






Optimized HDBSCAN clustering for reconstructing the merger history of the Milky Way: applications and limitations

Andrea Sante ¹*, Andrea S. Font ¹, Dharmesh Mistry ², Sandra Ortega-Martorell ² and Ivan Olier ²

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

²*School of Computer Science and Mathematics, Liverpool John Moores University, James Parsons Building, 3 Byrom Street, Liverpool L3 3AF, UK*

Accepted 2026 March 11. Received 2026 March 1; in original form 2025 September 8

ABSTRACT

Clustering algorithms can help reconstruct the assembly history of the Milky Way by identifying groups of stars sharing similar properties in a kinematical or chemical abundance space. Despite being promising tools, their efficiency has not yet been fully tested in a realistic cosmological framework. We investigate the effectiveness of the HDBSCAN clustering algorithm in the recovery of the progenitors of Milky Way-type galaxies, using several systems from the Auriga suite of simulations. We develop a methodology aimed at improving the efficiency of the algorithm and avoiding fragmentation: first, we use a 12-dimensional feature space including a range of chemodynamical properties and stellar ages; furthermore, we optimize the algorithm using information from the internal structure of the clusters of accreted stars. We show that our approach yields good results in terms of both purity and completeness of clusters for galaxies with different types of accretion histories. We also evaluate the decrease in efficiency due to contamination by *in situ* stars. While for accreted-only haloes the algorithm matches well the recovered clusters with the individual progenitors and is able to recover accretion events up to a redshift of accretion $z_{\text{acc}} \sim 3$, for accreted + *in situ* haloes it can only identify the more recent accretion events ($z_{\text{acc}} < 1$). However, the purity of the identified clusters remains remarkably high even in this case. Our results suggest that HDBSCAN can efficiently identify accreted debris in Milky Way-type galaxies in realistic conditions, however, it requires careful optimization to provide valid results.

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

1 INTRODUCTION

In a Lambda cold dark matter (Λ CDM) cosmological model, galaxies similar to the Milky Way grow hierarchically, by accretion and disruption of a multitude of satellite dwarf galaxies (L. Searle & R. Zinn 1978; S. D. M. White & M. J. Rees 1978). The remnants of the disruption events are expected to reside in the stellar halo, in the form of tidal streams of various degrees of phase-mixing (A. Helmi 2020). The identification of past merger events can be done, in principle, by using information retained in their kinematics using the assumption that certain kinematical properties are conserved, i.e. energy, angular momenta, actions (A. Helmi et al. 1999; A. Helmi & P. T. de Zeeuw 2000; A. Knebe et al. 2005; A. S. Font et al. 2006; H. L. Morrison et al. 2009; F. A. Gómez et al. 2010). Moreover, stars from a specific progenitor are also expected to retain a memory of their common star formation history and, implicitly, chemical evolution (K. Freeman & J. Bland-Hawthorn 2002; see also the reviews of J. Bland-Hawthorn & O. Gerhard 2016 and A. J. Deason & V. Belokurov 2024).

The increased availability of high-quality astrometric data from *Gaia* (Gaia Collaboration 2016, 2023) and complementary chemical abundances from spectroscopic surveys such as APOGEE (S. R. Majewski et al. 2017), *Gaia*-European Southern Observatory (ESO) (G. Gilmore et al. 2012; S. Randich, G. Gilmore & Gaia-ESO Consortium 2013), LAMOST (G. Zhao et al. 2012), H3 (C. Conroy et al. 2019), or GALAH (S. Buder et al. 2021, 2025) has allowed the discovery of several sub-structures in the Galaxy. The debris identified so far appear as distinct features in the chemodynamical space (e.g. R. Klement et al. 2009; V. Belokurov et al. 2018; A. Helmi et al. 2018; H. H. Koppelman et al. 2019; J. M. D. Kruijssen et al. 2019; A. P. Ji et al. 2020; R. P. Naidu et al. 2020; D. Horta et al. 2023; K. Malhan & H.-W. Rix 2024).

However, the identification of debris remains challenging in the presence of the background population and uncertainties introduced by instrumental effects (A. G. A. Brown, H. M. Velázquez & L. A. Aguilar 2005), and given that overdensities originating from multiple progenitors overlap in kinematical space and share similar chemical abundances (I. Jean-Baptiste et al. 2017; M. D. A. Orkney et al. 2023; A. Carrillo et al. 2024; N. Kizhuprakkat et al. 2024; A. Mori et al. 2024).

Machine learning methods, and clustering algorithms in particular, are well-suited for identifying halo sub-structures, given

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

the natural tendency of stars from disrupted progenitors to cluster in kinematical space or display similar chemical abundances, in correspondence to the properties of progenitor galaxies they originate from. Indeed, clustering methods have shown their usefulness by detecting several new stellar sub-structures in the Milky Way and confirming some of the existing ones (e.g. P. Re Fiorentin et al. 2015; H. H. Koppelman et al. 2019; G. C. Myeong et al. 2019; N. W. Borsato, S. L. Martell & J. D. Simpson 2020; L. Necib et al. 2020; Z. Yuan et al. 2020; T. Ruiz-Lara et al. 2022; E. Dodd et al. 2023; B. Kim et al. 2025).

However, clustering methods also have some limitations when applied to Galactic archaeology. One example is the family of partitioning algorithms, such as k -means (J. MacQueen 1967), which aim to divide the data into distinct groups by minimizing the intra-cluster variance. Their requirement to predefine the number of clusters makes them less ideal for identifying debris from Milky Way progenitors, where this number is not known a priori. These methods are more useful when the clusters correspond to known structures, such as multiple stellar populations of globular clusters (D. Fraix-Burnet & E. Davoust 2015), or the distinct Galactic components (J. T. Mackereth et al. 2019). Moreover, the k -means algorithm is biased towards identifying spherically clustered groups of stars, which makes it unsuited for picking out tidal debris that typically displays an asymmetric distribution. For example, instead of ‘clumps’ in the energy–angular momentum space, as expected from conservation of these two quantities for any progenitor, the overdensities in this space could be severely distorted as debris orbits in an evolving gravitational potential or interacts with other substructures in the host galaxy (e.g. A. S. Font et al. 2006; F. A. Gómez et al. 2010, 2013).

Hierarchical clustering methods (see review of F. Murtagh & P. Contreras 2012), in particular the so-called single-linkage methods, build a nested tree of clusters by progressively merging the two closest clusters in the parameter space (or feature space, in machine learning terminology) until every point belongs to the same cluster. Despite mimicking the hierarchical nature of galaxy formation, these methods incorporate every data point in a hierarchy which makes them sensitive to noise. Isolated points in the feature space can act as a bridge between clusters or distort the internal structure of clusters. Moreover, extracting a final, ‘flat’ partition from the hierarchy is often not straightforward, as meaningful clusters, representing coherent debris from distinct accretion events, can appear at various levels within the hierarchy, reflecting differences in their spatial extent and detectability. For example, a tightly clustered debris from a recently accreted progenitor might remain distinct at a higher level in the hierarchy, whereas a more dispersed debris from an early accretion might merge with the background noise at a lower level. Using the single linkage-based clustering methodology developed by S. S. Lövdal et al. (2022) on the Auriga simulations, G. F. Thomas et al. (2025) found that only recently accreted ($< 6\text{--}7$ Gyr ago) progenitors are efficiently identified. They also found that the method is prone to generate artificial detections, with a significant fraction of clusters being composed mainly of *in situ* stars.

Density-based clustering methods (e.g. DBSCAN, HDBSCAN; M. Ester et al. 1996; M. Ankerst et al. 1999; R. J. G. B. Campello, D. Moulavi & J. Sander 2013) explicitly account for the presence of noise points by defining clusters as dense regions of points separated by sparser regions of noise. These methods are well suited for the distribution of debris from individual progenitor

galaxies, which correspond to overdensities in the feature space against a background of unrelated *in situ* stars. Furthermore, density-based methods make no assumptions about the shape of a cluster, allowing for the identification of elongated filamentary structures characteristic of tidal streams. HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise; R. J. G. B. Campello et al. 2013) offers significant advantages over its predecessor, DBSCAN (M. Ester et al. 1996). Although DBSCAN relies on a single distance scale to define dense regions, HDBSCAN identifies clusters at various scales simultaneously, providing a hierarchical representation of structures in the feature space. This is important as different progenitor systems can have specific internal velocity dispersions or metallicity distributions, which result in structures of different overdensities in the feature space. Moreover, HDBSCAN automates the extraction of a flat partition by selecting the most persistent and significant clusters in the hierarchy based on a cluster stability metric.

However, as with other clustering algorithms, the configuration of the clusters retrieved by HDBSCAN varies considerably with the choice of its internal parameters. An improper parameter selection can lead to either the fragmentation of larger, coherent structures or the erroneous aggregation of unrelated stars into spurious clumps (B. Kim et al. 2025).

In this study, we evaluate the effectiveness of HDBSCAN on simulated stellar haloes, where the ‘ground truth’ is known (i.e. tidal debris is already identified and mapped to its corresponding progenitor galaxies, and contamination from *in situ* stars is also known). We use a sample of three Milky Way-mass galaxies from Auriga suite (R. J. J. Grand et al. 2017, 2024), which has been shown to form realistic disc-dominated galaxies similar to the Milky Way (R. J. J. Grand et al. 2017; M. D. A. Orkney et al. 2022). The three galaxies are chosen so that they have significantly different assembly histories, in order to evaluate the variations in the performance of HDBSCAN with different chemodynamical distributions of tidal debris. We also incorporate a wider range of chemodynamical properties (features) in the analysis than previously used in similar clustering algorithms.

To avoid cluster fragmentation, we implement a systematic optimization of the HDBSCAN parameters tested against various internal and external metrics. The internal metrics account for how the data are partitioned by the algorithm, while the external metrics take into account the similarity between the clusters and the actual progenitor groups. Detailed explanations of the HDBSCAN algorithm, the optimization of its parameters, and the metrics used to evaluate its effectiveness are provided in Section 3.

Furthermore, the stellar halo of the Milky Way is known to have two distinct components, accreted and formed *in situ* (D. Carollo et al. 2007; T. C. Beers et al. 2012; V. Belokurov & A. Kravtsov 2022). Cosmological simulations naturally produce dual haloes (A. Zolotov et al. 2009; A. S. Font et al. 2011; I. G. McCarthy et al. 2012; P. B. Tissera et al. 2013; A. P. Cooper et al. 2015). Because the two halo populations can have overlapping chemodynamical properties (M. Haywood et al. 2018; P. Di Matteo et al. 2019; A. Mori et al. 2024), the performance of the clustering algorithm generally decreases in the presence of *in situ* stars or can lead to the identification of spurious sub-structures. To rigorously validate the efficacy of our optimized HDBSCAN methodology, we first apply it to accreted-only components. This allows us to quantify the intrinsic ability of the algorithm in recovering accreted structures of varying morphologies and densities without the contamination from the *in situ* population, hence establishing

a baseline for interpreting the results. With this baseline established, we then apply the algorithm to the fully mixed stellar halo.

The paper is structured as follows. Section 2 describes the *Auriga* suite of simulations and the selection of our sample of galaxies, including a brief overview of their distinct assembly histories. Section 4 presents the results of the optimized HDBSCAN, separately for the accreted-only (Section 4.1) and accreted + *in situ* stellar haloes (Section 4.2). The implications of our findings for reconstructing the accretion history of the Milky Way and the potential future directions of research are discussed in Section 5.

2 THE *Auriga* SIMULATIONS

The *Auriga* suite comprises 39 magnetohydrodynamical (MHD) simulations of Milky Way-mass ($0.5 < M_{200}/(10^{12} M_{\odot}) < 2$) haloes run with the AREPO (V. Springel 2010) moving-mesh code at two resolution levels. In the following, we provide a brief description of the *Auriga* suite and refer the reader to R. J. J. Grand et al. (2017) and R. J. J. Grand et al. (2024) for more details.

The Milky Way-mass haloes were randomly selected from the most isolated quartile of their respective mass ranges at $z = 0$ in the parent L100N1504 EAGLE dark-matter only simulation (J. Schaye et al. 2015). Zoomed-in simulations were then performed, following the particles within a distance of $4 R_{200}$ ¹ from the centre of the haloes to $z = 127$ and sampling lower mass particles in the enclosed area with higher mass particles added at larger distances. A Λ cold dark matter cosmology with parameters taken from Planck Collaboration XVI (2014) was adopted. Baryons were introduced by splitting the particles into dark matter particle-gas cell pairs according to the cosmological baryon mass fraction. The evolution of the MHD equations in a cosmological context was followed by AREPO through a discretization of the simulation domain carried out with the Voronoi tessellation of mesh-generating points that move with the velocity of the fluid flow. Each cell was assigned a given mass so that high-density regions are resolved with a larger number of cells.

The *Auriga* galaxy formation model incorporates a wide range of subgrid physics, including primordial and metal line cooling with self-shielding corrections against a spatially uniform, redshift-dependent ultraviolet (UV) background field (C.-A. Faucher-Giguère et al. 2009) that completes reionization at $z \sim 6$ (M. Vogelsberger et al. 2013). The interstellar medium (ISM) is treated using the two-phase model of V. Springel & L. Hernquist (2003), with stochastic star formation occurring in gas denser than $n = 0.13 \text{ cm}^{-3}$ assuming a G. Chabrier (2003) initial mass function. The model also accounts for stellar evolution, tracking the associated mass-loss and chemical enrichment of the surrounding gas from asymptotic giant branch stars (A. I. Karakas 2010), Type Ia supernovae (F.-K. Thielemann et al. 2003; C. Travaglio et al. 2004) and Type II supernovae (L. Portinari, C. Chiosi & A. Bressan 1998). Feedback is implemented through stellar wind particles released from Type II supernovae, as well as with radiative and thermal energy outbursts from active galactic nuclei with both radio and quasar modes. The model also accounts for the seeding and growth of supermassive black holes

(V. Springel 2005) and includes a uniform, comoving primordial magnetic field of 10^{-14} G (R. Pakmor & V. Springel 2013).

To trace the formation history of galaxies, simulation outputs (‘snapshots’) were post-processed to identify structures and link them through time. First, Friends-of-Friends (FOF) groups were identified based on the dark matter particle spatial distribution. Within these FOF groups, the SUBFIND algorithm (V. Springel et al. 2001) was run on all particle types to identify gravitationally self-bound substructures, or sub-haloes. These sub-haloes were then linked across snapshots to construct merger trees. In this process, a sub-halo in a given snapshot was identified as the main progenitor of a sub-halo in a subsequent snapshot if it contributes the highest number of particles, weighted by their binding energy.

Star particles were tagged as either accreted or formed *in situ*, by tracing them back to their point of origin: accreted if they formed within a satellite dwarf galaxy (either before or after accretion on to the host) and are gravitationally bound to the main halo at $z = 0$, and *in situ* if they formed in the main galaxy (A. Monachesi et al. 2019). We use this information here, together with additional data on the accreted star particles available in the *Auriga* public database², namely the merger tree index of the sub-halo where each accreted star particle formed, `RootIndex`, and the index of the sub-halo to which the star particles were bound at the time when the sub-halo reached its maximum stellar mass, `PeakMassIndex` (see R. J. J. Grand et al. 2024).

In our analysis, a progenitor of a Milky Way-mass galaxy is defined as a group of accreted stars³ that share the same `PeakMassIndex`; henceforth, we will refer to these progenitors based on their respective numerical index, i.e. ‘prog. `PeakMassIndex`’. We motivate this definition by the fact that sub-haloes tend to reach their maximum stellar mass just before accretion on to their host, that is, when they first cross the R_{200} of their host. At around this time, which we denote τ_{infall} , most infalling satellites experience tidal and ram-pressure stripping, which results in a sharp decrease in their star formation (e.g. D. Kawata & J. S. Mulchaey 2008; C. M. Simpson et al. 2018; A. S. Font et al. 2022). We therefore evaluate the performance of the HDBSCAN algorithm by comparing the clusters of accreted stars identified by the algorithm with the `PeakMassIndex`-defined progenitor groups.

The Milky Way-mass galaxies in *Auriga* match the galaxy scaling relations, e.g. the stellar mass–halo mass and stellar mass–metallicity relations (R. J. J. Grand et al. 2017). They also have stellar haloes with similar stellar masses, surface brightnesses, and metallicity profiles of observed Milky Way analogues (A. Monachesi et al. 2019). Morphologically, they are also similar to the Milky Way at present time (i.e. they have rotationally supported discs). In particular, for our study, the suite of simulations includes a variety of assembly histories, which allows us to test the effectiveness of the HDBSCAN algorithm on different chronochemodynamical distributions of stellar haloes.

¹Galactocentric distance which defines a spherical region where the average matter density is 200 times the critical matter density of the universe. It is often used as a proxy for the virial radius of a halo.

²This information is provided for each simulation as part of the ‘Accreted particle lists’ catalogue downloadable at https://wwwmpa.mpa-garching.mpg.de/auriga/data_new.html.

³We note that the term ‘stars’ in the context of simulations refers to star particles, which are essentially single stellar populations.

2.1 The simulations sample

The three galaxies considered here, Au 7, Au 25, and Au 27, are selected from the original sample of 30 Milky Way-mass systems in the ‘level 4’ suite of simulations. The particle mass resolution is $\sim 4 \times 10^5 M_\odot$ for dark matter and, typically, $\sim 5 \times 10^4 M_\odot$ for baryons.

The three host galaxies display significantly different present-day distributions of accreted stars, in either kinematics, chemical abundances or ages. Au 7 has a predominantly early formation epoch, and consequently the accreted debris is mostly phase-mixed at $z = 0$; in contrast, Au 25 has a late formation, and its debris is mostly in the form of cold stellar streams; Au27 has an assembly history more similar to the Milky Way, including an early major merger similar to the *Gaia* Enceladus–Sausage (GES) (A. Fattahi et al. 2019). Au 27 also has a large stellar disc and a bar at $z = 0$.

In addition to the information from the merger trees, we also use the morphological classification of tidal debris in this Auriga suite performed by A. H. Riley et al. (2025) and N. Shipp et al. (2025). This classification groups accretion events into three categories, ‘intact’, ‘phase-mixed’, and ‘stellar streams’, based on their configuration at $z = 0$. An accretion event is classified as intact if the majority (> 97 per cent) of accreted stars are still bound to the system at present. The debris from disrupted satellites is considered either phase-mixed or as (cold) stellar stream based on a mass-dependent cut in the velocity dispersion, with the phase-mixed structures exhibiting a higher velocity dispersion than the streams.

For each galaxy, we consider only the accreted stars that contribute to the main halo, defined as those accretion events that contribute at least 0.1 per cent to the total number of accreted particles at $z = 0$. This corresponds to a cut-off in the number of particles per progenitor ranging between 100 and 1000, depending on the system. We also exclude particles that are bound to surviving satellites and those from merger events labelled as ‘intact’, under the assumption that the identification of these systems would not require the application of clustering techniques.

The assembly histories of the selected galaxies are briefly summarized below, with an overview of the number and properties of their most significant progenitors given in Table 1. The merger trees for the galaxies in the sample are shown in the first row of Fig. 7.

2.1.1 Au7

The majority (> 90 per cent) of the accreted stellar content in Au7 is in the form of phase-mixed structures. These structures originate mostly from three progenitors (prog. 3210, prog. 4401, and prog. 2197) of similar stellar mass ($\sim 10^{10} M_\odot$), which contribute 20 per cent, 34 per cent, and 36 per cent, respectively, to the accreted halo mass. These events succeeded each other almost regularly every ~ 3 Gyr starting from $z \sim 1$.

2.1.2 Au25

Most of the accreted star particles bound to the main halo at $z = 0$ in Au25 originate from a satellite of stellar mass $\sim 10^{10} M_\odot$, prog. 736, accreted onto the main halo < 2 Gyr ago. However, these comprise only a small fraction (~ 5 per cent) of the stellar content of the satellite, which is still present as a self-bound system at $z = 0$. The other massive accretion event experienced by the

Au25 was a satellite of stellar mass $\sim 10^{8.5} M_\odot$, prog. 41397, accreted at $z_{\text{acc}} > 2$. The stars disrupted from this progenitor form a phase-mixed structure at present, comprising ~ 34 per cent of the accreted stars of its host.

2.1.3 Au27

The accreted halo in Au27 is dominated by stars originating from a massive satellite (stellar mass of $\sim 10^{9.5} M_\odot$), prog. 7219, accreted on a radial orbit at $z_{\text{acc}} \sim 1.5$, similar to a GES event (V. Belokurov et al. 2018; A. Helmi et al. 2018). This contributed ~ 40 per cent of the accreted star particles bound to the main halo at $z = 0$. Two other significant, slightly less massive satellites were accreted just before, contributing 19 per cent and, respectively, 14 per cent to the mass of the accreted halo. Although most of the accreted debris are completely mixed in with *in situ* stars, 13 per cent of the accreted stars still present stellar stream-like configurations; such an example is the remnant of the massive satellite, prog. 205, accreted ~ 7 Gyr ago. Au27 also presents a companion of similar mass as the Large Magellanic Cloud at $z = 0$, which, however, is not considered in the analysis.

3 THE HDBSCAN METHODOLOGY

HDBSCAN (R. J. G. B. Campello et al. 2013) is a non-parametric clustering algorithm that identifies clusters of varying shapes and densities, while also identifying data points that do not belong to any cluster, termed noise. It extends the DBSCAN algorithm by producing a hierarchical clustering structure from which a simplified set of flat clusters can be extracted. The methodology begins by transforming the feature space to reflect the local density of data points through the concept of *core distance*. For each data point x , its core distance, $core_k(x)$, is the distance to its k th nearest neighbour, providing an estimate of the local density.

A new distance metric, the *mutual reachability distance*, is then defined between two points, a and b , as

$$d_{\text{mreach-}k}(a, b) = \max(\text{core}_k(a), \text{core}_k(b), d(a, b)),$$

where $d(a, b)$ is the original distance. This metric effectively smooths out the density of the data space by ensuring that points in dense regions remain close while pushing points in sparse regions further apart.

The data set is then treated as a fully connected graph, where the edge weights are the mutual reachability distances. From this graph, a *minimum spanning tree* (MST) is constructed. This is a sub-graph connecting all points with the minimum possible total edge weight and no cycles. The MST captures the underlying structure of the data.

Afterward, a cluster hierarchy is created by sorting the MST edges by weight in increasing order and iteratively merging the components connected by each edge. This process results in a tree of nested clusters. This full hierarchy is then simplified by a condensation process governed by the `min_cluster_size` parameter (see Section 3.1). Traversing the hierarchy from the root, if a split results in a new cluster with fewer points than `min_cluster_size`, that component is deemed to be made of points that have ‘fallen out of a cluster’ and is not considered for further partitioning. The final step is to extract a stable, flat set of clusters from this condensed hierarchy. A stability score is

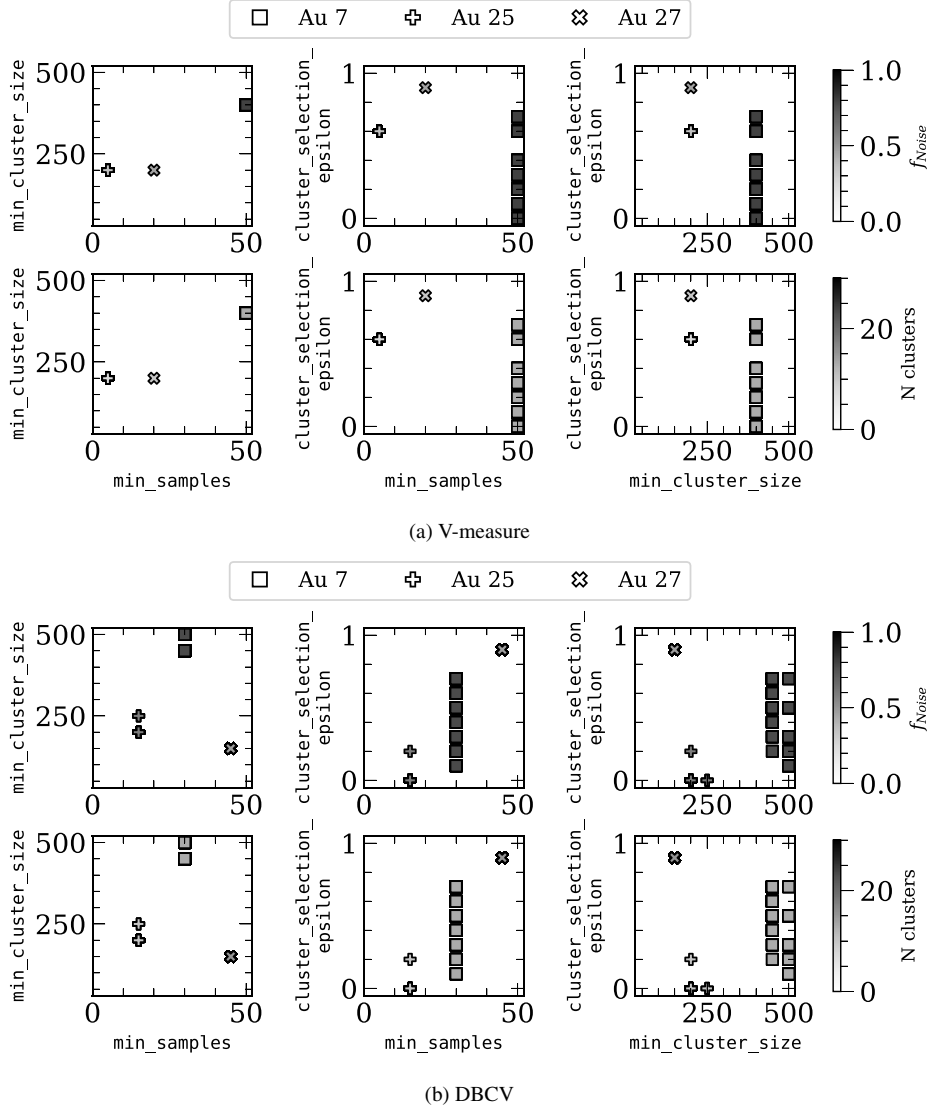


Figure 1. Optimal HDBSCAN parameter configurations found by the Optuna search. Panels (a) and (b) show the results of optimizations using the V-measure and DBCV metrics, respectively. In both panels, each symbol represents a different galaxy. The top sub-panels show the configurations coloured by the fraction of particles classified as noise, while the bottom sub-panels show the same configurations coloured by the number of clusters in the resulting partition.

calculated for each cluster based on its persistence or ‘lifespan’ within the hierarchy. The clusters with the highest stability are selected as the final output, and any point not belonging to a selected cluster is classified as noise.

3.1 Optimizing the HDBSCAN parameters with OPTUNA

For the analysis, we use the algorithm implementation provided by the HDBSCAN PYTHON package (L. McInnes, J. Healy & S. Astels 2017). The functioning of a HDBSCAN clustering model is governed primarily by the following parameters:

(i) `min_samples`, which sets the minimum number of samples required in a neighbourhood for a point to be considered a core point. This parameter influences the core distance calculation and dictates how conservatively the density is estimated. A higher value results in larger core distances, increasing the likelihood that points in low-density regions are classified as noise.

(ii) `min_cluster_size`, which defines the minimum number of points a cluster must have to be considered a genuine cluster during the hierarchy condensation step.

(iii) `cluster_selection_epsilon`, which specifies a distance threshold for cluster extraction. When traversing the condensed hierarchy, if a split from a parent to its children occurs over a distance greater than `cluster_selection_epsilon`, the parent cluster is selected instead of its children. This allows for the extraction of clusters at a specific density level.

The choice of these parameters can affect the way the data are grouped into clusters, which consequently impacts the efficiency of HDBSCAN in retrieving the progenitors of a galaxy. To ensure an objective and reproducible analysis, we have automated the parameter selection process. This process involves the definition of a search space for the relevant parameters, running HDBSCAN models for different configurations sampled from this space, and evaluating the resulting clusters with a given metric. We im-

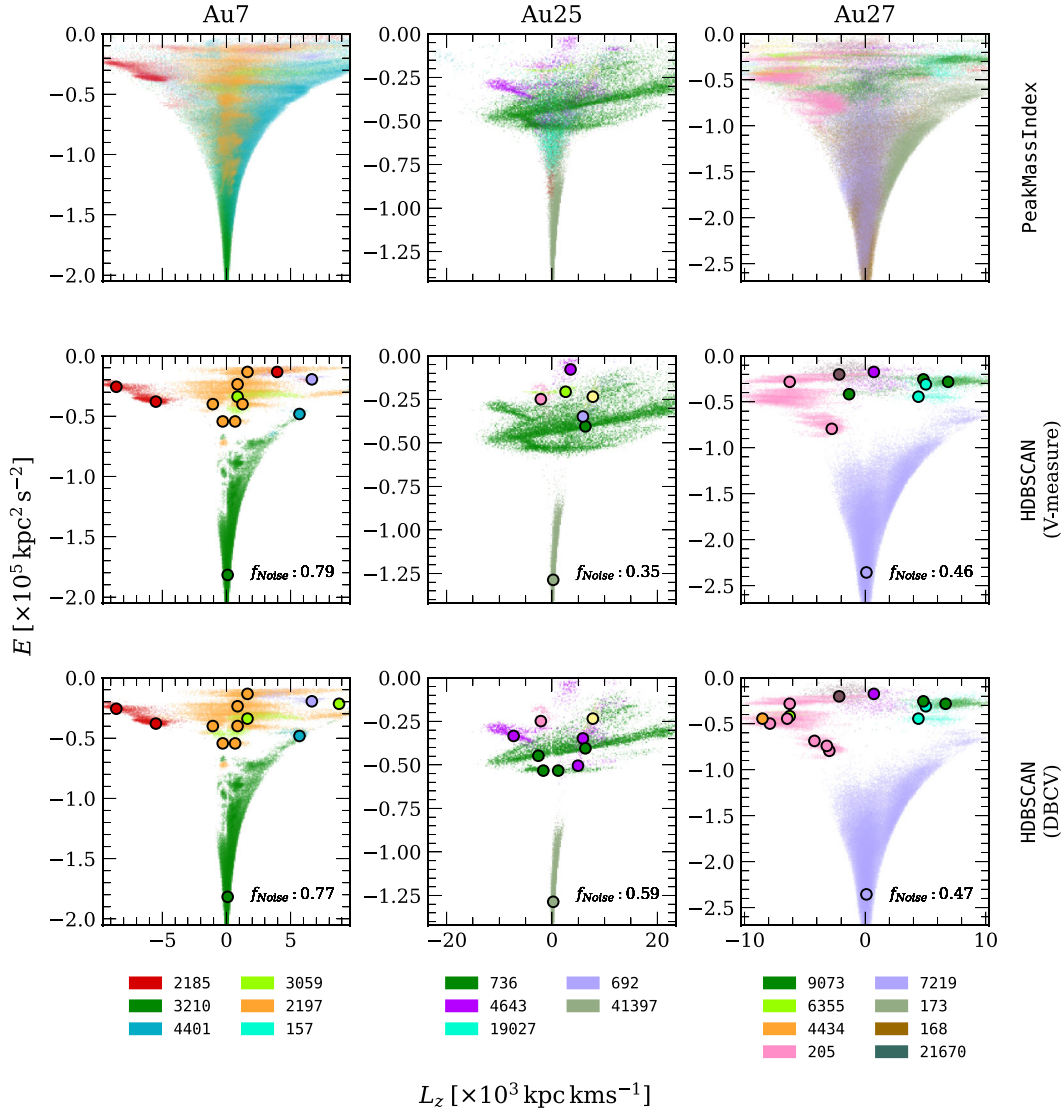


Figure 2. Density distribution of accreted star particles in the $E - L_z$ plane. The top row shows the ground-truth distribution, where particles are coloured by their progenitor group (PeakMassIndex). The middle and bottom rows show the HDBSCAN clustering results when optimized with the V-measure and DBCV metrics, respectively. In these lower panels, identified clusters are coloured according to their majority progenitor group, with noise particles shown in grey. The overplotted scatter points highlight the highest-density region of each cluster. The fraction of noise points is also included in each panel. At the bottom of each column, a legend shows the colours associated with the merger events reported in Table 1.

plemented this optimization using the Optuna hyperparameter search framework (T. Akiba et al. 2019). Optuna efficiently samples the parameter space by constructing a probabilistic model from past trials to predict which configurations are most likely to maximize the evaluation metric.

Specifically, the parameter ranges explored in the search are:

- (i) $\text{min_samples} \in [5, 50]$
- (ii) $\text{min_cluster_size} \in [50, 500]$, and
- (iii) $\text{cluster_selection_epsilon} \in [0.0, 1.0]$.

Two distinct evaluation metrics were considered to guide the optimization:

(i) the V-measure (A. Rosenberg & J. Hirschberg 2007), an external entropy-based metric that compares the particle IDs in the HDBSCAN clusters to the ones in the progenitor groups defined by PeakMassIndex . It is defined by a combination of two

criteria: homogeneity, which is satisfied if each cluster contains only particles of a single progenitor, and completeness, which is satisfied if all members of a given progenitor are assigned to the same cluster. The V-measure score ranges from 0 to 1, with higher scores indicating better clustering results. A score of 1 represents a perfectly homogeneous and complete clustering, meaning each cluster contains only particles from a single progenitor, and all particles of a given progenitor are grouped together in one cluster. A score of 0 indicates a completely random or poor clustering, where there is no relationship between the clusters and the true progenitor groups.

(ii) the Density-Based Clustering Validation index (DBCV; D. Moulavi et al. 2014), an internal metric that does not require information on the true progenitor groups. DBCV is specifically designed for density-based methods and evaluates the intrinsic quality of a clustering by assessing the density within clusters relative to the density between clusters. It rewards partitions where

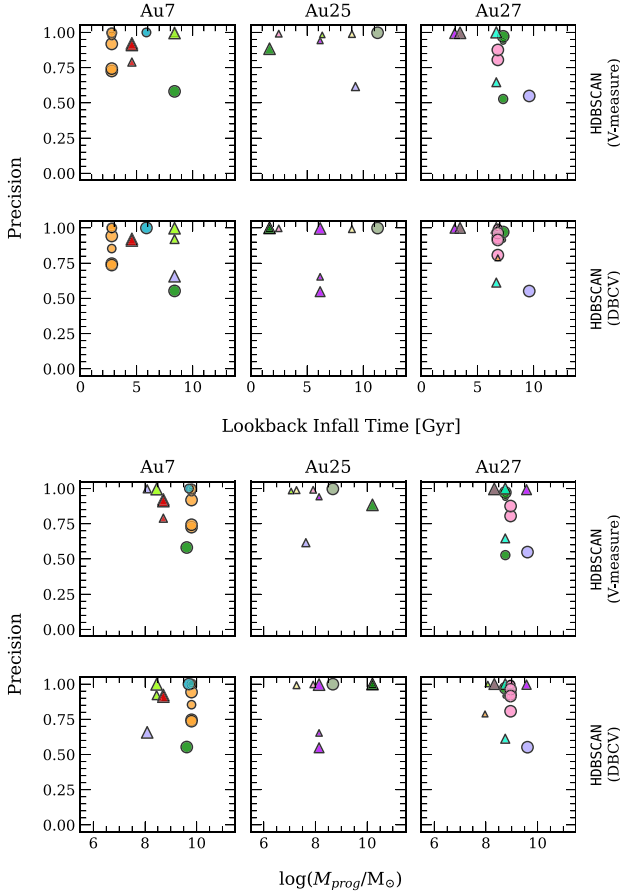


Figure 3. Precision scores for the clusters identified in accreted-only haloes by the HDBSCAN models optimized with the V-measure and DBCV index metrics. For each cluster, the precision is calculated with respect to the (dominant) progenitor with the highest number of cluster members, and plotted against the lookback infall time (top) and stellar mass at infall of the progenitor (bottom). The clusters are represented as triangles if the dominant progenitor is a stellar stream at $z = 0$ and by circles if a phase-mixed structure. The size of the symbols is proportional to the number of members in the cluster, and the colour scheme follows the one in Fig. 2 and indicates the dominant progenitor of each cluster.

clusters are internally dense (low sparseness) and well separated from other clusters by low-density regions.

By optimizing the HDBSCAN parameters using the V-measure, we establish a performance benchmark that represents the best possible clustering outcome with respect to the known progenitor labels. The DBCV metric is then used to investigate whether a comparable performance can be achieved in the more realistic scenario where the true origin of the accreted debris is unknown.

3.2 Evaluation metrics

While the V-measure provides a global assessment of the purity and completeness of clustering, we also investigate the purity and completeness for the individual clusters. To this end, we calculate the *precision* and *recall* for each identified cluster with respect to its dominant progenitor group (i.e. the progenitor with the most members in that cluster). Here, precision is defined as the fraction of the members of a cluster that belong to the dominant progenitor, while recall is the fraction of all stars from that dominant

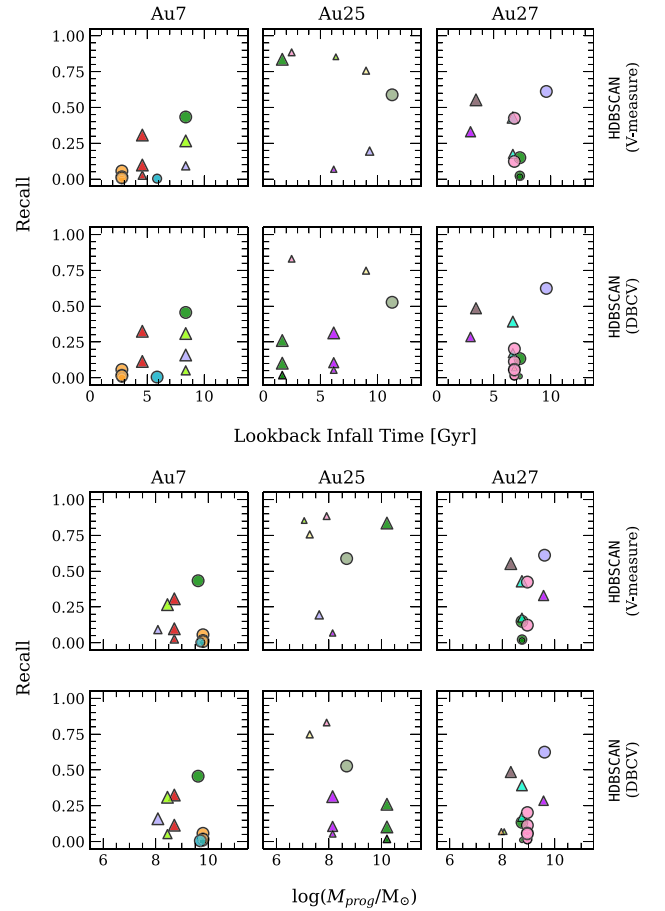


Figure 4. Recall scores for the clusters identified in accreted-only haloes by the HDBSCAN models optimized with the V-measure and DBCV index metrics. The symbols, colours, and sizes are the same as described in Fig. 3, but now for the recall scores.

progenitor that are contained within the cluster. Assuming \mathbf{Cl} is the set of particles of a cluster and \mathbf{Pr} is the set of particles from the dominant progenitor of that cluster, i.e. sharing the same PeakMass Index, the precision and recall scores for the cluster are calculated as:

$$\text{Precision} = P = \frac{|\mathbf{Pr} \cap \mathbf{Cl}|}{|\mathbf{Cl}|} \quad \text{and} \quad \text{Recall} = R = \frac{|\mathbf{Pr} \cap \mathbf{Cl}|}{|\mathbf{Pr}|},$$

where ‘|’ indicates the cardinality, i.e. the number of elements of the sets.

In order to quantify the *spread* of the stellar debris of a progenitor in a multidimensional parameter space (see Section 3.3), we calculate its covariance matrix (\mathcal{C}) (from now on we will refer to the covariance matrix as the ‘spread’). This parameter captures the information on both the dispersion along the single properties of the stellar debris and their inter-dimensional correlations. The trace of the covariance matrix is essentially the *total variance* (σ_T^2) and measures the spread of values for the particles in the distribution, i.e.

$$\sigma_T^2 = \text{tr}(\mathcal{C}) = \sum_{n=1}^D \sigma_n^2,$$

where σ_n^2 is the variance of the star particles along the dimension n .

Under the assumption that the distribution of the particles from a progenitor is a multivariate Gaussian in the parameter

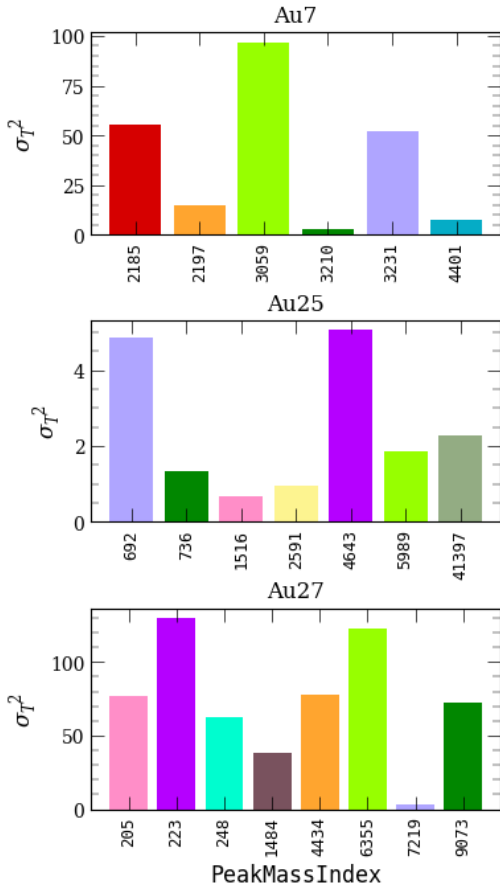


Figure 5. Total variance values for the progenitor galaxies appearing as dominant progenitors in the HDBSCAN clusters.

space, which is enforced by the transformation of variables in the pre-processing stage (see Section 3.3), the degree of similarity between the debris from two different progenitors (A and B) can be estimated using the Mahalanobis distance, i.e.

$$D_M = \sqrt{(\mu_A - \mu_B)^T (C_A + C_B)^{-1} (\mu_A - \mu_B)},$$

where μ_A and μ_B represent the mean vectors of the stellar properties for the particles from progenitors A and B . The Mahalanobis distance measures the difference between the mean stellar properties of the debris from the two progenitors in terms of their combined variance. We also use this metric to quantify the mixing of the debris of a progenitor with the *in situ* population.

3.3 The 12-dimensional parameter (feature) space

We apply the HDBSCAN algorithm on a multidimensional parameter space comprised of ages, chemical abundance ratios, and kinematics of halo stars. The kinematics are computed in a Cartesian coordinate system centred on the centre-of-mass of the host galaxy, and the z -axis aligned with the total angular momentum vector of the disc component. Two of the kinematical parameters are the integrals of motion (IoM): the energy E and L_z , the component of the angular momentum along the z -axis, which is perpendicular to the plane of the stellar disc. We also include the perpendicular component of the angular momentum, which is the component in the $x - y$ plane, $L_\perp = \sqrt{L_x^2 + L_y^2}$. This




Figure 6. Mahalanobis distance between the 12D debris distributions of the progenitor with the largest number of particles classified as noise and the other progenitors of the galaxies. The symbol and colours of the progenitors are the same one used in Fig. 3.

is often used, even though it is known to not be fully conserved (A. Helmi et al. 1999). To these, we add the Cartesian positions and velocities⁴ (x, y, z, v_x, v_y, v_z). For chemical abundances, we use the α -to-iron and iron-to-hydrogen chemical abundance ratios ($[\alpha/\text{Fe}]$, $[\text{Fe}/\text{H}]$), while for ages we use the lookback time since formation of star particles (τ_{form}). In summary, we choose a 12-dimensional (12D) parameter space comprised of:

$$\{E, L_z, L_\perp, x, y, z, v_x, v_y, v_z, [\text{Fe}/\text{H}], [\alpha/\text{Fe}], \tau_{\text{form}}\}$$

Each of these stellar properties provides complementary information in the identification of progenitors. The IoMs are conserved quantities over time, which makes them ideal for identifying early mergers that appear as phase-mixed sub-structures at $z = 0$ (A. Helmi & S. D. M. White 1999). The phase-space coordinates are most effective for recently accreted debris or progenitors

⁴We experimented with the removal of positions and velocities from the clustering features, but we found a decrease in the performance of the model compared with including them. We therefore decided to incorporate positions and velocities among the features as they may add information (for example, the location of particles in the host galaxy, which otherwise might be hard to disentangle from the angular momentum features).

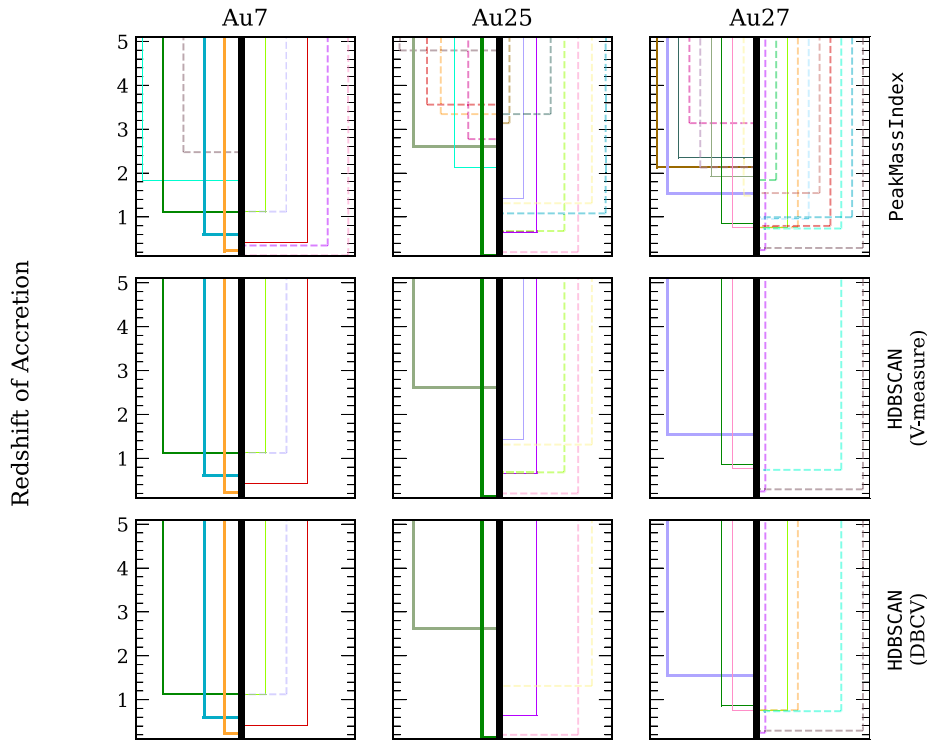


Figure 7. Merger trees for the galaxies in the sample as reconstructed by the V-measure (middle) and DBCVC (bottom) optimized HDBSCAN models for the accreted-only haloes. Top panel shows the full merger trees. Each colour represents a different progenitor galaxy. The thickness of each line is proportional to the number of accreted star particles from that progenitor. Dashed lines indicate progenitors that contribute < 1 per cent to the overall accreted population of the galaxy at $z = 0$.

that are actively being disrupted, where the positions and velocities can be highly correlated (K. V. Johnston et al. 2008). Chemical abundance ratios are excellent tracers of the common chemical evolution of stars within satellite galaxies prior to accretion on to their host (K. Freeman & J. Bland-Hawthorn 2002), while stellar ages allow the reconstruction of the star formation history from a single progenitor (J. Montalbán et al. 2021; C. Gallart et al. 2024).

Because the Auriga galaxies include a mix of stellar populations, from phase-mixed stars in the inner galaxy to dynamically cold stellar streams in the outer halo, the feature distributions are often skewed and heavy-tailed along the single dimensions in the 12D clustering space. To mitigate this, we preprocess the data to produce more Gaussian-like distributions. Specifically, we apply a cubic-root transformation to L_z , a logarithmic transformation to L_{\perp} , and an inverse hyperbolic sine transformation to the Cartesian coordinates (x, y, z) . Finally, to account for the different scales and units of the various features, the data set of each galaxy is standardized by subtracting the median from each feature and dividing by its interquartile range, making it less sensitive to outliers than the standard scaling.

After preprocessing the data, we determine the optimal configuration of the HDBSCAN model parameters for each galaxy by exploring 640 different configurations of model parameters, as described in Section 3.1.

4 RESULTS OF CLUSTERING

In this section, we present the results of clustering for both the accreted-only and accreted + *in situ* stellar haloes. These are pre-

sented separately, in Sections 4.1 and 4.2, respectively. The results on accreted-only haloes allow us to evaluate to what degree the assembly history of a galaxy can be reconstructed by HDBSCAN in the ‘best-case scenario’, without contamination from *in situ* stars. For the accreted + *in situ* case, we aim to exclude the stellar disc which otherwise would dominate over other overdensities in feature space. For this, we employ a velocity cut that separates halo stars from a rotationally supported disc (e.g. H. H. Koppelman et al. 2019); specifically, we define halo stars those with $|\mathbf{V} - \mathbf{V}_*| \geq 200 \text{ km s}^{-1}$, where $\mathbf{V}_* = (0, 200, 0) \text{ km s}^{-1}$ in a galactocentric cylindrical coordinate system $(\hat{v}_r, \hat{v}_\theta, \hat{v}_z)$. We note that even with this selection, the resulting data sets have considerably high *in situ* fractions, of 0.49, 0.77, and 0.72 for Au7, Au25, and Au27, respectively.

The structure of the two subsections is very similar; they both include a discussion of the purity and completeness of the retrieved clusters and a comparison between the merger trees inferred from the clustering analysis with the full merger trees constructed from the original simulations.

4.1 Clustering in accreted-only stellar haloes

From the optimization procedure described in Section 3.1, we find that there are several ‘optimal’ HDBSCAN parameter configurations that return similar partitioning of the data. This can be seen by the consistent number of clusters and the fraction of noise points in Fig. 1, which shows the HDBSCAN parameter values for the runs performed with both the V-measure (1a) and DBCVC (1b) metrics. Configurations corresponding to a given

Table 1. Merger events contributing at least 1 per cent of the accreted stars bound to the main halo at $z = 0$ for the galaxies in the sample. From the leftmost to the rightmost, the columns describe: (i) the index of the progenitor in the merger tree of the galaxy; (ii) the lookback infall time, i.e. the first time the progenitor crossed R_{200} of its host; (iii) the stellar mass of the progenitor at the time of infall; (iv) the fraction of star particles bound to the main halo at $z = 0$ which were accreted from the progenitor; and (v) the present-day configuration of the progenitor debris as classified by A. H. Riley et al. (2025).

PeakMassIndex	τ_{infall} (Gyr)	$\log(M_*/M_\odot)$	f_{acc}	Morphology
Au7				
prog. 2197	2.8	9.8	0.36	phase-mixed
prog. 2185	4.6	8.7	0.02	stream
prog. 4401	5.9	9.7	0.34	phase-mixed
prog. 3210	8.4	9.6	0.20	phase-mixed
prog. 3059	8.4	8.5	0.02	stream
prog. 157	10.2	8.9	0.04	phase-mixed
Au25				
prog. 736	1.7	10.2	0.44	stream
prog. 4643	6.2	8.1	0.08	stream
prog. 692	9.3	7.6	0.03	stream
prog. 19027	10.7	8.1	0.07	phase-mixed
prog. 41397	11.3	8.7	0.34	phase-mixed
Au27				
prog. 205	6.8	9.0	0.13	phase-mixed
prog. 6355	6.8	8.0	0.01	stream
prog. 4434	6.8	8.0	0.01	stream
prog. 9073	7.3	8.8	0.05	phase-mixed
prog. 7219	9.6	9.6	0.40	phase-mixed
prog. 173	10.4	9.1	0.14	phase-mixed
prog. 168	10.7	9.2	0.20	phase-mixed
prog. 21670	11.0	8.6	0.03	phase-mixed

host galaxy are shown with the same symbol. The number of clusters and the fraction of noise points associated with each clustering model are indicated on the greyscale bars. This figure also shows that a degeneracy in the clustering evaluation metrics arises once the optimal values for the `min_samples` and `min_cluster_size` parameters are identified. These two parameters establish an optimal condensed hierarchy, after which the final flat partition becomes less sensitive to the precise value of `cluster_selection_epsilon`, as the clusters can persist over a wide range of distance scales in the hierarchy. Based on these findings, we proceed to randomly select one configuration from this optimal set of parameter values to define the HDBSCAN model used for clustering.

The results of the clustering analysis are visualized in Fig. 2, which shows the density distribution of the accreted star particles in the $E - L_z$ plane for the three galaxies (one per column). The top row serves as the ground truth, with particles coloured according to their original progenitor group (i.e. their `PeakMassIndex`). The middle and bottom rows show the results of the HDBSCAN clustering when optimized using the V-measure and DBCV metrics, respectively. To facilitate the comparison, each identified cluster is assigned the colour of the ground truth progenitor group that constitutes the majority of its members; particles classified as noise are shown in grey. For a better visualization of the distribution of clusters, the highest-density region of each cluster is marked with scatter circles. This helps locate the smaller clusters that may be obscured by the background noise in this plane.

From the top row, it is clear that only a small number of progenitors contribute the majority of accreted stars in each galaxy. As a consequence, most of the clusters identified by the HDBSCAN models are dominated by particles from these major merger events, some progenitors being resolved into multiple distinct clusters. Although a significant fraction of accreted stars are classified as noise, samples from the major merging events are consistently recovered. For instance, in Au7, which has the highest fraction of noise points, the identified clusters contain samples from 6 of the 10 progenitors, which, however, together contributed ~ 95 per cent of the total accreted particles.

The cluster distributions obtained with the V-measure and DBCV optimizations (middle and bottom rows of Fig. 2) are remarkably similar. This suggests that a partition of the data maximizing the intrinsic density and separation of the clusters is also strongly linked to one with a high level of purity and completeness. Quantitatively, the V-measure scores for the clusters obtained by optimizing with the V-measure are 0.17, 0.49, and 0.26 for Au7, Au25, and Au27, respectively. The partitions found by optimizing the DBCV index achieve comparable V-measure scores of 0.17, 0.37, and 0.24 for the same galaxies. This compares to V-measure scores of 0.01, 0.03, and 0.14 for the three galaxies when the clustering is performed with a non-optimized model (i.e. run with the default parameter values). This result is significant, as it suggests that even without ground-truth information, an unsupervised metric like the DBCV index can effectively guide the HDBSCAN algorithm towards a physically meaningful partition of the data that reflects the underlying progenitor structure. However, we note that the model optimized with the DBCV index sometimes tends to fragment progenitor groups into a larger number of clusters than the one optimized with the V-measure (see the green clusters for Au25 and the pink clusters for Au27 in Fig. 2).

4.1.1 Purity and completeness of clusters

The precision and recall for the clusters identified in the sample galaxies are reported in Figs 3 and 4, respectively. In each figure, the top panel refers to the clusters from the V-measure optimized model, while the bottom panel refers to the ones from the DBCV-optimized. The metrics are plotted against the lookback infall time, τ_{infall} (top) and the stellar mass at infall (bottom) of the dominant progenitor. Each cluster is represented by a symbol that indicates whether its dominant progenitor is a stellar stream (triangles) or a phase-mixed sub-structure (circles) at $z = 0$. The sizes of symbols scale with the number of stars in each cluster, the largest symbols representing clusters with more than 1000 members.

Generally, HDBSCAN identifies clusters with a high-precision, which means that clusters are mostly composed of stars that belong to a single *dominant* progenitor. The lowest precision scores for the V-measure model are associated with specific physical scenarios. For Au7, the cluster linked to prog. 3210 (precision $P = 0.58$) corresponds to the oldest major merger. In Au25, the cluster for prog. 692 ($P = 0.62$) is a small stream dynamically similar to debris from a recent major merger, prog. 736. For Au27, low-precision scores are found for clusters associated with prog. 9073 ($P = 0.53$), prog. 248 ($P = 0.65$), and prog. 7219 ($P = 0.55$). The former two progenitors are fragmented into three and two clusters, respectively; the low-precision clusters are also those that overlap with other debris in the phase space. The clus-

ter associated with prog. 7219 is significantly contaminated by debris from the other two main merger events (prog. 168 and prog. 173), which slightly preceded the dominant progenitor and have not been recovered by the clustering model. The DBCV-optimized model identifies the same low-precision progenitors, with the addition of clusters for prog. 4643 in Au25, which also dynamically overlaps with debris from prog. 736.

On average, clusters associated with stellar streams show higher precision ($P = 0.91$ for V-measure and 0.90 for DBCV) than those associated with phase-mixed structures ($P = 0.84$ for V-measure and 0.88 for DBCV). These dynamically cold filamentary structures are less likely to overlap with debris from other progenitors than the phase-mixed structures, the latter often being contaminated by earlier mergers that remain undetected. This effect may also explain the higher precision found for clusters from the late mergers (that is, lookback time of infall ≤ 5 Gyr) compared to earlier ones; this effect is particularly evident for Au27 (see the rightmost sub-panel in top panel of Fig. 3). In terms of progenitor masses, clusters from the most massive mergers tend to have lower purity scores, which is likely due to the wider range in stellar properties (features) of their debris.

In contrast to the high-precision scores, the recall scores are significantly lower for both models. The average recall (R) values for clusters of stellar streams are 0.40 (for V-measure) and 0.24 (for DBCV), while for clusters of phase-mixed structures are 0.16 (for V-measure) and 0.13 (for DBCV), respectively. This is primarily due to two factors: a large number of particles are being classified as noise, and the fragmentation of single progenitors into multiple clusters. The latter effect explains the lower recall of the DBCV-optimized model, which tends to produce more clusters per progenitor. For example, for prog. 2197 in Au7, 85 per cent of its particles are classified as noise in both models, while 10 per cent are split into six separate clusters.

The fragmentation of the debris of a progenitor into multiple clusters is expected for the most massive progenitors. Those are typically accreted late and experience a prolonged chemical evolution; therefore, they tend to consist of multiple stellar populations and their debris spans a wide range in chronochemodynamical properties. To investigate this hypothesis, we calculated the total variance of the debris of each progenitor in the 12D parameter space (Fig. 5). Progenitors that are split into several clusters generally show higher variance than the rest. However, a large variance does not always lead to fragmentation. In cases where the debris are too widely dispersed, the algorithm either assigns many of those stars to other clusters or classifies them as noise. Stars are most often classified as noise when debris from different progenitors overlap in the parameter space. This effect is illustrated in Fig. 6, where the progenitors with the largest number of stars flagged as noise are also those with the smallest Mahalanobis distance from the progenitor that dominates the noise population.

4.1.2 The assembly histories inferred from clustering

We further investigate the assembly histories of the three galaxies, as represented by their merger trees constructed from the clusters recovered by the optimized algorithm. The full merger trees for the sample galaxies, computed directly from the simulations with the methods described in Section 2, are shown in the top panel of Fig. 7. Each branch in the tree represents a progenitor galaxy that joins the main branch (shown in black) when the progenitor

falls into the host. The thickness of each branch is proportional to the number of accreted particles that the progenitor satellite contributes to the host at $z = 0$. The dashed lines indicate progenitors that contribute less than 1 per cent of the total number of accreted particles at $z = 0$. The middle and bottom panels show the merger trees as inferred from the clustering results obtained with the V-measure-optimized and DBCV-optimized HDBSCAN models, respectively. Here, a branch is included only if its corresponding progenitor is dominant in at least one identified cluster.

As expected, both models are more successful in recovering more massive mergers (those contributing > 1 per cent of all accreted particles, shown as solid lines and reported in Table 1). The major merger events that are missed are predominantly those that occur at high redshift ($z_{\text{acc}} \gtrsim 2$), with the exception of prog. 6355 (light green) in Au27. This progenitor merged with the main galaxy around the same time as prog. 205, a satellite about 10 times as massive; thus, its stellar debris was likely masked in the feature space by the dominant signature of the larger satellite, preventing its identification by the clustering algorithm.

4.2 Clustering in accreted + *in situ* stellar haloes

Fig. 8 shows the distribution of the HDBSCAN clusters in the $E - L_z$ plane obtained from the data set contaminated with *in situ* stars, which is a similar analysis done in Fig. 2 for the accreted-only case. In this figure, however, the noise points are not shown to avoid confusion with the *in situ* stars which are shown in grey. However, we include the fraction of points classified as noise by both V-measure and DBCV models. The figure also adopts the same colour scheme as used previously in Fig. 2.

Several observations can be made from the results of Fig. 8. First, the fraction of points classified as noise decreases substantially in comparison with the accreted-only case, as now the dense *in situ* population forms its own clusters. At the same time, the number of recovered accreted progenitors also decreases, as the particles belonging to these missing progenitors end up in the *in situ* clusters. Despite this contamination, the majority of identified clusters are still dominated by stars from a single accreted progenitor, and, as in the accreted-only analysis, the partitions from the V-measure and DBCV optimizations are generally very similar. A notable exception is Au7, where the V-measure-optimized model yields a significantly different partition (20 clusters, $f_{\text{noise}} = 0.60$) compared to the DBCV-optimized model (5 clusters, $f_{\text{noise}} = 0.06$). The key difference in the optimal parameters for Au7 lies in the `min_samples` value: 10 for the V-measure-optimized model and 5 for the DBCV-optimized one. By increasing `min_samples`, the density estimate of the algorithm becomes more conservative, requiring more neighbours for a point to be considered part of a dense core. The different clustering configurations resulting from the two values of `min_samples` can be explained in terms of the overlap in the clustering space between the oldest debris mergers and the *in situ* population.

Fig. 9 shows the Mahalanobis distance between the 12D distributions of the debris from the progenitors and the *in situ* population. This shows that the oldest major mergers (e.g. prog. 3210 in Au7, prog. 41397 in Au25, and prog. 7219 in Au27) have the largest overlap with the *in situ* population in the clustering space. Therefore, the particles from these mergers tend to be associated with the same cluster of the *in situ* population. A stricter density requirement (i.e. a higher value of `min_samples`) causes a fragmentation of the sparser regions of the *in situ* cluster. Because

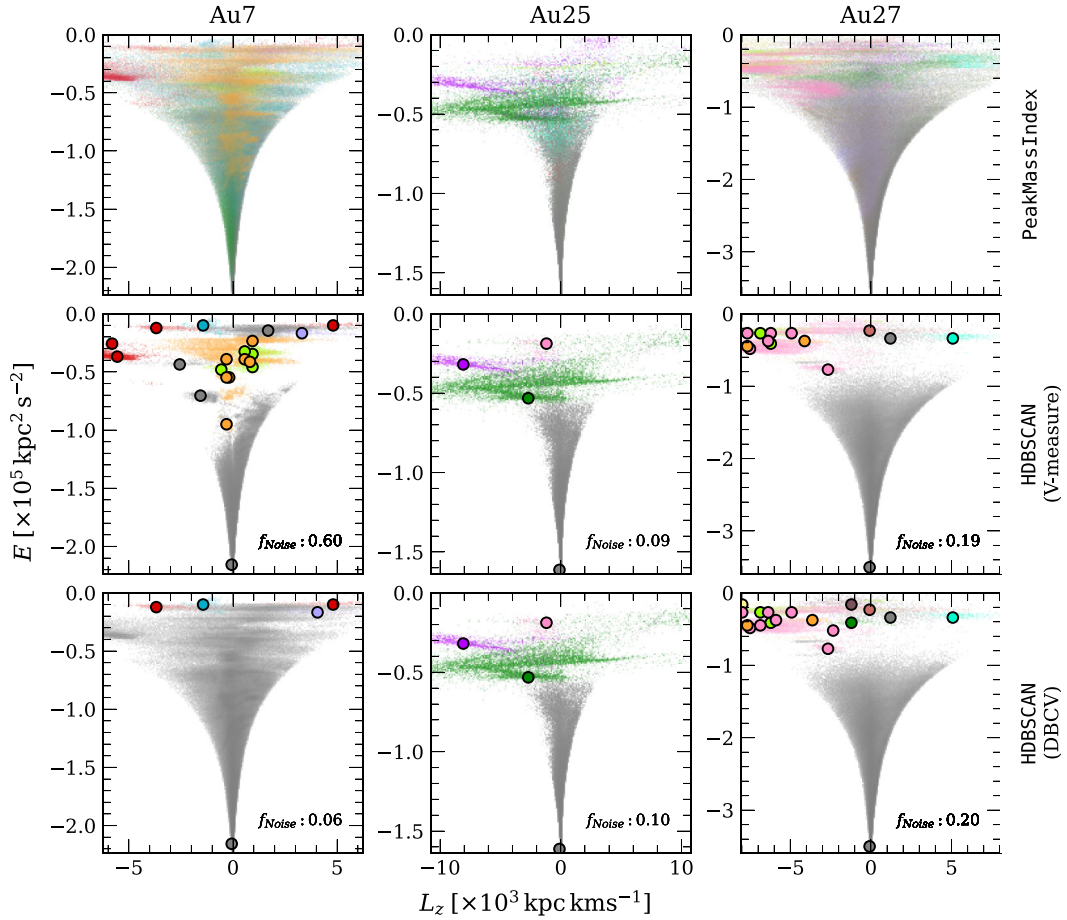


Figure 8. Density distributions in the $E - L_z$ plane for the accreted stars (similar to Fig. 2), but now including the *in situ* stars which are shown in grey. The clustering results from the optimized models are shown in the middle panel for the V-measure and the bottom panel for DBCV. The noise distributions are not shown, however we report the fraction of points classified as noise in each panel.

a larger number of progenitors in Au7 are mixed with the *in situ* ones than in the other galaxies, the fragmentation of the *in situ* clusters leads to the formation of groups dominated by accreted particles, hence increasing the overall V-measure score of the clustering configuration despite of a decrease in the cluster density.

4.2.1 Purity and completeness of clusters

The purity and completeness of the clusters identified in the *in situ* contaminated data sets are described through the precision and recall scores in Figs 10 and 11. Similarly to Figs 3 and 4, the clusters are represented by triangles if the debris from the corresponding dominant progenitor is in a form of stellar stream at $z = 0$, or by circles if it is in a phase-mixed structure. The colour of the cluster symbol is the same as used for the dominant progenitor in Fig. 8, while the size indicates the number of members in the cluster. We also show the clusters dominated by *in situ* particles as black crosses. In Figs 10 and 11, the precision and recall scores are plotted against the lookback infall time of their dominant progenitors; for *in situ* clusters, this is taken as the 90th percentile of the age distribution of the particles in the cluster. If an *in situ* cluster were mistakenly to be considered as a sample of

accreted stars, this statistic could be used as a proxy to the infall time of the relative progenitor galaxy.

As shown in Fig. 10, the contamination from *in situ* particles only slightly degrades the precision of the identified accreted clusters, with most still being composed of well over 50 per cent of stars from their dominant progenitor. The average precision for clusters associated with stellar streams and phase-mixed substructures is 0.91 and 0.81, respectively, for the V-measure optimized model, and 0.87 and 0.83 for the DBCV one. Conversely, the recall scores decreased on average: 0.23 (V-measure) and 0.27 (DBCV) for streams, 0.03 (V-measure) and 0.04 (DBCV) for phase-mixed structures. This is due to increased fragmentation, as shown in Fig. 11. A notable exception is Au25, whose larger progenitors, which appear as stellar streams, are recovered with similar or even higher recall scores than in the accreted-only analysis. This counterintuitive result likely arises from a contrast effect: the introduction of the dense, clustered *in situ* population makes the relatively less dense but coherent stellar streams stand out more clearly, allowing the algorithm to capture a larger fraction of their members. The *in situ* dominated clusters are also particularly pure and have high recall scores. This is probably due to both the efficiency of the algorithm of isolating the dense *in situ* component into one or a few coherent groups, and the significantly larger number of *in situ* particles compared to the accreted

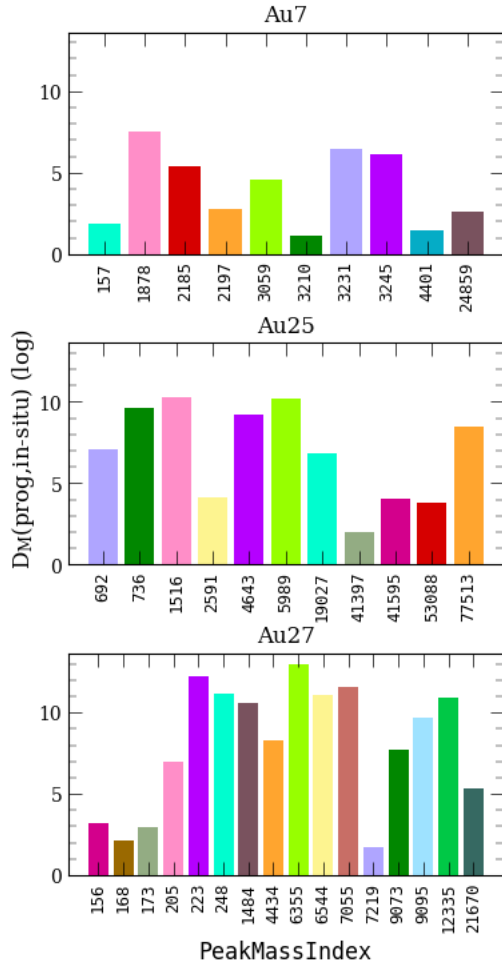


Figure 9. Mahalanobis distance between the particle distributions of the progenitor galaxies and the *in situ* population. Colours are associated to progenitors using the same colour scheme as in Fig. 2.

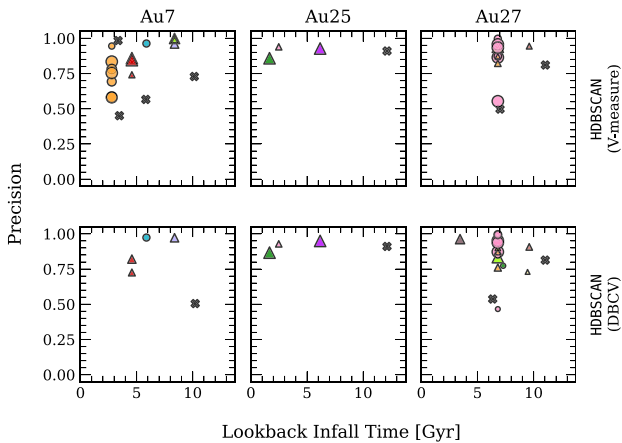


Figure 10. Precision scores for the clusters identified by the HDBSCAN models optimized with the V-measure and DBCV index metrics for the accreted + *in situ* haloes. Symbols and colours are the same as in Fig. 3, for the accreted-only haloes. Crosses indicate clusters that are dominated by *in situ* particles and their associated ‘lookback infall time’ is the 90th percentile of the age distribution of the cluster.

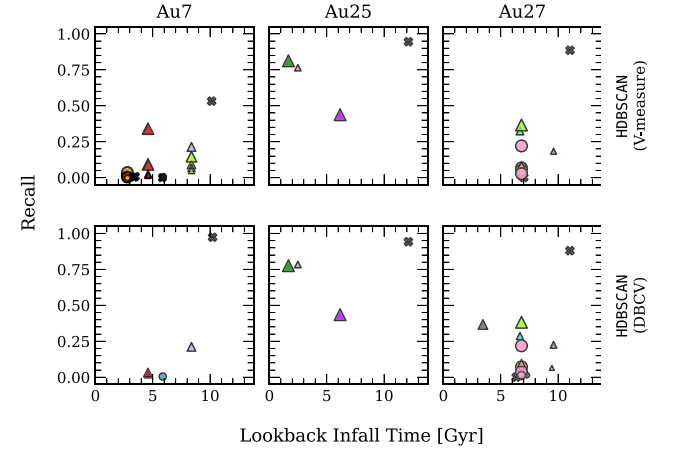


Figure 11. Recall scores for the clusters identified in accreted + *in situ* haloes by the HDBSCAN models optimized with the V-measure and DBCV index metrics. Symbols, sizes, and colours are the same as in Fig. 10.

ones in these regions of the clustering space. The ‘inferred’ lookback infall time for the *in situ* clusters, i.e. the 90th percentiles of the cluster age distributions, appears to vary significantly, which indicates that particles that formed in different star formation episodes also have different chrono-chemodynamical properties. In all galaxies, the only clusters with an associated lookback infall time > 10 Gyr are *in situ* dominated.

4.2.2 The assembly histories inferred from clustering

A comparison of the results with and without *in situ* contamination reveals that the primary impact of the *in situ* particles is not a degradation in the purity of recovered clusters, but rather the complete non-detection of certain progenitor groups. To better understand which mergers are more likely to be affected, we analyse the merger tree plots in Fig. 12. Similarly to Fig. 7, the actual merger history of the galaxies is shown in the top panel, while the middle and bottom panels only display those merger events which have been identified as the dominant progenitors of the HDBSCAN clusters. Again, the continuous lines indicate merger events that contribute at least 1 per cent of the total accreted stars in the galaxy (after the velocity selection cut) at $z = 0$, and the thickness of the line is proportional to the actual number.

This figure shows that none of the major progenitors accreted before $z_{\text{acc}} \sim 1$ (i.e. ~ 8 Gyr ago) are retrieved as clusters by the optimized HDBSCAN model. This suggests that the clustering model may identify only the cold stellar streams, as also seen in Figs 3 and 4. Therefore, in the presence of *in situ* stars, the model can retrieve only the most recent fraction of the overall merger history of the galaxy. This result is in clear contrast to the results obtained in Section 4.1 for accreted-only haloes, where the model could retrieve some of the most ancient merger events. Hence, reducing the contamination of *in situ* stars in the sample of halo stars is crucial for an effective recovery of old ($z_{\text{acc}} > 1$) merger events when using HDBSCAN.

4.3 The impact of stellar ages

As already mentioned, in our analysis clusters are obtained from the full chrono-chemodynamical properties of the accreted particles. However, replicating this parameter space for stars in the

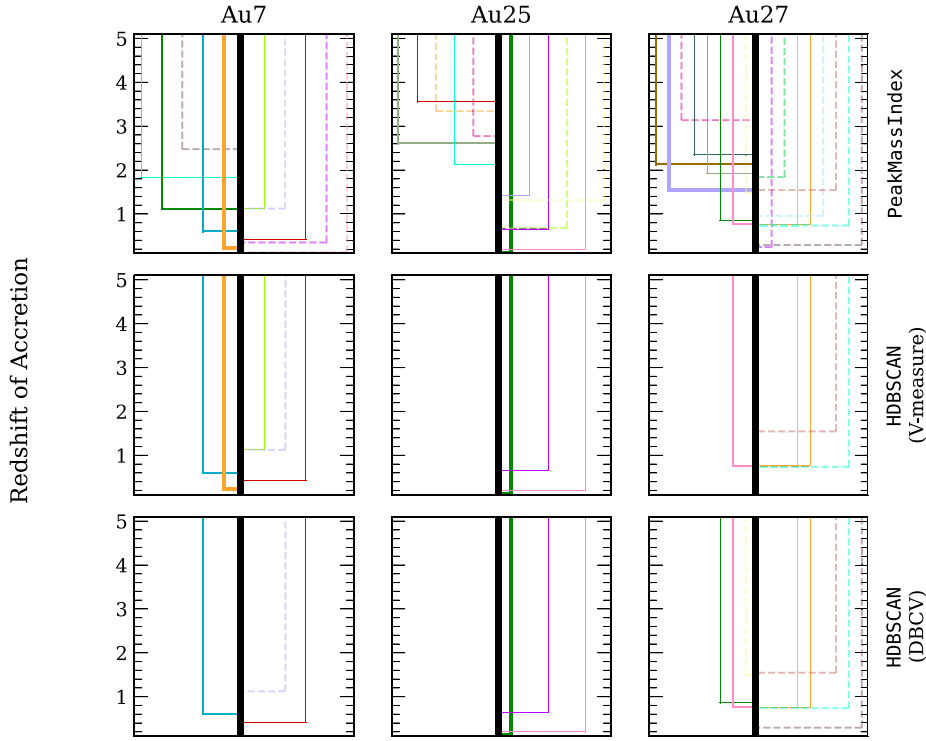


Figure 12. Merger trees for the galaxies in the sample as reconstructed by the V-measure (middle) and DBCV (bottom) optimized HDBSCAN models when contamination from *in situ* stars is included. Top panel shows the full merger trees. The correspondence between lines and progenitors is the same as described in Fig. 7 for the accreted-only case. The style and thickness of the lines reflects the new fraction of accreted stars from each progenitor after the velocity selection cut is applied.

Milky Way might be challenging. Although kinematics measurements are accurately provided for tens of millions of stars by *Gaia* and self-consistent chemical abundances are becoming available for a growing number of stars through the combination of data from different spectroscopic surveys (G. F. Thomas et al. 2024; D. Horta et al. 2025), stellar ages remain accessible only for a fraction of them. Therefore, a more realistic application of our methodology would involve only the dynamical and chemical properties of the stars in the Milky Way. In order to assess the quality of the merger tree reconstruction under these conditions, we repeat the analysis discussed in Sections 4.1 and 4.2 excluding the ages of the star particles (τ_{form}) from the vector of stellar properties describing the clustering space. For the sake of conciseness, we consider only Au27, which is the galaxy with the assembly history most similar to the Milky Way in our sample, for this part of the analysis.

Similarly to Figs 2 and 8, the distribution in the $E - L_z$ plane of the star particles in Au27 coloured based on their HDBSCAN group are reported in Fig. 13 for both the clustering considering only accreted stars (left plots) and including *in situ* contamination (right plots). In both cases, the HDBSCAN parameters have been fine-tuned to adapt to the new clustering space without the age information. The top plots show the results of the clustering model optimised using the V-measure metric, while the bottom ones refer to the optimisation obtained with the DBCV metric. By comparing Fig. 13 with Figs 2 and 8, it can be noticed that excluding the stellar ages results in a larger number of clusters. The clustering model optimized with V-measure returns 30 and 38 clusters in this scenario for the accreted and accreted + *in situ*

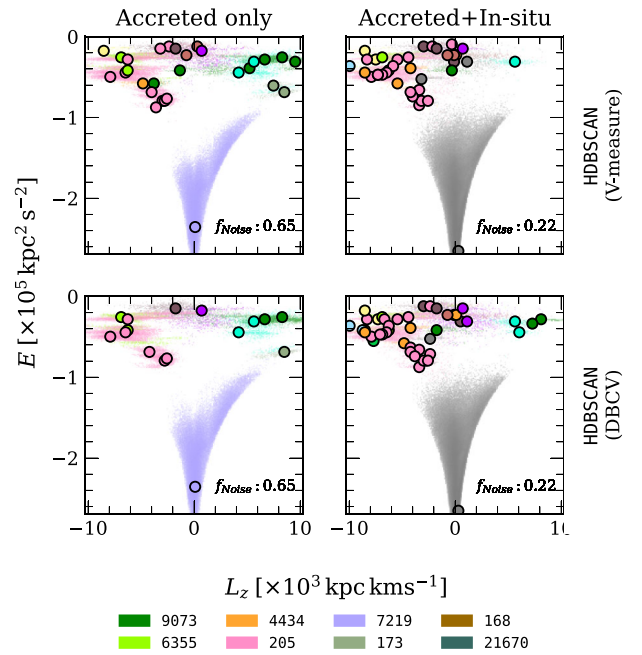


Figure 13. Density distribution in the $E - L_z$ plane for the stars in Au27 colour-coded by their HDBSCAN label. The clustering analysis is repeated excluding the age of the star particles from both the accreted-only and accreted + *in situ* data sets. The plot format follows Figs 2 and 8.

data sets, while only 11 and 15 clusters, respectively, are identified by the same model when the age information is provided. Similarly, the DBCV-optimized model returns a larger number of clusters for both the accreted-only (17 versus 16) and accreted + *in situ* (46 versus 20) cases. This effect can be explained as fragmentation of the progenitors identified in the analysis with stellar ages into a larger number of clusters and a consequent decrease in the average completeness of the retrieved groups in both the accreted-only (0.15 versus 0.29, V-measure and 0.24 versus 0.19, DBCV) and accreted + *in situ* (0.08 versus 0.16, V-measure and 0.09 versus 0.15, DBCV) data sets. The average purity of the groups is similar whether the clustering is performed with (0.78, V-measure, and 0.86, DBCV, for the accreted-only stars and 0.88, V-measure, and 0.87, DBCV, for the accreted + *in situ* stars) or without the age (0.90, V-measure, and 0.84, DBCV, for the accreted-only stars and 0.87, V-measure, and 0.87, DBCV, for the accreted + *in situ* stars) of the star particles. Therefore, the identification of pure samples of stars accreted from a single progenitors is still possible without the knowledge of their formation time. However, the inclusion of such information is important for grouping the accreted stars in a number of sub-structures more representative to the number of actual progenitor galaxies. Therefore, if stellar ages cannot be determined, the development of a methodology for further grouping the star samples identified by HDBSCAN is crucial for a correct inference of the Galaxy assembly history.

5 DISCUSSION

Our analysis shows that an optimized HDBSCAN model can produce pure clusters of accreted stars even in environments strongly contaminated by the *in situ* population (see Fig. 10). This is in contrast with the results of G. F. Thomas et al. (2025) who use the single-linkage clustering methodology developed by S. S. Lövdal et al. (2022) and find that it predominantly returns clusters composed in majority by *in situ* stars or by a combination of debris from different progenitors (see section 4.1 in G. F. Thomas et al. 2025). We believe that the optimization of the clustering algorithm is the crucial factor in correctly identifying the clusters made of accreted stars. Also, while the S. S. Lövdal et al. (2022) clustering algorithm uses only properties related to IoM, our study uses an extended, 12D parameter (feature) space that combines chemistry, kinematics, and ages. There are, of course, other factors that may explain the discrepancies between the results of the two methods. The S. S. Lövdal et al. 2022 methodology imposes the creation of ‘background haloes’ in order to determine the significance score of the clusters, which may introduce artificial density variations in the clustering space. There are also differences between the two Auriga data sets in terms of their numerical resolution, as G. F. Thomas et al. (2025) use the ‘level 3’ suite, which are higher in resolution than the ‘level 4’ suite used in our analysis. As the stellar masses of galaxies, both *in situ* and accreted, increase at higher resolution (see e.g. R. J. J. Grand et al. 2021 and A. H. Riley et al. 2025), it is possible that the contamination by *in situ* stars in the sample also increases. Moreover, although both studies use samples of star particles with similar halo-type kinematics, G. F. Thomas et al. (2025) focus on solar neighbourhood regions, whereas we study the entire haloes.

Our optimized HDBSCAN model can retrieve mergers up to $z_{\text{acc}} \sim 3$ (see Fig. 7). However, when applied to a data set contaminated by *in situ* stars, only galaxies accreted after $z_{\text{acc}} \simeq 1$ are retrieved as dominant progenitors of the clusters. From an

observational point of view, this suggests that accretion events > 10 Gyr old may not be found by clustering methods, or that any clusters from this type of events are likely dominated by *in situ* material. This suggests that observational samples representing early accretions in the Galaxy, such as GES (V. Belokurov et al. 2018; A. Helmi et al. 2018), Kraken (J. M. D. Kruijssen et al. 2019), or Shiva & Shakti (K. Malhan & H.-W. Rix 2024) could be significantly contaminated by *in situ* stars.

These findings are similar to the results of G. F. Thomas et al. (2025), who show that their methodology is effective for mergers that occurred in the last 6–7 Gyr. However, because the purity of the HDBSCAN clusters does not change significantly between the accreted-only and accreted + *in situ* data sets, we argue that the presence of the *in situ* population in the clustering data set mainly impacts the number and type of detectable mergers. We note that the problem could be attenuated when using observations, as the Auriga simulated haloes may have a higher fraction of *in situ* stars than measured in the Milky Way (A. Monachesi et al. 2019). Therefore, we expect that a higher fraction of the assembly history of the Milky Way could be retrieved with our methodology than implied by the results from clustering on accreted + *in situ* stars in Section 4.2. Moreover, there are other machine learning methods that can very efficiently separate *in situ* stars from accreted stars (B. Ostdiek et al. 2020; A. Sante et al. 2024), before applying any clustering methods. Using these methods in combination, one would expect a higher rate of success in finding sub-structures, even in the case of heavy contamination by *in situ* stars.

In Sections 4.1–4.2, the clusters were obtained from the full (12D) chrono-chemodynamical properties of the accreted particles. However, replicating this 12D parameter space for stars in the Milky Way may be challenging. Although the kinematical measurements are accurately provided for tens of millions of stars by *Gaia* and self-consistent chemical abundances are becoming available for a growing number of stars through the combination of data from different spectroscopic surveys (G. F. Thomas et al. 2024; D. Horta et al. 2025), stellar ages remain accessible only for a fraction of them. In Section 4.3, we examined a more realistic scenario where the ages were not available in the set of features. We found that, on average, the purity of the clusters was not affected, but it was a slight decrease in their completeness. Overall, the cluster became ever more fragmented.

Furthermore, the methodology presented in this paper does not account for any uncertainty or incompleteness in the stellar properties. Hence, more studies are needed to understand the implications of this missing information on the capability of HDBSCAN to recover past mergers.

6 CONCLUSIONS

We have developed and validated a new methodology for optimizing the parameters of the HDBSCAN clustering algorithm for identifying tidal debris in Milky Way-mass galaxies. Our approach is designed to maximize the recovery of coherent structures in chrono-chemodynamical space, and we demonstrate that this optimization simultaneously enhances both the purity and completeness of the resulting stellar groups (see Figs 2 and 8), ensuring they correspond to the debris of single progenitors.

To test the power of this method, we applied it to three simulated Milky Way-mass galaxies from the Auriga suite, which feature diverse assembly histories. We first established a baseline performance in an idealized scenario, using only

accreted star particles. In these ideal conditions, we found that

(i) The vast majority of clusters identified by HDBSCAN are dominated by stars originating from a single accreted progenitor, confirming the ability of the algorithm to trace individual merger events.

(ii) Cluster members represent only a small fraction of the total number of the accreted particles associated to a progenitor, as particles are lost as noise points or, for massive systems, split into different clusters due to the wide range in the chronochemodynamical properties of their debris.

(iii) The major mergers of a galaxy can be retrieved as clusters up to a redshift of accretion of $z_{\text{acc}} \sim 3$ unless a more recent massive merger has occurred (see Fig. 7).

When we advance to a more realistic scenario that includes contamination of the *in situ* population, the performance of the algorithm remains high. The purity and completeness of the recovered clusters decrease only marginally (see Figs 10 and 11). The primary challenge introduced by the *in situ* stars is a reduction in the lookback time for merger detection. The HDBSCAN algorithm confidently identifies only the most recent major accretion events ($z_{\text{acc}} < 1$), as older, more phase-mixed structures become difficult to distinguish in the chrono-chemodynamical plane (see Fig. 12).

In conclusion, our work demonstrates that density-based clustering can successfully identify the main merger events of a Milky Way-like galaxy, provided that comprehensive information on its stellar content is available and the algorithm parameters are optimized. The application of this methodology to the Milky Way itself looks encouraging; however, significant challenges remain. Future work must rigorously quantify the performance of the algorithm with incomplete data sets, which reflect the current observational reality, where, for instance, radial velocities or detailed chemical abundances are not available for every star. Furthermore, developing more sophisticated techniques to distinguish the *in situ* and accreted populations a priori remains critical to uncover the earliest epoch of Galaxy formation.

ACKNOWLEDGEMENTS

We thank the reviewer for their suggestions and constructing feedback which improved the clarity of the paper. We also thank the Auriga team for providing public access to the particle and merger trees data for their simulations. AS acknowledges a Science Technologies Facilities Council (STFC) Ph.D. 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.

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, in Proc. 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.
- Ankerst M., Breunig M. M., Kriegel H.-P., Sander J., 1999, *SIGMOD record*, 28, 49
- Beers T. C. et al., 2012, *ApJ*, 746, 34
- 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
- Bland-Hawthorn J., Gerhard O., 2016, *ARA&A*, 54, 529
- Borsato N. W., Martell S. L., Simpson J. D., 2020, *MNRAS*, 492, 1370
- Brown A. G. A., Velázquez H. M., Aguilar L. A., 2005, *MNRAS*, 359, 1287
- Buder S. et al., 2021, *MNRAS*, 506, 150
- Buder S. et al., 2025, *PASA*, 42, e051
- Campello R. J. G. B., Moulavi D., Sander J., 2013, in Pei J., Tseng V. S., Cao L., Motoda H., Xu G. eds, Advances in Knowledge Discovery and Data Mining. Springer Berlin Heidelberg, Berlin, Heidelberg, p. 160
- Carollo D. et al., 2007, *Nature*, 450, 1020
- Carrillo A., Deason A. J., Fattahi A., Callingham T. M., Grand R. J. J., 2024, *MNRAS*, 527, 2165
- Chabrier G., 2003, *PASP*, 115, 763
- Conroy C. et al., 2019, *ApJ*, 883, 107
- Cooper A. P., Parry O. H., Lowing B., Cole S., Frenk C., 2015, *MNRAS*, 454, 3185
- Deason A. J., Belokurov V., 2024, *New Astron. Rev.*, 99, 101706
- Di Matteo P., Haywood M., Lehnert M. D., Katz D., Khoperskov S., Snaith O. N., Gómez A., Robichon N., 2019, *A&A*, 632, A4
- Dodd E., Callingham T. M., Helmi A., Matsuno T., Ruiz-Lara T., Balbinot E., Lövdal S., 2023, *A&A*, 670, L2
- Ester M., Kriegel H.-P., Sander J., Xu X., 1996, in Proc. of 2nd International Conference on Knowledge Discovery. p. 226
- 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., Crain R. A., Theuns T., Schaye J., Wiersma R. P. C., Dalla Vecchia C., 2011, *MNRAS*, 416, 2802
- Font A. S., McCarthy I. G., Belokurov V., Brown S. T., Stafford S. G., 2022, *MNRAS*, 511, 1544
- Fraix-Burnet D., Davoust E., 2015, *MNRAS*, 450, 3431
- Freeman K., Bland-Hawthorn J., 2002, *ARA&A*, 40, 487
- Gaia Collaboration, 2016, *A&A*, 595, A1
- Gaia Collaboration, 2023, *A&A*, 674, A1
- Gallart C. et al., 2024, *A&A*, 687, A168
- Gilmore G. et al., 2012, *The Messenger*, 147, 25
- Gómez F. A., Helmi A., Brown A. G. A., Li Y.-S., 2010, *MNRAS*, 408, 935
- Gómez F. A., Helmi A., Cooper A. P., Frenk C. S., Navarro J. F., White S. D. M., 2013, *MNRAS*, 436, 3602
- Grand R. J. J. et al., 2017, *MNRAS*, 467, 179
- Grand R. J. J. et al., 2021, *MNRAS*, 507, 4953
- Grand R. J. J., Fragkoudi F., Gómez F. A., Jenkins A., Marinacci F., Pakmor R., Springel V., 2024, *MNRAS*, 532, 1814
- Haywood M., Di Matteo P., Lehnert M. D., Snaith O., Khoperskov S., Gómez A., 2018, *ApJ*, 863, 113
- Helmi A., 2020, *ARA&A*, 58, 205
- Helmi A., White S. D. M., 1999, *MNRAS*, 307, 495
- Helmi A., de Zeeuw P. T., 2000, *MNRAS*, 319, 657
- 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
- Horta D. et al., 2023, *MNRAS*, 520, 5671
- Horta D., Price-Whelan A. M., Hogg D. W., Ness M. K., Casey A. R., 2025, *AJ*, 169, 314
- Jean-Baptiste I., Di Matteo P., Haywood M., Gómez A., Montuori M., Combes F., Semelin B., 2017, *A&A*, 604, A106

- Ji A. P. et al., 2020, *AJ*, 160, 181
- Johnston K. V., Bullock J. S., Sharma S., Font A., Robertson B. E., Leitner S. N., 2008, *ApJ*, 689, 936
- Karakas A. I., 2010, *MNRAS*, 403, 1413
- Kawata D., Mulchaey J. S., 2008, *ApJ*, 672, L103
- Kim B. et al., 2025, *MNRAS*, 540, 264
- Kizhuprakkat N. et al., 2024, *MNRAS*, 531, 4108
- Klement R. et al., 2009, *ApJ*, 698, 865
- 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., Reina-Campos M., Crain R. A., Bastian N., 2019, *MNRAS*, 486, 3180
- Lövdal S. S., Ruiz-Lara T., Koppelman H. H., Matsuno T., Dodd E., Helmi A., 2022, *A&A*, 665, A57
- MacQueen J., 1967, in Proc. 5th Berkeley Symposium on Mathematical Statistics and Probability. p. 281
- Mackereth J. T. et al., 2019, *MNRAS*, 482, 3426
- Majewski S. R. et al., 2017, *AJ*, 154, 94
- Malhan K., Rix H.-W., 2024, *ApJ*, 964, 104
- McCarthy I. G., Font A. S., Crain R. A., Deason A. J., Schaye J., Theuns T., 2012, *MNRAS*, 420, 2245
- McInnes L., Healy J., Astels S., 2017, *J. Open Source Softw.*, 2, 205
- 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
- Morrison H. L. et al., 2009, *ApJ*, 694, 130
- Moulavi D., Jaskowiak P. A., Campello R. J. G. B., Zimek A., Sander J., 2014, in *SDM*.
- Murtagh F., Contreras P., 2012, *WIREs: Data Min. Knowl. Discov.*, 2, 86
- 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
- Necib L. et al., 2020, *Nat. Astron.*, 4, 1078
- Orkney M. D. A. et al., 2022, *MNRAS*, 517, L138
- Orkney M. D. A. et al., 2023, *MNRAS*, 525, 683
- Ostdiek B. et al., 2020, *A&A*, 636, A75
- Pakmor R., Springel V., 2013, *MNRAS*, 432, 176
- Planck Collaboration XVI, 2014, *A&A*, 571, A16
- Portinari L., Chiosi C., Bressan A., 1998, *A&A*, 334, 505
- Randich S., Gilmore G., *Gaia-ESO Consortium*, 2013, *The Messenger*, 154, 47
- Re Fiorentin P., Lattanzi M. G., Spagna A., Curir A., 2015, *AJ*, 150, 128
- Riley A. H. et al., 2025, *MNRAS*, 542, 2443
- Rosenberg A., Hirschberg J., 2007, in Eisner J. ed., *EMNLP-CoNLL. ACL*, p. 410
- Ruiz-Lara T., Matsuno T., Lövdal S. S., Helmi A., Dodd E., Koppelman H. H., 2022, *A&A*, 665, A58
- 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
- Shipp N. et al., 2025, *MNRAS*, 542, 1109
- 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., 2005, *MNRAS*, 364, 1105
- 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
- Thielemann F.-K. et al., 2003, *Nucl. Phys. A*, 718, 139
- Thomas G. F. et al., 2024, *A&A*, 690, A54
- Thomas G. F., Battaglia G., Grand R. J. J., Aguiar Álvarez A., 2025, *A&A*, 704, A40
- Tissera P. B., Scannapieco C., Beers T. C., Carollo D., 2013, *MNRAS*, 432, 3391
- Travaglio C., Hillebrandt W., Reinecke M., Thielemann F.-K., 2004, *A&A*, 425, 1029
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Yuan Z. et al., 2020, *ApJ*, 891, 39
- Zhao G., Zhao Y.-H., Chu Y.-P., Jing Y.-P., Deng L.-C., 2012, *Res. Astron. Astrophys.*, 12, 723
- Zolotov A., Willman B., Brooks A. M., Governato F., Brook C. B., Hogg D. W., Quinn T., Stinson G., 2009, *ApJ*, 702, 1058

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