\odot}] < 2$) and 136 from the low mass ones ($0.5 < M_{200}/[10^{12} M_{\odot}] < 1$). The majority of them (648) are completely disrupted at $z = 0$, and appear as accreted stars; the rest of the stars are from existing satellites (280) orbiting within R_{200} .

An example of the separation of accreted star particles into the progenitor galaxies in which they formed is reported in Fig. 1, which shows – for the Au21 simulation – the $z = 0$ distribution of the accreted star particles within R_{200} in the integrals-of-motion (top) and the alpha-iron (bottom) planes. The former is described by the total energy and the z -component (perpendicular to the plane of the disc) of angular momentum of particles, while the latter is defined by the iron-to-hydrogen, $[\text{Fe}/\text{H}]$, and alpha-to-iron, $[\alpha/\text{Fe}]$, chemical abundance ratios. In the leftmost column, plots show star particles colour-coded by the progenitor galaxy in which they formed based on their `PeakMassIndex` label. In the middle column, the star particles are coloured depending on whether their system of origin is disrupted (grey) or appears as an existing satellite (black) at $z = 0$. In the rightmost column, the star particles are coloured based on the lookback infall time at which their progenitor of origin (as defined by the associated `PeakMassIndex`) first crossed R_{200} . Although 30 progenitor galaxies were identified for Au21, the majority of accreted stars come from only a few systems, as can be visually distinguished in the plots on the leftmost column. Particles from early accretion events tend to have higher abundance of α -elements and are more dispersed in the lower region of the integrals-of-motion plane than their later accreted counterparts.

3 SIMULATION-BASED INFERENCE

The underlying idea of the GalactiKit methodology is that some properties of a merging galaxy – such as the time of infall or the halo mass – shape the present-day chemodynamical distribution of the

¹Galactocentric radius enclosing a mass density 200 times higher than the critical density of the universe, often taken as a proxy of the virial radius.

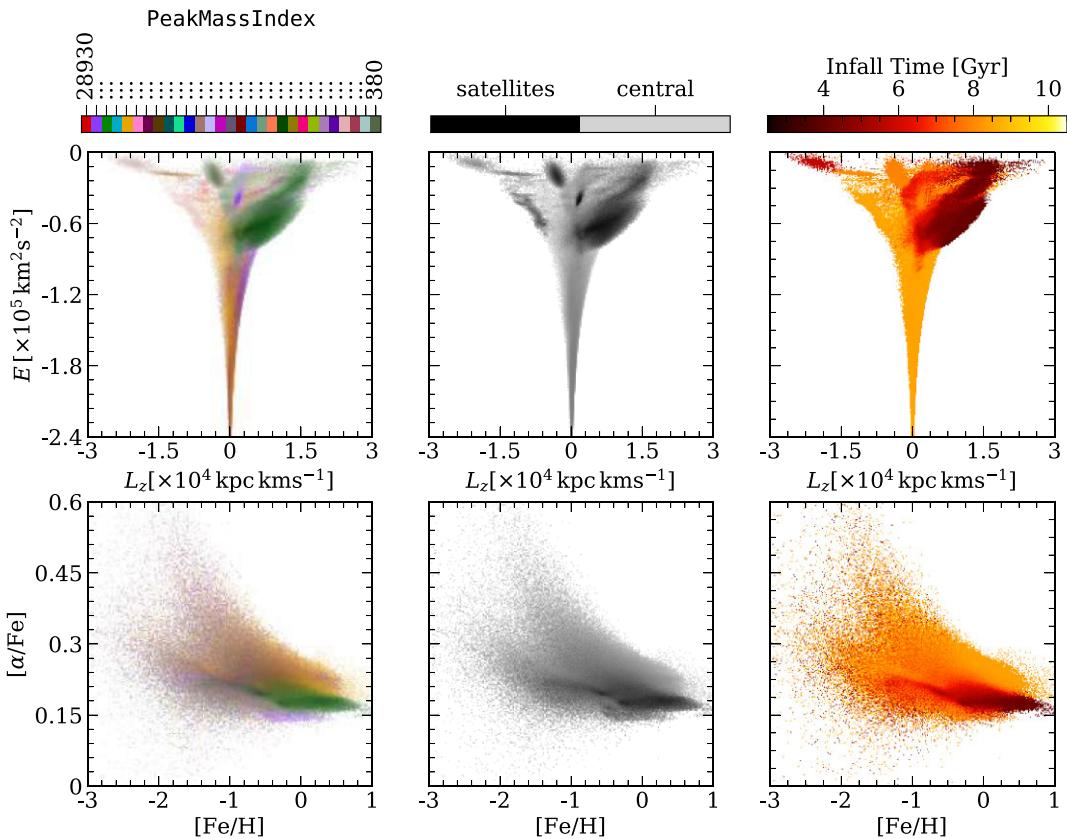


Figure 1. Accreted star particles located within R_{200} in the Au21 simulation at $z = 0$. The top row shows the integrals-of-motion distribution of the particles, while the $[\text{Fe}/\text{H}] - [\alpha/\text{Fe}]$ plane is reported in the bottom. For the plots on the left, the same colour is given to particles coming from each of the 30 progenitor galaxies contributing with at least 100 star particles. Only the `PeakMassIndex` of two progenitor galaxies is explicitly shown in the figure, while ellipses are reported for the colours corresponding to the remaining progenitors. In the plots in the middle column, star particles are colour-coded differently if their progenitors are disrupted in the main halo (black) or are orbiting sub-haloes within R_{200} (grey). In the rightmost column, the plots show the distribution of star particles colour-coded based on the lookback time at which their associated progenitor galaxy crossed R_{200} for the first time.

corresponding debris. Hence, we assume there is an undefined model \mathcal{M} which, given the vector of properties of a satellite at accretion $\boldsymbol{\theta}$, predicts the stellar properties \mathbf{x} of the corresponding debris at $z = 0$. An analytical definition of \mathcal{M} is practically challenging as it should account for complex, coupled physical processes such as the chemical evolution and mass assembly of the host and satellite galaxies in a cosmological context and the phase-space mixing of the accreted debris through dynamical friction within the evolving gravitational potential of the host. However, most of the relevant galaxy formation processes are currently described in the Auriga cosmological simulations at a level of detail that allows for direct comparison between simulated galaxies and their observational counterparts (Monachesi et al. 2019). Therefore, simulations can be thought to already encode a realistic representation of \mathcal{M} . In order to extract the information encoded in the simulations and make predictions on the properties of a galaxy at infall from the $z = 0$ properties of its debris, GalactiKit adopts the SBI framework, which provides a Bayesian description of \mathcal{M} .

In the SBI formalism, given a set of parameters $\boldsymbol{\theta}$ and observations \mathbf{x} from an undefined model or with an intractable likelihood function, if the data generation process can be simulated, the posterior distribution $p(\boldsymbol{\theta} | \mathbf{x})$ can be evaluated by either:

(i) iteratively comparing the observed data with simulations of the model run with different parameters until a similarity threshold

is met. This technique is called approximate Bayesian computation (ABC; Rubin 1984); and

(ii) approximating the likelihood or posterior distributions with density estimation techniques.

Once the posterior distribution is known, inference on the model parameters for a given observation is performed by sampling the posterior distribution and computing summary statistics.

Because ABC requires simulations to be run at arbitrary points in the $\boldsymbol{\theta}$ -parameter space and the sample of merger events $\{\mathbf{x}, \boldsymbol{\theta}\}$ in the Auriga simulations is limited by the assembly history of the simulated haloes, we adopt approach (ii) and implement SBI by developing a deep learning model to approximate the posterior distribution $p(\boldsymbol{\theta} | \mathbf{x})$.

In this analysis, we focus on the following properties $\boldsymbol{\theta}$ of the progenitor galaxies:

- (i) infall time (τ), defined as the lookback time at which the satellite first crossed R_{200} ;
- (ii) stellar mass and halo mass (M_* and M , respectively) at infall; and
- (iii) the merger mass ratio (MMR), defined as the ratio between the halo mass of the satellite and that of the main galaxy at infall.

To ensure a strong correlation between \mathbf{x} and $\boldsymbol{\theta}$, we consider the properties of the debris that are known to be physically related to the merger variables:

- (i) total energy (E), which is related to the infall time through the infall time–binding energy relation of sub-haloes and debris (Rocha, Peter & Bullock 2012; García-Bethencourt et al. 2023);
- (ii) total specific angular momentum (L) which, like the total energy or the radial action, is an integral of motion for a star in a spherical potential; and
- (iii) [Fe/H], which, as a proxy of metallicity, is correlated to the mass of a galaxy through the mass–metallicity relation (Harmsen et al. 2017).
- (iv) $[\alpha/\text{Fe}]$, which relates to the infall time of the progenitor as stars from early accreted galaxies have, in general, higher abundance of α -elements compared to stars from late mergers (Font et al. 2006).

The posterior estimation is performed with the masked autoregressive flow (MAF; Papamakarios, Pavlakou & Murray 2017) technique. Briefly, MAFs are generative models that approximate complex probability distributions defined by the training data. By drawing a random number from a normal distribution, MAFs generate a new data point through a series of parametrized bijective transformations; this is equivalent to sampling from the approximated probability distribution of the training data. The models are trained by tuning the parameters of the transformations such that the negative log-probability of the training data is minimized.

A more detailed mathematical description of the SBI framework and MAFs is reported in Section 3.1, while details on the implementation and training of the model are presented in Section 3.2. The reader who is not interested in the detailed methodology can skip these sections and refer directly to the results of the analysis in Section 4.

An overview of the GalactiKit methodology is shown in Fig. 2. Initially, the Auriga simulations were separated into a training set, used to train MAF models that estimate $p(\boldsymbol{\theta} | \mathbf{x})$, and a test set, used for validation. For simulations in both sets, the information from the merger tree was used to define a set of merger–debris pairs $\{\mathbf{x}, \boldsymbol{\theta}\}$ by grouping the accreted stars orbiting within R_{200} of the main halo at $z = 0$ into the corresponding progenitor galaxies as defined in Section 2.1. The resulting training and test datasets were then processed (see Section 3.2) to become valid inputs for the MAF models. The data used to train the model comprised a variable number of samples of 100 stars drawn randomly from each progenitor in the training simulations. The number of samples was calculated considering the total number of stars belonging to each progenitor, such that the debris distribution of larger mergers is sampled more times. Similarly, the test data consisted of samples of 100 stars drawn randomly from the progenitors in the test galaxy. However, only one sample was drawn per merger to avoid results being biased toward major mergers. After training, the quality of the estimation of $p(\boldsymbol{\theta} | \mathbf{x})$ was determined by comparing the actual properties of the mergers in the test galaxy with the ones inferred by sampling the MAF models. To obtain a statistical representation, 1000 $\boldsymbol{\theta}$ -samples per merger were generated using the debris data as input to the models. To assess how the model performs on unseen data, we performed a leave-one-out cross-validation, where each of the 39 haloes was excluded one at a time from the training set and used for testing. Hence, for each iteration of the cross-validation procedure, the model was trained on the merger–debris pairs of the 38 simulated galaxies in the training set and tested on the merger events of the galaxy left out in the test set. This ensures that the model was tested on unseen data, while its evaluation still accounts for all the mergers in the simulations.

3.1 Density estimation with masked autoregressive flows

Consider the model \mathcal{M} that relates the properties of a satellite galaxy at the time of infall ($\boldsymbol{\theta}$) with the chemodynamical properties at $z = 0$ of the corresponding debris (\mathbf{x}). Applying Bayes’ theorem, the degree of belief that a certain $\boldsymbol{\theta}$ leads to a realization \mathbf{x} of the model can be measured by evaluating the posterior probability

$$p(\boldsymbol{\theta} | \mathbf{x}) = \frac{p(\mathbf{x} | \boldsymbol{\theta}) p(\boldsymbol{\theta})}{\int p(\mathbf{x} | \boldsymbol{\theta}) d\boldsymbol{\theta}}, \quad (1)$$

where $p(\boldsymbol{\theta})$ is the prior, encoding the data-independent information on the possible merger parameters, and $p(\mathbf{x} | \boldsymbol{\theta})$ is the likelihood, measuring the probability of observing a certain configuration of the merger debris distribution given certain merger parameters. Hence, for an observed distribution of merger debris \mathbf{x}_{obs} , the most likely set of properties of the corresponding progenitor galaxy $\boldsymbol{\theta}$ can be inferred by sampling the posterior probability $p(\boldsymbol{\theta} | \mathbf{x}_{\text{obs}})$. However, as the likelihood is unknown, the posterior cannot be evaluated explicitly. This problem can be overcome adopting the SBI framework, where simulations of the model \mathcal{M} are used to obtain an approximation of the posterior distribution. In particular, the posterior distribution is modelled as a distribution function whose parameter values depends on the vector of debris properties \mathbf{x}_{obs} used for conditioning the posterior, i.e.

$$p(\boldsymbol{\theta} | \mathbf{x}_{\text{obs}}) \approx p_{\boldsymbol{\phi}}(\boldsymbol{\theta} | \mathbf{x}_{\text{obs}}), \quad (2)$$

where $\boldsymbol{\phi} = h(\mathbf{x}_{\text{obs}})$ is the vector of parameters of the *variational* distribution $p_{\boldsymbol{\phi}}$, and the function h represents the mapping between the observations of the model \mathbf{x} and the $\boldsymbol{\phi}$ parameters. Because only the internal parameters change depending on the conditioning input, $p_{\boldsymbol{\phi}}$ is the same distribution function for any value of \mathbf{x}_{obs} ; hence, the posterior distribution is *amortized*, i.e. once the mapping between distribution parameters and data is defined, inference on the merger properties $\boldsymbol{\theta}$ can be performed for any value of the debris properties \mathbf{x} .

In order to implement the posterior approximation described in equation (2), we use MAF, a neural network technique which is based on the idea that generating data with an autoregressive model is equivalent to a normalizing flow (Kingma et al. 2016). A brief description of the mathematical formalism of MAF applied to the inference of merger parameters is presented below; we refer to Papamakarios et al. (2017) for details on the algorithm.

Autoregressive models (Uria et al. 2016) are artificial neural networks used for density estimation, and are based on the idea that any distribution in D dimensions can be represented using the probability product rule as a combination of the conditional probabilities in the single dimensions. Hence, the posterior distribution $p(\boldsymbol{\theta} | \mathbf{x})$ of a merger defined by a galaxy whose properties are $\boldsymbol{\theta} = (\tau, M_*, M, \text{MMR})$ and debris with a certain vector of properties \mathbf{x} at $z = 0$, can be defined as:

$$p(\boldsymbol{\theta} | \mathbf{x}) = \prod_{d=1}^{D=4} p(\theta_d | \boldsymbol{\theta}_{<d}, \mathbf{x}), \quad (3)$$

where $\boldsymbol{\theta}_{<d} = (\theta_1, \dots, \theta_{d-1})$ is a vector whose components are the first $d - 1$ components of the four-dimensional merger parameter vector $\boldsymbol{\theta}$.

In MAF, the conditional probabilities are modelled as single Gaussians whose parameters are computed using Masked Autoencoders for Distribution Estimation (MADE; Germain et al. 2015) models:

$$p(\theta_d | \boldsymbol{\theta}_{<d}, \mathbf{x}) \approx p_{\phi_d}(\theta_d | \boldsymbol{\theta}_{<d}, \mathbf{x}) = \mathcal{N}(\theta_d | \mu_d, (\exp \alpha_d)^2), \quad (4)$$

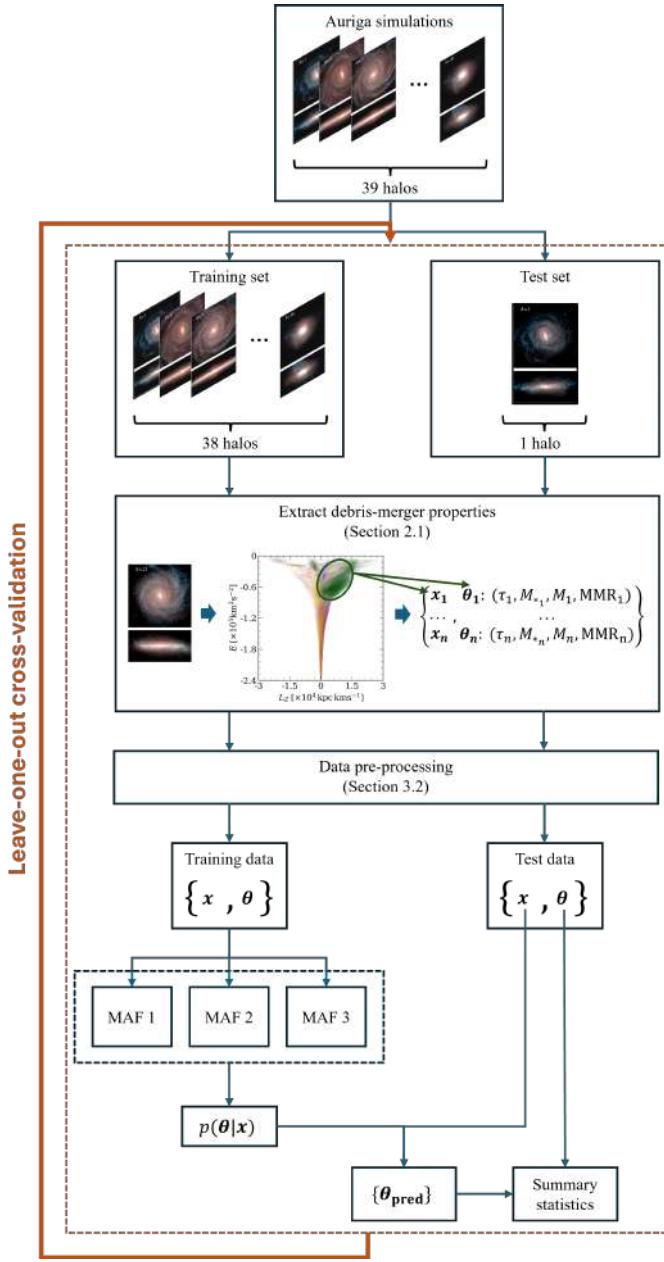


Figure 2. Overview of the GalactiKit methodology. The Auriga simulations were first split into a training and test sets. For both cases, the accreted stars within R_{200} of the main halo were associated to the corresponding progenitor galaxy creating a set of merger–debris pairs $\{\mathbf{x}, \boldsymbol{\theta}\}$; this was then pre-processed to be used as input to the MAF models approximating $p(\boldsymbol{\theta} | \mathbf{x})$. The quality of the estimation was assessed through summary statistics comparing the inferred values to the actual ones for the properties of mergers in the test. To ensure robust performance assessment, the procedure was repeated while systematically leaving out one Auriga simulation for testing at a time. The images of the Auriga galaxies are taken from Grand et al. (2017).

with $\mu_d = h_{\mu_d}(\boldsymbol{\theta}_{<d}, \mathbf{x})$ and $\alpha_d = h_{\alpha_d}(\boldsymbol{\theta}_{<d}, \mathbf{x})$, where h_{μ_d} and h_{α_d} represent the MADE outputs as the mean and log-standard deviation of the d th conditional probability. Hence, the $p(\theta_d | \boldsymbol{\theta}_{<d}, \mathbf{x})$ conditionals are approximated by Gaussian variational distributions whose parameters, $\phi_d = \{\mu_d, \alpha_d\}$, are computed from an artificial neural network which takes as input the $d - 1$ previous components of $\boldsymbol{\theta}$ and the conditions \mathbf{x} . In this framework, sampling is equivalent to generating a new point in the $\boldsymbol{\theta}$ -parameter space by explicitly computing each dimension as

$$\theta_d = u_d \exp \alpha_d + \mu_d, \quad (5)$$

where $u_d \sim \mathcal{N}(0, 1)$. Hence, equivalently to a normalizing flow (Papamakarios et al. 2021), the model is mapping a point $\mathbf{u} = (u_1, \dots, u_d)$ in the base distribution to the merger parameter space through an invertible and differentiable transformation: $\boldsymbol{\theta} = f(\mathbf{u})$. Hence, the variational distribution approximating the posterior can be expressed in terms of the base distribution p_u as

$$p_{\phi}(\boldsymbol{\theta} | \mathbf{x}) = p_u(f^{-1}(\boldsymbol{\theta})) \left| \frac{\partial f^{-1}}{\partial \boldsymbol{\theta}} \right|, \quad (6)$$

where $\left| \frac{\partial f^{-1}}{\partial \boldsymbol{\theta}} \right|$ is the determinant of the Jacobian matrix of f^{-1} , which for the transformation defined in equation (5) is $\left| \frac{\partial f^{-1}}{\partial \boldsymbol{\theta}} \right| =$

$\exp\left(-\sum_d^{D=4} \alpha_d\right)$. Considering the model as a normalizing flow, it is possible to increase the expressivity of MAF by stacking multiple instances as a sequence of transformations. Thus, a sample of merger parameters generated by a MAF model with N transformations is computed as

$$\boldsymbol{\theta} = \mathcal{F}(\mathbf{u}), \text{ with } \mathcal{F} = f_1 \circ \dots \circ f_N \text{ and } \mathbf{u} \sim \mathcal{N}(0, \mathbf{I}), \quad (7)$$

where \mathbf{I} is the identity matrix and each f represents a single transformation, which includes an ensemble of Gaussian conditionals and the MADE model, which take the vector of debris properties \mathbf{x} and the relevant components of $\boldsymbol{\theta}$ as inputs and return the corresponding mean and log-standard deviation. Because a combinations of invertible functions is also invertible and the overall Jacobian determinant is the product of the Jacobian determinants of the single transformations, equation (6) can be adapted to the case of an arbitrary number of transformation defined in equation (7), and written in log-form as

$$\log p_{\boldsymbol{\theta}}(\boldsymbol{\theta} | \mathbf{x}) = \log p_{\mathbf{u}}(\mathbf{u}) + \sum_{i=1}^N \log \left| \frac{\partial f_i^{-1}}{\partial \boldsymbol{\theta}_i} \right|, \quad (8)$$

where $\boldsymbol{\theta}_i = f_i \circ \dots \circ f_1(\mathbf{u})$ is an intermediate state of the flow \mathcal{F} relating the base distribution to the one of the merger properties $\boldsymbol{\theta}$. During the training routine, the internal parameters of the MADE models, which are used to compute the vectors $\boldsymbol{\mu}_i$ and $\boldsymbol{\alpha}_i$ defining each transformation f_i , are updated in order to maximize the total log likelihood $\sum_n \log p_{\boldsymbol{\theta}}(\boldsymbol{\theta}_n | \mathbf{x}_n)$ of the training data, where $\mathbf{x}_n, \boldsymbol{\theta}_n \sim p(\boldsymbol{\theta}, \mathbf{x})$, which is the joint probability distribution of merger and debris properties defined by the progenitor galaxies in the Auriga simulations.

Once the model has been trained, a new sample of the posterior distribution of the progenitor parameters given a certain observation of debris properties at $z = 0$ can be obtained by drawing a random number from a Gaussian distribution and applying the series of data-dependent transformations specified by the MADE models as described in equation (7).

3.2 Data pre-processing and training implementation

Before being used to train the MAF model, the progenitors $\boldsymbol{\theta}$ and merger debris \mathbf{x} properties were processed as follows:

(i) The distributions of the debris in E and L were transformed by taking the square root of the corresponding absolute values, to compensate for the skewness of the original distributions. We also considered scaling the distributions of the debris based on the structural, present-day properties of the galaxies; namely, the maximum rotational velocity, scale radius, R_{200} and M_{200} . However, we find that this leads to little or no improvements in the posterior estimation.

(ii) Outliers from the interquartile (i.e. 25th and 75th percentile) ranges of the debris distributions in E , L , [Fe/H], and $[\alpha/\text{Fe}]$ were detected and discarded. This is performed in preparation of the next step, as particles whose properties lie at the margin of the distributions might skew the training data while not representing the majority of the debris from a given merger. After this procedure, all merger events that fell below the 100 star threshold were also removed from the analysis. To ensure all inputs of the models are on the same scale, the resulting distributions were standardized by subtracting the mean and dividing by the standard deviation in each dimension.

(iii) Because the number of particles representing the merger debris depends on the progenitor mass, while MAF requires an input of fixed dimensionality, N samples of 100 stars were drawn randomly

from each debris distribution, allowing for stars to potentially be included in multiple samples. In order to account for the intrinsic variability in \mathbf{x} when sampling large progenitors, the number of samples for a given merger was calculated in terms of the number of associated accreted stars S , as $N = \min(10, \lfloor S/100 \rfloor)$.

(iv) To provide a balanced representation of all mergers in the training set, the debris distributions of mergers with uncommon parameters were oversampled. For this purpose, a k-means model was applied to the $\boldsymbol{\theta}$ -parameter space to form 20 groups with similar properties. Then, the number of samples from mergers in each group was counted. By imposing that each group has to contribute the same number of samples, a balanced data set was achieved by drawing more samples from the debris of the mergers from underrepresented groups, with each contributing with an equal number of extra samples.

The above process returns a data set composed of merger-debris pairs which more uniformly covers the $\boldsymbol{\theta}$ -parameter space, is normally distributed in \mathbf{x} , and has homogeneous samples (i.e. all sharing the same dimensionality $D \times 100$, where D is the number of stellar properties considered).

The posterior distribution $p(\boldsymbol{\theta} | \mathbf{x})$ was modelled with an ensemble of three MAFs to avoid overfitting, which is more likely to affect a single density estimator. All the MAF models share the same initial configuration (i.e. number of transformations, number of neurons in the hidden layers, etc.), but are trained separately; because the optimization of the internal parameters of an artificial neural network is a stochastic process, the resulting three MAFs are effectively three different models of the posterior distribution. The implementation of the SBI framework was performed using the Learning the Universe Implicit Likelihood Inference (LtU-ILI) pipeline (Ho et al. 2024) with the `sbi` backend. The number of transformations in each MAF, the length of the hidden layers in the MADE models, and the training parameters (batch size and learning rate) were derived using the `Optuna` hyperparameter optimization framework (Akiba et al. 2019) setting the log-likelihood of a fixed validation data set as objective function. The resulting MAF models used in the analysis comprised 22 MADE layers with 463 neurons per hidden layer and were trained on batches of 2569 examples with a learning rate of 0.001.

4 RESULTS

Three implementation of GalactiKit were developed to estimate the posterior distribution of the model \mathcal{M} described in Section 3 for different combinations of debris properties \mathbf{x} , namely: (a) E, L ; (b) [Fe/H], $[\alpha/\text{Fe}]$; and (c) $E, L, [\text{Fe/H}], [\alpha/\text{Fe}]$. The target merger parameters were fixed to $\boldsymbol{\theta} = (\tau, \log(M_*/M_{\odot}), \log(M/M_{\odot}), \log(\text{MMR}))$ for all cases. The aim was to investigate to what extent the properties of a merger can be reconstructed when different degrees of debris information are available. For each combination of \mathbf{x} , the results of the leave-one-out cross-validation procedure are reported in Fig. 3 as the inferred against the actual parameters for the mergers in the test galaxies. In this paper, ‘true’ indicates the parameter values as directly measured in the Auriga simulations, while ‘predicted’ refers to the values inferred from the estimated posterior distributions. Results from each model are shown in separate rows with estimates from model (a), (b), and (c) reported in the top, middle, and bottom rows, respectively. In the top plots of each row, the points represent the median of the 1000 samples drawn from the inferred posterior distribution of the merger parameters, whereas the 34th–68th percentile ranges (σ_{θ}) are

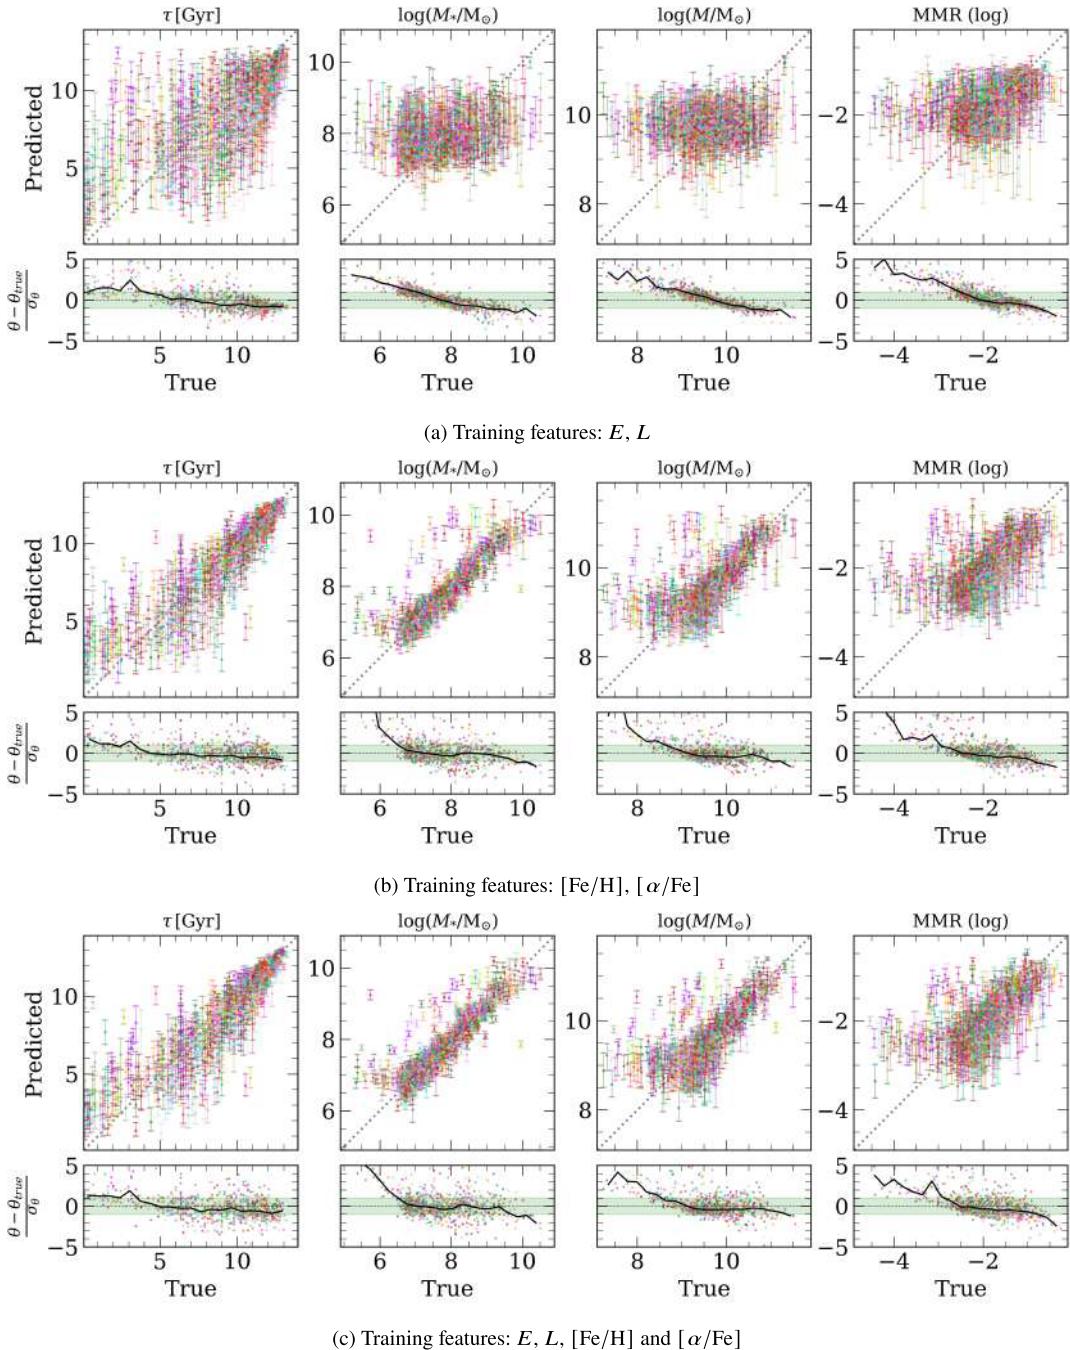


Figure 3. Merger parameters inferred by SBI models trained on different combinations of $z = 0$ debris properties: (a) E, L (top), (b) $[\text{Fe}/\text{H}], [\alpha/\text{Fe}]$ (middle), and (c) $E, L, [\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ (bottom). Each point represents the median value of the parameter samples, while the error bar shows the extent of the 34th–68th percentile range of the distribution. A unique colour is associated to merger events associated to a specific galaxy, for which case the SBI models used for inference were trained on the mergers from the remaining galaxies in the suite.

shown as error bars. The dotted line indicates perfect prediction of the merger parameters. Different colours are associated to mergers from different test galaxies included in the cross-validation procedure. The distribution of the fractional deviations between the predicted and true parameters compared to σ_θ is reported as a sub-panel in each plot. A deviation of $1\sigma_\theta$ from the true parameter is highlighted as a green-shaded region, while the median distribution of the deviations is shown as black continuous line. By visually inspecting Fig. 3, it can be noticed that the colour distribution is uniform in all the plots,

suggesting that the model performances are similar and do not depend on the assembly history of the test galaxy.

As intuitively expected from the infall time–binding energy relation of sub-haloes (Rocha et al. 2012) and conservation principles, the energy and angular momentum distributions of debris appear to contain information on the accretion time of the associated progenitor, while being unrelated to the mass of the system. The energy distribution of debris is strongly dependent on the gravitational potential of the main galaxy; debris from early mergers have had

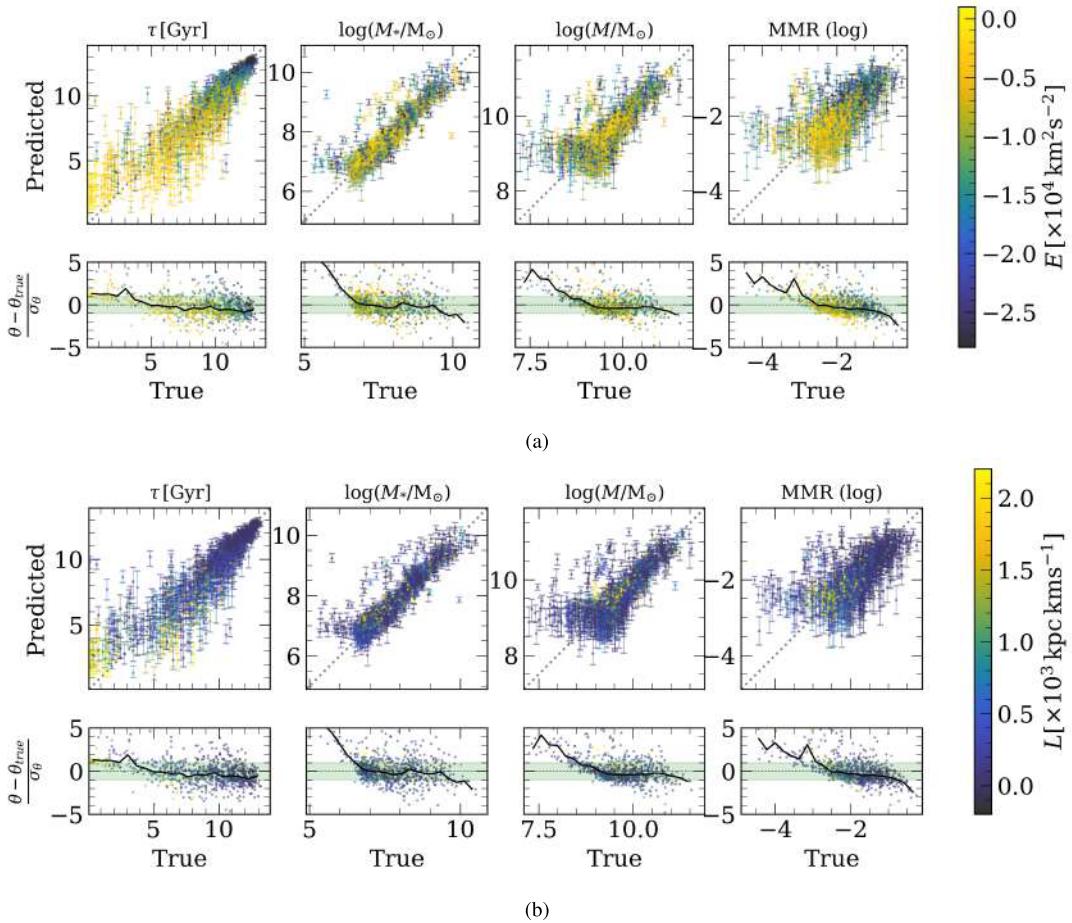


Figure 4. Same as Fig. 3(c), but mergers are colour-coded by the median value of the total energy (a, upper panel) and total angular momentum (b, bottom panel) of the corresponding debris distribution.

the time to complete a few orbits around the galaxy and sink into the bottom of the gravitational potential, hence having a higher binding energy (i.e. lower total energy) than debris from later mergers. Moreover, as the host halo accretes more mass, the gravitational potential increases overtime contributing to shifting early-accreted stars at lower energies. This is shown in Fig. 4(a), which replicates the plots in Fig. 3(c) colour-coding the mergers by the median total energy of the associated debris at $z = 0$.

A drop in the performance of model (a) can be noticed for mergers more recent than five billion years ago; rather than being physically motivated, this limitation is likely caused by a lack of training examples in that region as the infall time merger distribution, shown in the relevant sub-panels, is denser at high τ values. A potential solution could be complementing the current data set with merger events from other cosmological simulations. Alternatively, a machine learning model could be trained as a surrogate of the simulation process (see e.g. Viterbo & Buck 2024) to generate an arbitrary number of training examples in any region of the θ -parameter space.

The gravitational potential of the main galaxy is closely related to its mass density profile, which is not significantly affected by the accretion of satellites (with the exception of major mergers) as galaxies tend to grow ‘inside-out’ through *smooth accretion* (L’Huillier, Combes & Semelin 2012; Pérez et al. 2013). Hence, the energy distribution of debris is not an appropriate proxy for the mass of accreted satellites as shown in Fig. 3(a). The same conclusion can be drawn from Fig. 4, which shows that both the median values

of the debris E and L do not correlate with the mass of the associated mergers.

Interestingly, despite failing at predicting exact values, the energy of debris appears to be broadly informative for distinguishing between massive ($MMR > 1/20$) and less-massive ($MMR < 1/20$) mergers as shown in Fig. 4(a). When limited to this task, model (a) is able to correctly identify 66 of the 139 massive mergers in the suite, with a 0.48 purity.

Providing information on the chemical abundances of debris significantly improves the posterior fitting, as shown by the middle and bottom panels in Fig. 3. The $[Fe/H]$ and $[\alpha/Fe]$ abundance ratios are closely related to the star formation history of the progenitor galaxy, which is affected both by the accretion time, when it stops due to cold gas being stripped away by the main galaxy, and by the mass of the system, as more massive satellites form multiple generations of stars, hence producing debris with higher metallicity values. Fig. 5 shows samples from model (c) trained on the E , L , $[Fe/H]$, and $[\alpha/Fe]$ debris information and colour-coded by the median values of the actual $[Fe/H]$ (Fig. 5a) and $[\alpha/Fe]$ (Fig. 5b) distributions. The $[Fe/H]$ abundance of debris correlates well with the mass of the associated progenitor, which is expected from the mass–metallicity relation, as more massive systems formed multiple generations of stellar populations which contributed to enrich the interstellar medium with metals. However, the mass–metallicity relation presents an intrinsic scatter which depends on galaxy-specific properties such as the gas fraction, inflow, and outflow, or the star formation rate

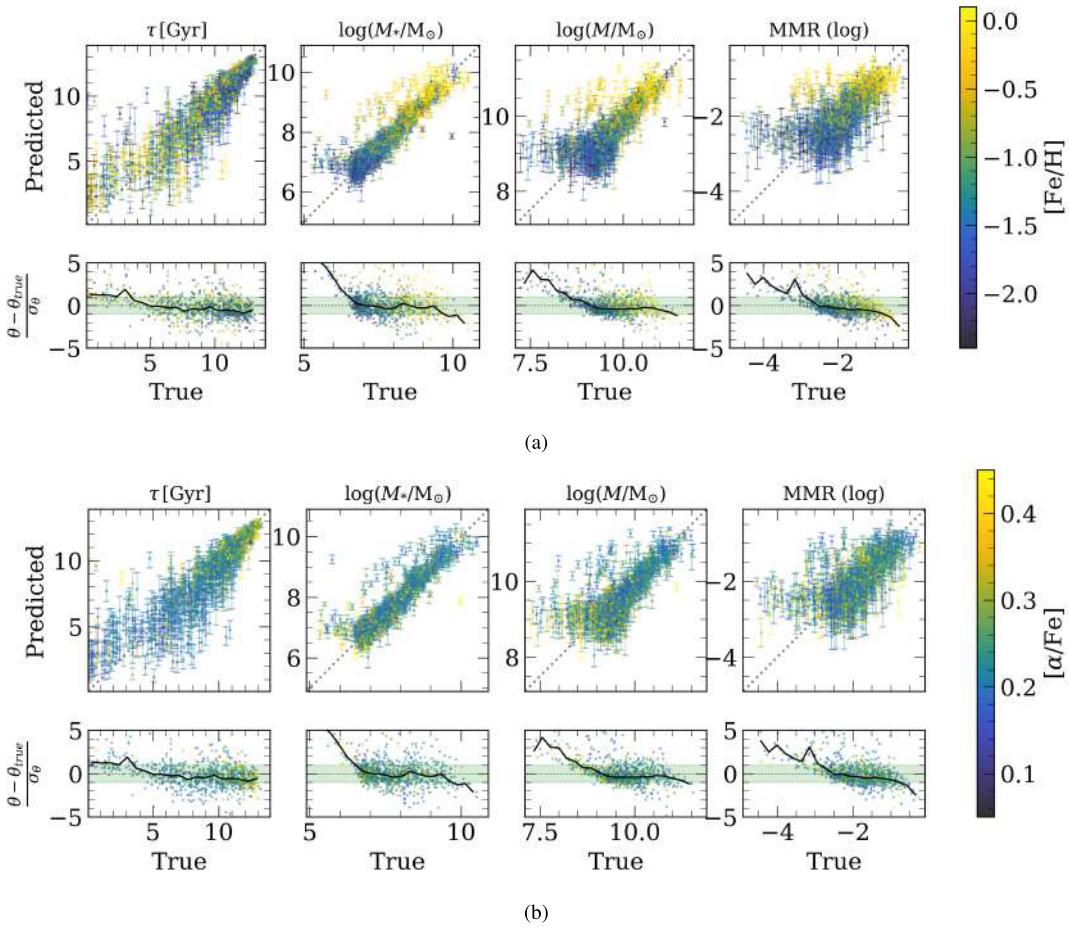


Figure 5. Same as Fig. 3(c), but mergers are colour-coded by the median value of the [Fe/H] (a, upper panel) and $[\alpha/\text{Fe}]$ (b, bottom panel) distributions of the corresponding debris.

(van Loon, Mitchell & Schaye 2021). This could be the cause of the significant deviation between the mass estimates for some metal-rich mergers as can be seen comparing the $\log(M/M_\odot)$ plots in Figs 3(b) and 5(a).

Fig. 5 suggests that the $[\alpha/\text{Fe}]$ ratio appears to correlate better with the accretion time of mergers than [Fe/H]. This could be explained in terms of the time-scales involved in the two enrichment channels and the effect of the gravitational interaction of the host. α -elements are mostly produced in Type II supernovae on a time-scale of few million years, whereas Fe atoms are mostly released during Type Ia supernovae, which require a timescale of \sim billion year (McWilliam 1997). When a satellite is accreted, its star formation stops due to the cold gas being stripped away by the host galaxy. Hence, satellites accreted at late times would appear at a later stage of their chemical evolution with stars characterized by a lower abundance of α -elements compared to stars from early mergers (Font et al. 2006).

The infall time estimates for late mergers ($\tau < 5$ Gyr) obtained from models (b) and (c) trained with [Fe/H] and $[\alpha/\text{Fe}]$ are closer to the actual values than the ones of model (a) trained only on the E , L debris distributions; however, the model accuracy is still not uniform across the whole range of the parameter space. This, alongside with the plateaus at the edges of the predicted $\log(M/M_\odot)$ and $\log(\text{MMR})$ distributions, is probably related to the limited number of training examples in those ranges. A significant improvement is also seen when distinguishing massive mergers with models (b) and

(c) identifying 74 and 88 massive mergers reaching purity values of 0.61 and 0.62, respectively.

A comparison of the performance of the GalactiKit models defined by the three combinations of debris properties can also be performed quantitatively in terms of the root-mean-squared error (RMSE) and mean-relative-uncertainty (MRU). The RMSE is measured for a single merger by directly comparing the inferred values of the parameters θ to the actual θ_{true} as

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^N (\theta_i - \theta_{\text{true}})^2}{N}}, \quad (9)$$

where $N = 1,000$ is the number of i -samples drawn from the approximated posterior distribution of each merger event. The RMSE estimates the accuracy of the model by measuring how close the average prediction of the inferred merger parameters is to the true value.

The MRU is defined comparing the size of a given confidence interval, defined by a percentile range $Q_a(\theta) - Q_b(\theta)$, to the extent of the prior distributions of each dimension of θ . Each prior distribution informs on the probability that a merger property can assume a certain value, hence excluding the ones that are not physical. In this study, the four merger parameters that define the properties of the progenitors at infall are assumed to have uniform prior distributions such that, for a given merger, the true parameters are known to have equal probability to lie within any given point of the corre-

Table 1. Quantitative comparison between the merger parameter estimates inferred through the posterior models trained on different combinations of debris information. From left to right, the columns report: (i) the inferred merger parameter; (ii) the root-mean-squared error (RMSE) between the actual and inferred merger parameter; (iii) the fraction of mergers whose actual parameters are within the 34th–68th percentile range of the samples drawn from the posterior model; (iv) the mean-relative uncertainty (MRU) for the 34th–68th percentile range; (v) and (vi) are the same as (iii) and (iv) but considering the 5th–95th percentile range as confidence interval.

θ	RMSE	within 34th–68th	MRU(34,68)	within 5th–95th	MRU(5,95)
<i>E, L</i>					
τ	3.44	0.28	0.15	0.81	0.56
$\log(M_*/M_\odot)$	1.25	0.28	0.14	0.88	0.47
$\log(M/M_\odot)$	1.00	0.36	0.14	0.91	0.51
MMR(log)	0.94	0.39	0.12	0.91	0.48
[Fe/H], $[\alpha/\text{Fe}]$					
τ	1.92	0.30	0.08	0.88	0.33
$\log(M_*/M_\odot)$	0.58	0.39	0.05	0.92	0.24
$\log(M/M_\odot)$	0.73	0.41	0.08	0.92	0.38
MMR(log)	0.87	0.38	0.10	0.91	0.46
<i>E, L, [Fe/H], $[\alpha/\text{Fe}]$</i>					
τ	1.83	0.29	0.08	0.86	0.30
$\log(M_*/M_\odot)$	0.54	0.35	0.04	0.91	0.22
$\log(M/M_\odot)$	0.69	0.34	0.08	0.90	0.35
MMR(log)	0.81	0.34	0.10	0.89	0.41

sponding ranges, i.e. $\tau / \text{Gyr} \in \{0, 13.5\}$, $\log(M_*/M_\odot) \in \{5, 11\}$, $\log(M/M_\odot) \in \{7, 12\}$, and $\log(\text{MMR}) \in \{-5, 0\}$. Thus, the MRU for the predictions of the properties of a given merger event is calculated as

$$\text{MRU}(a, b) = \left| \frac{Q_a(\theta) - Q_b(\theta)}{\max(\theta_{\text{true}}) - \min(\theta_{\text{true}})} \right|, \quad (10)$$

where $Q_a(\theta)$ and $Q_b(\theta)$ are the values of θ at the a th and b th percentile of the predicted posterior distribution, and $\max(\theta_{\text{true}}) - \min(\theta_{\text{true}})$ refers to the difference between the maximum and minimum values of θ as defined by the prior ranges. The MRU provides an estimate of the precision of the predictions, where MRU values approaching 0 indicate a model that consistently makes similar predictions for a given merger (i.e. the model is confident about the estimated value of θ), whereas MRU values close to 1 suggest that, despite the information of the debris properties, the model is not able to constrain the merger parameters to a specific region of the prior space.

The RMSE and MRU were computed for each of the 928 merger events considering the θ samples used for producing Figs 3, 4, and 5. The average of the RMSE and MRU values are reported in Table 1, alongside with the average fraction of test examples for which the true merger parameters fall within the 34th–68th and 5th–95th percentile ranges of the predicted posterior distributions.

Compared to model (a) informed only through dynamical data, the GalactiKit models (b) and (c) trained including [Fe/H] and $[\alpha/\text{Fe}]$ provide the most accurate predictions of the merger parameters with an average decrease in RMSE of about 45, 55, 29, and 12 per cent for the estimates of τ , $\log(M_*/M_\odot)$, $\log(M/M_\odot)$, and $\log(\text{MMR})$, respectively. As previously discussed, the time estimates of all models are considerably more accurate for merger events occurred earlier than five billion years ago with average RMSE values of 3.30, 1.75, and 1.64 Gyr compared to 4.45, 3.12, and 3.09 Gyr for late mergers as predicted by models (a), (b) and (c), respectively. This would be less problematic for Galactic archaeology studies as significant satellites of the Milky Way with $\tau < 5$ Gyr are expected to be either still orbiting the Galaxy or traceable from stellar streams.

Estimates for the stellar and halo masses of mergers are also more accurate when provided by the models (b) and (c). In both cases, the stellar mass is better predicted than the halo mass because of the intrinsic scatter in the stellar mass–halo mass relation (SMHR), which could be due, for instance, to haloes of same total mass but different concentration forming a different number of stars as more concentrated haloes are associated to earlier formation times (Matthee et al. 2017). The scatter in the SMHR is found to be larger at smaller halo masses, which could also be the cause to the decrease in accuracy for the halo mass estimates of mergers with $\log(M/M_\odot) < 9$.

Table 1 indicates that all models predict less than half of the merger parameters within the 34th–68th percentile range; as the RMSE are reasonably low, this is a sign of overconfidence probably associated to both the reduced variety of merging events in the training set and the repetition of samples from the same progenitors. The small size of the confidence interval compared to the parameter space, as reported by the MRU values, points towards the same conclusion. More conservative estimates of the merging parameters could be obtained taking into account a larger confidence interval, as shown in the last two columns of Table 1; however, there is the risk of obtaining an uninformative analysis as the size of the considered percentile range approaches the extent of the prior ranges of the parameters.

Overall, model (b) trained exclusively on [Fe/H] and $[\alpha/\text{Fe}]$ has a very similar performance compared to model (c) trained including E and L . Thus, the information contained in the chemical abundance ratios of debris appears to be enough for inferring the infall time and mass of the associated merger. The extra information provided by the dynamical quantities seems to contribute to the overconfidence of the model as the decrease in MRU values comes with a slight decrease in the fraction of parameters correctly inferred for both the 34th–68th and 5th–95th percentile ranges. However, the information on the E and L distributions of debris also appears to help constraining the mass predictions of metal-rich, but relatively low-mass progenitor galaxies, as can be noticed comparing the $\log(M_*/M_\odot)$, $\log(M/M_\odot)$, and $\log(\text{MMR})$ plots in Figs 3(b), (c),

and 5(a). This is also reflected with a difference in RMSE (0.54 and 0.58; 0.69 and 0.73; 0.81 and 0.87 dex) between model (c) and (b) as shown in Table 1. Hence, the combination of chemical and dynamical debris properties appears to be preferred for the development of models of the merger parameters, if the observed overconfidence is addressed through changes in the data, e.g. increasing the variety of training examples, or methodology, e.g. changing the technique for the posterior estimation.

5 DISCUSSION

GalactiKit provides a quantitative description of the assembly history of Milky Way-like galaxies by returning predictions within confidence intervals of the infall time, stellar mass, halo mass, and mass ratio with the host of their progenitors. The precision of the prediction depends on the amount of information available on the distribution of the merger debris $z = 0$, with models trained on chemical abundances significantly outperforming those based only on dynamical data. Complementing chemical with dynamical information seems to mostly contribute to reducing the scatter in the mass estimates of metal-rich mergers, while maintaining a similar performance to the models relying on chemical abundances only in the rest of cases.

Hence, there appears to be a strong link between the properties of a galaxy at infall and the chemical abundance distribution of its debris at $z = 0$. This can be motivated by known physical phenomena such as the observed relation between mass and metallicity of galaxies (Harmsen et al. 2017) and the expected difference in α -element abundances expected in late and early accreted satellites (Font et al. 2006; Grizzetti, Font & De Rossi 2024). However, the GalactiKit models leverage the knowledge of these galaxy formation processes encoded in the Auriga chemical evolution model, which is limited by theoretical uncertainties in stellar evolution such as the nucleosynthetic yields of the different enrichment channels or the characteristics of the intermediate mass function. Therefore, potential drops in the predictive ability of the models could be observed if applied to a cosmological context where the chemical evolution of galaxies significantly differs from the one described in Auriga. Testing on a suite of simulations with different chemical evolution prescription could offer a framework to evaluate the potential drop in performance before the application to observations. For instance, it would be interesting to investigate the impact of a galaxy's cosmological environment on the predictive performance of GalactiKit. The haloes in the Auriga simulations have been selected to be relatively 'isolated' at redshift zero; however, the Milky Way is thought to be orbiting within a Local Group (LG) of galaxies including a companion with similar mass (M31) and other hundreds of dwarf galaxies (McConnachie 2012). The High-resolutions Environmental Simulations of The Immediate Area (HESTIA; Libeskind et al. 2020), which are designed to reproduce the properties of the galaxy population in the LG, would represent an interesting testbed for investigating the GalactiKit predictive capacity on environmentally constrained assembly histories. However, we leave this investigation for further studies, as GalactiKit is primarily aimed at inferring the properties of accreted satellites in the Milky Way, which Auriga haloes have been shown to reproduce both in terms of morphology and assembly history (Fattahi et al. 2019; Monachesi et al. 2019).

Another limitation of the GalactiKit methodology is the assumption that accreted stars can be grouped into their progenitors of origin even in galaxies whose merger tree information is not accessible. Based on the idea that stars which formed in the same system share similar orbital and chemical properties, several methods have

been developed to identify accreted sub-structures in the Milky Way with clustering algorithms (e.g. Koppelman et al. 2019; Naidu et al. 2020; Dodd et al. 2023). However, the purity of the resulting merger debris groups is unknown and can be affected by the significant contamination of *in situ* stars (Thomas et al. 2025). Although there are methods to mitigate the presence of the *in situ* population (Sante et al. 2024), it remains to be investigated how well the GalactiKit models can infer the properties of disrupted progenitors based on incomplete or contaminated distributions of merger debris. Similarly, incompleteness could represent an issue if the samples of merger debris are limited to certain parts of their actual chemodynamical distribution. Therefore, we caution against the application of these models on observational data sets in which: (i) there is little confidence in the region of the chemodynamical space in which the merger debris are located; (ii) the overall chemodynamical distribution of stars from the progenitor is believed to not be effectively represented by the considered debris sample; and (iii) there are significant uncertainties on the measurements of the stellar parameters, as they are not taken into account by GalactiKit during the posterior density estimation.

The accuracy of posterior distributions in multidimensional SBI models, i.e. how similar the approximated distributions are to the real ones, can be assessed using the 'Tests of Accuracy with Random Points' (TARP) coverage testing method (Lemos et al. 2023). The test involves comparing the expected coverage probabilities of the posterior to the corresponding credibility levels at which they are evaluated. However, the effectiveness of the TARP test relies on having a large number of test simulations, as more granular regions of the posterior space can be probed. In this study, due to the leave-one-out cross-validation, each test set contains only the merger events relative to one galaxy. Therefore, a meaningful evaluation of the posterior coverage probabilities could not be performed using the TARP test for the three models, which is an important validation step that must be conducted before the model can be confidently applied to real data.

In spite of the limitations mentioned above, the GalactiKit methodology is promising for future applications to the Milky Way. Combined with the information on the merger debris distribution of the accreted substructures already discovered in the Milky Way, GalactiKit can be used to provide a quantitative picture of the formation history of the Galaxy by estimating the mass and infall time of the various accretion events. A similar endeavour was attempted by Kruijssen et al. (2019b), who used the statistical correlations between globular cluster and galaxy properties in the E-MOSAICS simulations (Kruijssen et al. 2019a) to reconstruct the merger tree of the Milky Way from the observed globular cluster population. GalactiKit provides a framework to extend this analysis from the globular cluster population to the millions of single stars observed by *Gaia* and to the present and upcoming spectroscopic surveys, hence depicting a more complete picture of the assembly history of the Milky Way.

A detailed reconstruction of the turbulent past of the Galaxy could be useful for a better understanding of the evolution of the Milky Way in a cosmological context. The Galaxy is considered to have undergone an unusually quiet accretion history within the Λ CDM cosmology framework (Hammer et al. 2007; Evans et al. 2020). Theoretical predictions from the dark matter mass function (Tinker et al. 2008) and the stellar-to-halo mass relation (Brook et al. 2014) suggest that the majority of merger events for a Milky Way-mass galaxy consists of the accretion of low-mass dark haloes (Purcell, Bullock & Zentner 2007), with the majority of stars in the stellar halo coming from only \sim three major (mass ratio $> 1/3$), luminous satellites (Fakhouri, Ma & Boylan-Kolchin 2010). In the case of the

Galaxy, merging events in this mass range correspond to the accretion of the GS/E dwarf galaxy and the Large Magellanic Cloud, which are more than 10 billion years apart. Adding further complexity, Hammer et al. (2024) showed that the accretion time of the GS/E dwarf galaxy, as measured using the age–metallicity relation (AMR) of the globular cluster population identified by Kruijssen et al. (2019b), would have had to be shifted at even earlier times to be consistent with the satellite population of the Milky Way under the infall time–binding energy relation (Rocha et al. 2012). Thus, the application of GalactiKit models could provide a clearer perspective by estimating the accretion times of all the disrupted satellites of the Galaxy, including those without an associated globular cluster population.

Studying the properties of the progenitors of the Milky Way can also be useful outside the Galactic archaeology context. Limberg (2024) suggested that the evidence for an intermediate mass black hole (IMBH) in ω Centauri (Häberle et al. 2024), arguably the old nuclear star cluster of the GS/E galaxy (Limberg et al. 2022), could extend the central black hole mass–host stellar mass relation (Greene, Strader & Ho 2020) to the dwarf galaxies regime. Hence, more IMBHs could potentially be located within the proximity of debris from disrupted progenitors of the Milky Way. Therefore, GalactiKit could be used to estimate the mass of the disrupted progenitors of the Milky Way directly from the distributions of the merger debris, hence informing on the potential presence of undetected IMBHs among the accreted stars. Moreover, if a black hole is detected and linked to the debris of an accretion event, such as the case of the IMBH in ω Centauri, GalactiKit could be used to infer the stellar mass of the progenitor galaxy, which could in turn be used to indirectly infer the black hole mass from the black hole mass–host stellar mass relation.

6 CONCLUSIONS

GalactiKit is a data-driven methodology for determining the look-back infall time, stellar mass, halo mass, and mass ratio with the host of disrupted satellite galaxies of Milky Way analogues. Using simulation-based inference to leverage the information on galaxy formation encoded in the merger trees of the Milky Way-mass haloes from the Auriga cosmological simulations, GalactiKit estimates the posterior distribution of the merger parameters as conditioned by the $z = 0$ distribution of the corresponding debris.

In order to investigate how accurate is the inference of merger properties from various degrees of information on the merger debris, we developed three GalactiKit models considering the following chemodynamical properties of the accreted stars: (a) E and L , (b) $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$, and (c) E , L , $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$. The results of our analysis can be summarized as follows:

(i) The total energy of accreted stars is a tracer of the accretion time of progenitor galaxies (Fig. 3a); this result has already been known for surviving satellites (Rocha et al. 2012) and is extended to phase-mixed sub-structures here. Moreover, energy and angular momentum information can be used for an approximate identification of massive mergers; model (a) correctly predicts a median MMR above 1/20 for 47 per cent of the total massive mergers, with a purity of 0.48.

(ii) Including the chemical information of the debris provides a significant improvement in the accuracy of the estimates for all merger parameters, and is a necessary ingredient for the prediction of the stellar and halo masses of progenitor galaxies at infall (Fig. 3). For both models (b) and (c), stellar masses are better predicted than halo masses. This is probably related to the intrinsic scatter in the

stellar mass–halo mass relation which depends on the chemical and stellar evolution specific to each progenitor galaxy.

(iii) Model (c), informed by both the chemical and dynamical properties of the merger debris, is the most accurate at the estimation of the merger parameters with an average RMSE of 1.83 Gyr, 0.54 dex, 0.69 dex, and 0.81 dex for τ/Gyr , $\log(M_*/\text{M}_\odot)$, $\log(M/\text{M}_\odot)$, and $\log(\text{MMR})$, respectively. Interestingly, as shown in Table 1, predictions of similar accuracy can be achieved with model (b) trained exclusively on the chemical abundance ratios distribution of debris, which is also found to provide slightly more conservative estimates. However, as can be seen comparing the $\log(M_*/\text{M}_\odot)$, $\log(M/\text{M}_\odot)$, and $\log(\text{MMR})$ plots in Figs 3(b), (c), and 5(a), model (b) tends to overestimate the mass of high metallicity and relatively low-mass progenitor galaxies, which is moderated in model (c) by the addition of dynamical information.

(iv) All models show a large scatter in the parameter estimates for mergers that are not well represented in the training set. This is due to the limited number of merger events in the simulations which was not successfully addressed in the data pre-processing step. Potential solutions involve combining data from multiple cosmological simulations or implementing a surrogate model for data generation. A significant difference in accuracy is observed between the infall time predictions of early ($\tau > 5$ Gyr) and late ($\tau < 5$ Gyr) mergers as more accurate predictions are made for the former.

In conclusion, we have shown that GalactiKit is an efficient method in studying the assembly history of Milky Way-like systems. This methodology can be efficiently applied to the analysis of the chemodynamical data of millions of stars observed by current and upcoming surveys in the Milky Way, to better constrain the infall times and the masses of disrupted progenitors.

ACKNOWLEDGEMENTS

We thank the reviewer for their constructive feedback. AS acknowledges a Science Technologies Facilities Council (STFC) PhD studentship at the LIV.INNO Centre for Doctoral Training ‘Innovation in Data Intensive Science’ and funding from The Alan Turing Institute through the ‘Enrichment Scheme Award’. This study made use of Prospero high-performance computing facility at Liverpool John Moores University. This work was also partly supported by the UK’s Science & Technology Facilities Council (STFC grant ST/S000216/1, ST/W001136/1). RJG was supported by 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).

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 and models developed in the analysis are available under reasonable request to the authors.

REFERENCES

Akiba T., Sano S., Yanase T., Ohta T., Koyama M., 2019, Optuna: A Next-generation Hyperparameter Optimization Framework. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining

Arentsen A. et al., 2020, *MNRAS*, 491, L11

Belokurov V., Kravtsov A., 2022, *MNRAS*, 514, 689

Belokurov V., Erkal D., Evans N. W., Koposov S. E., Deason A. J., 2018, *MNRAS*, 478, 611

Bignone L. A., Helmi A., Tissera P. B., 2019, *ApJ*, 883, L5

Binney J., Tremaine S., 2008, *Galactic Dynamics*, 2nd edn., Princeton University Press, Princeton, NJ

Bonaca A., Price-Whelan A. M., 2025, *New Astron. Rev.*, 100, 101713

Brook C. B., Kawata D., Gibson B. K., Flynn C., 2003, *ApJ*, 585, L125

Brook C. B., Di Cintio A., Knebe A., Gottlöber S., Hoffman Y., Yepes G., Garrison-Kimmel S., 2014, *ApJ*, 784, L14

Buck T., 2020, *MNRAS*, 491, 5435

Buder S. et al., 2021, *MNRAS*, 506, 150

Chabrier G., 2003, *PASP*, 115, 763

Chiba M., Beers T. C., 2000, *AJ*, 119, 2843

Ciucă I. et al., 2024, *MNRAS*, 528, L122

Cranmer K., Brehmer J., Louppe G., 2020, *PNAS*, 117, 30055

Cui X.-Q. et al., 2012, *Res. Astron. Astrophys.*, 12, 1197

Cunningham E. C. et al., 2022, *ApJ*, 934, 172

Deason A. J., Belokurov V., 2024, *New Astron. Rev.*, 99, 101706

Dodd E., Callingham T. M., Helmi A., Matsuno T., Ruiz-Lara T., Balbinot E., Lövdal S., 2023, *A&A*, 670, L2

Eggen O. J., Lynden-Bell D., Sandage A. R., 1962, *ApJ*, 136, 748

Evans T. A., Fattahi A., Deason A. J., Frenk C. S., 2020, *MNRAS*, 497, 4311

Fakhouri O., Ma C.-P., Boylan-Kolchin M., 2010, *MNRAS*, 406, 2267

Fattahi A. et al., 2019, *MNRAS*, 484, 4471

Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416

Font A. S., Johnston K. V., Bullock J. S., Robertson B. E., 2006, *ApJ*, 638, 585

Font A. S., McCarthy I. G., Belokurov V., Brown S. T., Stafford S. G., 2022, *MNRAS*, 511, 1544

Freeman K., Bland-Hawthorn J., 2002, *ARA&A*, 40, 487

Gaia Collaboration, 2016, *A&A*, 595, A2

Gaia Collaboration, 2018, *A&A*, 616, A1

Gallart C. et al., 2024, *A&A*, 687, A168

García-Bethencourt G., Brook C. B., Grand R. J. J., Kawata D., 2023, *MNRAS*, 526, 1190

Germain M., Gregor K., Murray I., Laroche H., 2015, in Bach F., Blei D., eds, *Proceedings of Machine Learning Research* Vol. 37, *Proceedings of the 32nd International Conference on Machine Learning*. PMLR, Lille, France, p. 881

Grand R. J. J. et al., 2017, *MNRAS*, 467, 179

Grand R. J. J. et al., 2020, *MNRAS*, 497, 1603

Grand R. J. J., Fragkoudi F., Gómez F. A., Jenkins A., Marinacci F., Pakmor R., Springel V., 2024, *MNRAS*, 532, 1814

Green S. R., Simpson C., Gair J., 2020, *Phys. Rev. D*, 102, 104057

Greene J. E., Strader J., Ho L. C., 2020, *ARA&A*, 58, 257

Grimozzi S. E., Font A. S., De Rossi M. E., 2024, *MNRAS*, 530, 95

Häberle M. et al., 2024, *Nature*, 631, 285

Hahn C., Melchior P., 2022, *ApJ*, 938, 11

Hahn C. et al., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 010

Hammer F., Puech M., Chemin L., Flores H., Lehnert M. D., 2007, *ApJ*, 662, 322

Hammer F. et al., 2024, *A&A*, 692, L1

Harmsen B., Monachesi A., Bell E. F., de Jong R. S., Bailin J., Radburn-Smith D. J., Holwerda B. W., 2017, *MNRAS*, 466, 1491

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

Hernández-Martínez E., Genel S., Villaescusa-Navarro F., Steinwandel U. P., Lee M. E., Lau E. T., Spergel D. N., 2025, *ApJ*, 981, 170

Ho M. et al., 2024, *Open J. Astrophys.*, 7, 54

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

Jeffrey N., Alsing J., Lanusse F., 2021, *MNRAS*, 501, 954

Kawata D., Mulchaey J. S., 2008, *ApJ*, 672, L103

Kawata D., Grand R. J. J., Hunt J. A. S., Ciucă I., 2024, preprint (arXiv:2412.12252)

Kingma D. P., Salimans T., Jozefowicz R., Chen X., Sutskever I., Welling M., 2016, in Lee D., Sugiyama M., Luxburg U., Guyon I., Garnett R., eds, *NeurIPS Proceedings* Vol. 29, *Advances in Neural Information Processing Systems* 29. Curran Associates, Inc.

Knebe A., Gill S. P. D., Kawata D., Gibson B. K., 2005, *MNRAS*, 357, L35

Koppelman H. H., Helmi A., Massari D., Price-Whelan A. M., Starkenburg T. K., 2019, *A&A*, 631, L9

Kruijssen J. M. D., Pfeffer J. L., Crain R. A., Bastian N., 2019a, *MNRAS*, 486, 3134

Kruijssen J. M. D., Pfeffer J. L., Reina-Campos M., Crain R. A., Bastian N., 2019b, *MNRAS*, 486, 3180

L'Huillier B., Combes F., Semelin B., 2012, *A&A*, 544, A68

Lane J. M. M., Bovy J., Mackereth J. T., 2023, *MNRAS*, 526, 1209

Lemos P., Coogan A., Hezaveh Y., Perreault-Levasseur L., 2023, in *Proceedings 40th International Conference on Machine Learning*, Vol. 202. 19256

Libeskind N. I. et al., 2020, *MNRAS*, 498, 2968

Limberg G., 2024, *ApJ*, 977, L2

Limberg G., Souza S. O., Pérez-Villegas A., Rossi S., Perottoni H. D., Santucci R. M., 2022, *ApJ*, 935, 109

Mackereth J. T., Bovy J., 2020, *MNRAS*, 492, 3631

Majewski S. R. et al., 2017, *AJ*, 154, 94

Malhan K., Rix H.-W., 2024, *ApJ*, 964, 104

Malhan K., Ibata R. A., Martin N. F., 2018, *MNRAS*, 481, 3442

Massari D., Koppelman H. H., Helmi A., 2019, *A&A*, 630, L4

Matthee J., Schaye J., Crain R. A., Schaller M., Bower R., Theuns T., 2017, *MNRAS*, 465, 2381

McConnachie A. W., 2012, *AJ*, 144, 4

McWilliam A., 1997, *ARA&A*, 35, 503

Meza A., Navarro J. F., Abadi M. G., Steinmetz M., 2005, *MNRAS*, 359, 93

Miglio A. et al., 2017, *Astron. Nachr.*, 338, 644

Monachesi A., Bell E. F., Radburn-Smith D. J., Bailin J., de Jong R. S., Holwerda B., Streich D., Silverstein G., 2016, *MNRAS*, 457, 1419

Monachesi A. et al., 2019, *MNRAS*, 485, 2589

Montalbán J. et al., 2021, *Nat. Astron.*, 5, 640

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., 2022, preprint (arXiv:2204.09057)

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 Proceedings* Vol. 30, *Advances in Neural Information Processing Systems* 30 (NIPS 2017). Curran Associates, Inc.

Papamakarios G., Nalisnick E., Rezende D. J., Mohamed S., Lakshminarayanan B., 2021, *J. Mach. Learn. Res.*, 22, 1

Pérez E. et al., 2013, *ApJ*, 764, L1

Planck Collaboration XVI, 2014, *A&A*, 571, A16

Purcell C. W., Bullock J. S., Zentner A. R., 2007, *ApJ*, 666, 20

Queiroz A. B. A. et al., 2018, *MNRAS*, 476, 2556

Radburn-Smith D. J. et al., 2011, *ApJS*, 195, 18

Rocha M., Peter A. H. G., Bullock J., 2012, *MNRAS*, 425, 231

Rubin D. B., 1984, *Ann. Statist.*, 12, 1151

Sanders J. L., Das P., 2018, *MNRAS*, 481, 4093

Sante A., Font A. S., Ortega-Martorell S., Olier I., McCarthy I. G., 2024, *MNRAS*, 531, 4363

Schaye J. et al., 2015, *MNRAS*, 446, 521

Searle L., Zinn R., 1978, *ApJ*, 225, 357

Simpson C. M., Grand R. J. J., Gómez F. A., Marinacci F., Pakmor R., Springel V., Campbell D. J. R., Frenk C. S., 2018, *MNRAS*, 478, 548

Springel V., 2010, *MNRAS*, 401, 791

Springel V., Hernquist L., 2003, *MNRAS*, 339, 289

Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726

Springel V., Di Matteo T., Hernquist L., 2005a, *MNRAS*, 361, 776

Springel V. et al., 2005b, *Nature*, 435, 629

Thomas G. F., Battaglia G., Grand R. J. J., Aguiar Álvarez A., 2025, preprint (arXiv:2504.10398)

Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709

Uria B., Côté M.-A., Gregor K., Murray I., Larochelle H., 2016, *JMLR*, 17, 1

van Loon M. L., Mitchell P. D., Schaye J., 2021, *MNRAS*, 504, 4817

Viterbo G., Buck T., 2024, preprint (arXiv:2411.17269)

Wang L., Dutton A. A., Stinson G. S., Macciò A. V., Penzo C., Kang X., Keller B. W., Wadsley J., 2015, *MNRAS*, 454, 83

Wang J., Hammer F., Yang Y., 2022, *MNRAS*, 510, 2242

White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341

Widmark A., Johnston K. V., 2025, *ApJ*, 987, 120

This paper has been typeset from a *TeX/LaTeX* file prepared by the author.