

Applying machine learning to Galactic Archaeology: how well can we recover the origin of stars in Milky Way-like galaxies?

Andrea Sante,¹★ Andrea S. Font,^{1,2} Sandra Ortega-Martorell^{1,2}, Ivan Olier^{1,2} and Ian G. McCarthy^{1,2}

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

²*Data Science Research Centre, Liverpool John Moores University, 3 Byrom Street, Liverpool L3 3AF, UK*

Accepted 2024 May 31. Received 2024 May 23; in original form 2024 April 25

ABSTRACT

We present several machine learning (ML) models developed to efficiently separate stars formed *in situ* in Milky Way-type galaxies from those that were formed externally and later accreted. These models, which include examples from artificial neural networks, decision trees, and dimensionality reduction techniques, are trained on a sample of disc-like, Milky Way-mass galaxies drawn from the ARTEMIS cosmological hydrodynamical zoom-in simulations. We find that the input parameters which provide an optimal performance for these models consist of a combination of stellar positions, kinematics, chemical abundances ([Fe/H] and [α /Fe]), and photometric properties. Models from all categories perform similarly well, with area under the precision–recall curve (PR-AUC) scores of $\simeq 0.6$. Beyond a galactocentric radius of 5 kpc, models retrieve > 90 per cent of accreted stars, with a sample purity close to 60 per cent, however the purity can be increased by adjusting the classification threshold. For one model, we also include host galaxy-specific properties in the training, to account for the variability of accretion histories of the hosts, however this does not lead to an improvement in performance. The ML models can identify accreted stars even in regions heavily dominated by the *in-situ* component (e.g. in the disc), and perform well on an unseen suite of simulations (the AURIGA simulations). The general applicability bodes well for application of such methods on observational data to identify accreted substructures in the Milky Way without the need to resort to selection cuts for minimizing the contamination from *in-situ* stars.

Key words: methods: data analysis – methods: numerical – Galaxy: abundances – Galaxy: kinematics and dynamics – solar neighbourhood – Galaxy: stellar content.

1 INTRODUCTION

In a lambda-cold dark matter (Λ CDM) cosmological model, large galaxies like the Milky Way (MW) form through a hierarchical process, with smaller structures merging progressively into larger ones (Searle & Zinn 1978; White & Rees 1978). In this framework, signatures of past accretion and disruption events in the Galaxy are left imprinted in a multidimensional parameter space, composed of positions, and kinematics and chemical abundances of stars (e.g. Helmi 2020). From the information gathered about tidal streams in this multidimensional parameter space, one can reconstruct the assembly history of the MW (Freeman & Bland-Hawthorn 2002), that is, determine the timing of the accretion events, the masses of the progenitor galaxies, their star formation histories or the orbital properties.

A multitude of methods have been devised to find tidal stellar streams. Streams from massive accretions can be usually detected from photometry, as they tend to be brighter and to be spatially extended. Those from low-mass progenitors may sometimes appear as coherent structures in physical space (i.e. in positions and velocities of stars), particularly if the accretions were recent (e.g. Bullock & Johnston 2005; Johnston et al. 2008) or in a special geometry

(Johnston, Hernquist & Bolte 1996). Over longer dynamical time-scales, however, the streams tend to disperse due to phase-mixing (Tremaine 1999) and thus they become increasingly difficult to distinguish from background field stars. Information about them may be still retained in the integrals of motion related to angular momenta and total energies of their orbits (see Binney & Tremaine 2008). Methods have been devised to identify tidal streams as ‘clumps’ in the energy (E) and angular momentum (L_z) space (Helmi & White 1999; Gómez et al. 2010), as these quantities are (quasi-)conserved through time. Methods that rely on finding specific patterns of tidal streams in the velocity space (Johnston, Spergel & Haydn 2002; Gómez et al. 2010; Koppelman & Helmi 2021), or in the angles, actions or frequencies (McMillan & Binney 2008; Bovy 2014; Sanders & Binney 2016; Malhan et al. 2022) have also been used to identify streams in the Galaxy. However, many of these methods require knowledge of the gravitational potential of the MW, although more recent techniques, such as STREAMFINDER (Malhan & Ibata 2018), do not rely on such assumptions.

Other parameters have also been used to improve the detection. For example, methods based on match filters that weigh the colour–magnitudes of stars (Grillmair et al. 1995; Rockosi et al. 2002; Balbinot et al. 2011) have proven useful in the detection of new streams (Shipp et al. 2019). More widely used are methods employing the chemical abundances of stars, building on the expectation that stars formed in a given progenitor share similar chemical ‘finger-

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

prints’ even if their spatial and kinematical information has become phase-mixed. In combination with 6D physical space parameters (positions and kinematics), the principle of ‘chemical tagging’ (Freeman & Bland-Hawthorn 2002) has been applied successfully in the discovery or characterization of tidal debris (Belokurov et al. 2018; Helmi et al. 2018; Koppelman et al. 2019; Kruijssen et al. 2019; Das, Hawkins & Jofré 2020; Ji et al. 2020; Naidu et al. 2022; Horta et al. 2023; Malhan & Rix 2024), or in the reconstruction of the early stages of the MW (Belokurov & Kravtsov 2024).

In the era of large Galactic surveys, such as *Gaia*, new methods based on machine learning (ML) techniques have been developed and have proven viable. Clustering algorithms, such as density-based clustering non-parametric algorithms (see Ester et al. 1996) have been applied to chemodynamical data to confirm existing discoveries or to reveal new ones (Koppelman et al. 2019; Borsato, Martell & Simpson 2020). An unsupervised learning method, called VIA MACHINAE, was also developed, using conditional density estimation and sideband interpolation to find local overdensities (Shih et al. 2022). Veljanoski et al. (2019) developed a gradient boosted trees model to identify halo stars based on astrometric and photometric data from the *Gaia* Universe Model Snapshot (Robin et al. 2012).

The discovery of new tidal streams is becoming increasingly more difficult, in part due to current limitations in the current methods. For example, many stream-finding methods assume that the halo is mostly of accreted origin. This is a reasonable assumption only for the outer parts of the Galaxy ($\gtrsim 20$ kpc), which are less contaminated by disc stars. Observations indicate that the stellar halo has a ‘dual nature’ (Carollo et al. 2007, 2010; Beers et al. 2012), where the two components, accreted and *in situ*, overlap over some distance. Disentangling the two components is important in order to construct a relatively clean sample of accreted stars on which to apply the stream detection methods. The two components differ in spatial distribution, kinematics and metallicity (Carollo et al. 2007); specifically, the *in-situ* halo is more centrally concentrated than the accreted and tends to have an overall rotating motion prograde with the disc, whereas the orbits of accreted stars are more randomly distributed; generally, the *in-situ* halo is also more metal-rich than the accreted. This suggests that the two halo components could be, in principle, clearly separated. In practice, however, most observational samples include some selection criteria which are meant to minimize the contamination from both the disc and *in-situ* halo.

The observational selection cuts usually relate to spatial location, kinematics or chemistry (or a combination thereof). In some cases, the criteria are purposefully conservative, for example, selecting stars only on retrograde orbits to search for debris. This kinematical cut has proven beneficial for the discovery of many substructures in the solar neighbourhood (Koppelman et al. 2019), including a debris from a massive progenitor, called *Gaia* Enceladus/Sausage (Belokurov et al. 2018; Helmi et al. 2018), which merged with the Galaxy ≈ 8 – 9 Gyr ago. However, many tidal streams are predicted to still remain hidden (Shipp et al. 2023). Therefore, relaxing the selection criteria for observational samples could lead to more debris discoveries, especially in the less explored regions of the Galaxy, such as the heated stellar disc (Mackereth et al. 2019; Belokurov et al. 2020).

The dual nature of stellar haloes is retrieved naturally in cosmological hydrodynamical simulations (Zolotov et al. 2009; McCarthy et al. 2012; Tissera et al. 2013; Cooper et al. 2015; Pillepich, Madau & Mayer 2015; Monachesi et al. 2016; Pillepich et al. 2018; Brook et al. 2020). However, these simulations are too general to inform the precise selection cuts that can be applied to observations in the MW, as they model systems with a variety of different accretion

histories none of which is expected to exactly match that of the MW. Moreover, depending on the implementation of subgrid physical prescriptions, simulations may predict different properties for the *in-situ* halo component. This is related to the different formation channels of *in-situ* halo stars in simulations: either as stars ejected from the galaxy disc, or formed within filaments of cold gas, or in the wakes of stripped gas from infalling satellites.

Rather than using simulations to inform selection cuts, one can use them to train ML models to separate the accreted from *in-situ* stars more accurately and hopefully in a way which is sufficiently general to apply to the observational data. ML provides an ideal framework to find out the relations between objects belonging to different classes by leveraging the information hidden in large data sets. This is particularly useful in those regions of the parameter space where the two halo components overlap, for example, in the case of accreted stars that overlap with the disc (e.g. Hawkins et al. 2015), or of the old, *in-situ* halo stars that may overlap in metallicity with some of the accreted substructure.

Several ML techniques have been developed recently to separate the two components. For example, by analysing a *Gaia* mock catalogue constructed from the FIRE simulations (Sanderson et al. 2020), Ostdiek et al. (2020) trained an artificial neural network (ANN) to classify accreted and *in-situ* stars based on 5D kinematics and then fine-tuning the model on a *Gaia* DR2/RAdial Velocity Experiment (RAVE) data set. This has led to the discovery of a new substructure in the MW, called Nyx (Necib et al. 2020). Recently, Tronrud et al. (2022) developed an ANN to separate accreted and *in-situ* stars, and trained it on the chemical abundances and ages of stars in AURIGA simulations (Grand et al. 2017). A similar method has been developed by Trujillo-Gomez et al. (2023) to classify accreted and *in-situ* globular clusters in the E-MOSAICS simulations (Pfeffer et al. 2018), using as inputs 17 observable properties, including some of the associated host galaxies.

Rather than focusing on the description of a single methodology, here we perform an extensive comparison of different ML algorithms, in order to decide which ones are more suitable for the classification of accreted and *in-situ* stars. We also include a wide range of input parameters the models, chosen as stellar parameters which can be directly observed from MW surveys, such as positions, kinematics, ages, chemical abundances ([Fe/H] and $[\alpha/\text{Fe}]$), and photometric properties. Our aim is to identify the optimal, data-driven model that can automatically identify accreted stars in observational samples of the MW. The ML algorithms we consider can be grouped broadly into three categories: ANNs, decision trees, and dimensionality reduction methods. The ANNs and decision trees are used for developing the classification models, whereas the dimensionality reduction technique is used for visualizing the data and providing insights into the output of the models. As a benchmark for comparison of models, we adopt an ANN model which resembles the ‘Galactic Archaeology Neural Network’ (GANN) of Tronrud et al. (2022). For training and testing the ML models, we use a sample MW-mass galaxies from the ARTEMIS simulations (Font et al. 2020), selected to be disc-like today, but with different accretion histories. Since in simulations the origin of each star particle is already known (i.e. whether it formed *in situ* or was accreted), we can assess the performance of the techniques more accurately than by testing them on observations.

The paper is organized as follows. Section 2 provides a brief description of the ARTEMIS simulations and of the sample of disc-like galaxies which are used for training and testing. The selection of the physical parameters considered as inputs for the models is described in Sections 2.1 and 2.2. In Section 3, we present the ML methods;

the metrics used for evaluating the classification are discussed in Section 3.1, while the methods are described individually in Section 3.2, including a description of how we determine the optimal set of input parameters (in Section 3.2.2). Section 4 includes a comparison of the performance of the ML models (Section 4.1) and shows how the *in-situ* and accreted stars identified by these methods are separated in a chemodynamical phase space (Section 4.2); it also shows how ML models may improve the detection of accreted stars in areas omitted by observational selection cuts (Section 4.3); and illustrates how the predictions of the models can be visualized with UMAP (Uniform Manifold Approximation and Projection, Section 4.4). In Section 5, we apply our ML methods on a different suite of simulations (AURIGA), to further test their performance. The conclusions of our study are summarized in Section 6.

2 THE ARTEMIS SIMULATIONS

ARTEMIS is a suite of zoomed-in, high-resolution cosmological hydrodynamical simulations of 45 MW-mass systems (Font et al. 2020; Font, McCarthy & Belokurov 2021) in a flat Λ CDM cosmological model with the following parameters measured from the Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP9) observations (Hinshaw et al. 2013): $\Omega_m = 0.2793$, $\Omega_b = 0.0463$, $h = 0.70$, $\sigma_8 = 0.8211$, and $n_s = 0.972$. The MW-mass systems have total masses ranging between $8 \times 10^{11} < M_{200}/M_\odot < 2 \times 10^{12}$, where M_{200} is the mass enclosing a mean density of 200 times the critical density of the universe at present time. The dark matter particle masses are $1.17 \times 10^5 M_\odot h^{-1}$, the initial gas particle masses are $2.23 \times 10^4 M_\odot h^{-1}$, and the (Plummer equivalent) force resolution is $125 \text{ pc } h^{-1}$.

The simulations were run with the GADGET-3 code (Springel et al. 2005), including an updated hydrodynamical solver and subgrid physical prescriptions developed for the EAGLE project (Schaye et al. 2015). The physical prescriptions for subgrid physics include metal-dependent radiative cooling in the presence of a photoionizing ultraviolet background, star formation, stellar and chemical evolution, formation of supermassive black holes, and feedback from supernova, stellar winds, and active galactic nuclei (see Crain et al. 2015 and Schaye et al. 2015 for details). The chemical enrichment model follows 11 element species taking in account the mass loss from asymptotic giant branch stars, stellar winds, and both core collapse and thermonuclear supernovae. The heavy elements relevant for our study, are Fe (produced mainly in Type Ia SNe) and Mg (an α -element, produced in Type II SNe).

Font et al. (2020) also computed, in post-processing, the optical properties of the simulated galaxies in ARTEMIS. These include luminosities, magnitudes, and colours in various passbands, assuming each star particle is an a single stellar population (SSP), by using the PARSEC v1.2S + COLIBRI PR16 isochrones (Bressan et al. 2012; Marigo et al. 2017) and a Chabrier (Chabrier 2003) initial mass function. In this study, we are mainly interested in *Gaia* photometric properties, which were not computed originally. For this, we use the available magnitudes for the Sloan Digital Sky Survey (SDSS) and convert them to the *Gaia* G , G_{BP} , and G_{RP} passbands equivalents, using the photometric relationships from Busso et al. (2022).

The origin of each star particle (accreted/*in situ*) was determined in post-processing. Here, we use the labels from Font et al. (2020). In brief, the redshift of formation is recorded for each star particle during the simulation and in post-processing the simulation snapshot that this is closest to (in lookback time) was identified. If at the time of its formation the star particle was gravitationally bound to the main progenitor of the MW-mass galaxy, it was labelled as *in situ*, otherwise, as accreted. Note that by this definition, stars that are born

from gas stripped from an infalling satellite, but inside the main halo, are also labelled *in situ*. This is the case only for a small percentage of stars, however.

Since we aim to devise ML models suited for the Galactic observations, we focus our training and testing of models on a subset of galaxies from ARTEMIS which have a disc-like component similar to that of the MW.¹ This ensures that there are sufficient examples in the learning set which encapsulate the distribution of accreted and *in-situ* stars in a disc-like galaxy. For the disc-like selection criteria, we use the kinematics, specifically the corotational parameter κ_{co} (Abadi et al. 2003; see also Font et al. 2020; Dillamore et al. 2022). Here, we define it as $\kappa_{\text{co}} = \sum_{r=0}^{30\text{kpc}} L_z / \sum_{r=0}^{30\text{kpc}} L_{z,\text{circ}}$, where L_z is the total stellar angular momentum along the z -axis, and $L_{z,\text{circ}}$ is the total angular momentum of star particles with the same energy but in a corotating circular orbit. For this computation, we only consider star particles within an aperture of 30 kpc, and impose a cut-off of $\kappa_{\text{co}} \geq 0.50$ to select galaxies with the most prominent disc components. This results in a sample of 16 galaxies listed in Table 1, together with their main physical properties: the total accreted stellar fraction, the corotation parameter, the total stellar mass, half-(stellar) mass radius, maximum circular velocity, and average chemical abundances ([Fe/H] and [α /Fe]).

Galaxies are further split into two sets: a training and a test data set, respectively. The training set is used to provide examples of accreted and *in-situ* stars to the ML models, while the test data set is used to assess the classification performance. The test data set is composed of galaxies with an assembly history more similar (although not exactly the same) to the one inferred for the MW, specifically those where the most massive accreted progenitor (MMAP) was accreted more than 8 Gyr ago, and the stellar mass ratio of this MMAP to the total stellar mass of the host is ≥ 0.4 (see fig. 7 and table A1 of Dillamore et al. 2022 for details). The two data sets, comprising of 12 galaxies for training and 4 for testing, are listed separately in Table 1.

For each galaxy in the training data set, we also reserve 20 per cent of stars,² for the validation data set. This comprises data that are used during the training routine for evaluating the neural networks on unseen data and detecting overfitting, that is, the modelling of the noise contained in the training examples.

2.1 Stellar parameters as features

For training and testing of the ML models, each star particle is described as a vector of physical parameters (also known as features) which are expected to be relevant for the distinction between

¹We note that, in an initial phase, we used for training simulated galaxies irrespective of their morphological type (i.e. both ellipticals and discs). However, we found that the performance of models was significantly reduced in this case, particularly when the training set was composed of mostly different morphological types than the testing set. Since here we are mainly interested in optimizing the performance of the ML models, we choose to use a subsample composed of only disc-like systems. This restricts somewhat the accretion histories that are included in the data sets to those that are more quiescent. However, this approach is justified as: (1) the accretion history of the MW is partially known, and therefore we do not need to be completely agnostic about this feature; and (2) disc-like galaxies can form with a variety of accretion histories (see e.g. Font et al. 2017) and our disc-like subsample covers a broad range of these scenarios.

²Although our ML models are developed with the aim of applying them on individual stars in the Galaxy, the simulations can only track star particles, which are essentially individual SSP. For brevity, throughout the paper, we will refer to the star particles as ‘stars’.

Table 1. Sample of disc-like galaxies in ARTEMIS selected based on their corotation parameter, κ_{co} . These galaxies are separated into two datasets used for training and test the performance of the ML models, respectively. The columns are: (1) galaxy label; (2) fraction of accreted stellar component (defined as the mass fraction of accreted star particles over the total stellar mass (*in situ* + accreted)); (3) corotational parameter; (4) the total stellar mass; (5) half stellar mass radius (defined as the radius enclosing 50 per cent of the total stellar mass); (6) maximum circular velocity; (7) average [Fe/H] abundance; and (8) average $[\alpha/\text{Fe}]$ abundance, where α is tracked by Mg abundance. Apart for the fraction of accreted stars, all quantities are computed within 30 kpc from the centre of the MW-mass galaxy.

Galaxy	f_{acc}	k_{co}	$M_{*}[10^{10} M_{\odot}]$	$r_{1/2}$ [kpc]	$v_{\theta_{\text{MAX}}}$ [kms $^{-1}$]	$\langle[\text{Fe}/\text{H}]\rangle$	$\langle[\alpha/\text{Fe}]\rangle$
Training data set							
G01	0.14	0.60	3.64	4.86	199	−0.14	0.19
G15	0.11	0.61	3.57	5.88	170	−0.16	0.22
G17	0.10	0.69	3.74	7.26	198	−0.21	0.25
G18	0.22	0.59	2.78	4.38	184	−0.15	0.20
G19	0.04	0.67	2.57	4.92	177	−0.16	0.24
G23	0.11	0.56	2.87	2.77	197	−0.07	0.21
G24	0.11	0.55	3.63	3.90	185	−0.16	0.22
G25	0.18	0.63	2.57	5.52	172	−0.24	0.26
G27	0.21	0.57	2.57	5.40	160	−0.19	0.22
G38	0.04	0.81	2.97	8.46	176	−0.13	0.20
G40	0.16	0.64	2.02	4.50	155	−0.15	0.20
G44	0.12	0.63	4.28	5.22	204	−0.27	0.31
Test data set							
G29	0.08	0.65	2.95	2.60	210	−0.08	0.18
G30	0.28	0.55	2.12	4.20	172	−0.10	0.23
G34	0.05	0.78	2.76	6.20	183	−0.16	0.21
G42	0.13	0.65	2.10	3.10	174	−0.18	0.25

accreted and *in-situ* stars. Although the simulations provide more information on each star particle (including, e.g. the mass or the gravitational potential), we focus only on stellar parameters which can be observed, to facilitate future applications on survey data. These features are divided into four categories, resembling (very broadly) the focus of different types of observational Galactic surveys:

(i) *Positions and kinematics.* For many MW stars, positions and velocities are readily available, for example, from *Gaia* and RAVE. Accreted stars are expected to differ from *in-situ* ones both in terms of their locations and of their overall motions. Accreted stars extend much further out into the halo where they tend to appear as kinematically cold tidal streams, and their orbits tend to be more randomly distributed, whereas the *in-situ* stars are more prominent in the inner region of a galaxy (Font et al. 2011), follow a more flattened distribution, and tend to have a prograde rotation with the disc (McCarthy et al. 2012). Therefore, for positions, we choose as features the Galactocentric radius in the plane of the disc, R , and the distance perpendicular to the plane of the disc, z , while for kinematics measures we use the rotational velocity in the plane of the disc, v_{θ} , and the velocity dispersion in the plane perpendicular to the disc, σ .

(ii) *Chemical abundances.* Chemical abundances are related to the star formation history of the associated progenitors (Freeman & Bland-Hawthorn 2002). Here, we focus on the stellar metallicities, defined as [Fe/H], and on the $[\alpha/\text{Fe}]$ abundances, where α is tracked by Mg. We expect stars of accreted origin to have, on average, lower [Fe/H] and higher [Mg/Fe] values than those *in situ*.

(iii) *Stellar ages.* Stellar ages, τ , can also be used in tandem with the kinematic and chemical properties of stars to trace populations formed in the same galactic environment (Helmi 2020). [Fe/H] values are expected to correlate well with ages and they are often used as proxies for the latter. Old accreted stars are also expected to have higher $[\alpha/\text{Fe}]$ abundances due to the short, bursty star formation episodes in their parent dwarf galaxies at high redshift (Robertson et al. 2005). In addition, ages can provide complementary

information about the rate of chemical enrichment at different epochs (e.g. Hawkins et al. 2014).

(iv) *Photometry.* Tidal streams from different disrupted satellite galaxies are expected to stand out in terms of their surface brightness; specifically, the brightness of a stream tends to correlate with the stellar mass of its dwarf progenitor (Font et al. 2006b; Johnston et al. 2008; Gilbert et al. 2009; Cooper et al. 2010). Similarly, accreted stars can also be distinguished from *in-situ* stars in terms of their photometric properties. Accreted debris tends to be fainter than the *in-situ* component, due to its lower mass and larger spatial extent. For photometric properties, we choose the absolute magnitude in the *Gaia* G passband, M_G , and the colour evaluated in the G_{BP} and G_{RP} passbands, $\text{BP} - \text{RP}$. We note, however, that our results are not dependent to the specific *Gaia* passbands or survey. These photometric properties also correlate with stellar ages (which be inferred from colour–magnitude diagram fitting, e.g. Gallart, Zoccali & Aparicio 2005), thus offering an alternative to direct age measurements, which are usually more difficult to obtain.

To summarize, we choose as possible input features for ML models the following stellar parameters:

$\{R, z, v_{\theta}, \sigma, [\text{Fe}/\text{H}], [\alpha/\text{Fe}], \tau, M_G, \text{BP} - \text{RP}\}$.

As described later, in Section 3.2.2, the ML models ultimately include an optimal set of features, which provides the best performance for our benchmark model. The optimal set is the same as the set of parameters above, but excluding the stellar ages (τ), for which the benchmark model is able to retrieve the information from the other parameters, mainly from [Fe/H] and $[\alpha/\text{Fe}]$.

2.2 Galaxy-specific features

In addition to stellar features, one of our models (see Section 3.2.4) includes a set of galaxy-specific features, devised to account for the accretion histories of MW-mass hosts. These are listed in Table 1, and

comprise of: the stellar mass half-radius, the corotational parameter, the total stellar mass, the maximum circular velocity, and the average $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ abundances of the respective MW-mass galaxy.

In general, we expect galaxies that experienced more massive mergers to have less well-defined stellar discs (e.g. smaller sizes, lower κ_{co}) and also tend to be more massive. Additionally, we expect that systems with higher masses would be more chemically enriched (higher $[\text{Fe}/\text{H}]$ and lower $[\alpha/\text{Fe}]$). Therefore, some indication about the accretion histories can be inferred from the present-day properties of stellar populations in the MW-mass hosts (see also Grimozzi, Font & De Rossi 2024).

We note that, although the merger histories of the simulated galaxies are already known (e.g. in the form of merger trees, or the properties of the MMAPs), we choose not to use this type of information and focus instead on observable parameters, as we do in the case of star particle features. In Section 4.1, we discuss the performance of models with and without these galaxy-specific features.

3 MACHINE LEARNING MODELS

To devise an appropriate ML model for our task, we need to address two challenges: (i) to identify ML models which can learn effectively the underlying patterns in the data; and (ii) to determine the input parameters that optimize the performance of these models.

For the first task, we consider models representative of two main families of supervised ML methods, namely ANNs and decision trees. From the ANN type, we consider a feedforward model called a multilayer perceptron (MLP). Given the relatively small number of features that describe the accreted and *in-situ* stars, we choose to consider shallower architectures than the ones selected in Tronrud et al. (2022), so as to limit the risk of overfitting. We then augment this MLP model with domain inputs, that is, we include additional galaxy-specific features. With the extended set of features, both stellar and galaxy-specific, we aim to mitigate the potential decrease in the performance caused by the variability introduced by the specific assembly histories of galaxies. For the same purpose, we also consider the transformational machine learning (TML) technique (Olier et al. 2021). From the category of decision tree-like systems, we consider the eXtreme Gradient Boosting (XGBoost) model (Chen & Guestrin 2016).

To better visualize the data set used for training the models, as well as to understand the functioning of these models, we use the UMAP dimensionality reduction technique (McInnes, Healy & Melville 2018a). This method consists of mapping the accreted and *in-situ* stars into a new, lower-dimensional plane, while maintaining the global and local structures; thus, stars are clustered in structures which make relations hidden in physical space more visually evident.

ANNs have been investigated recently for this task (Ostdiek et al. 2020; Tronrud et al. 2022; Trujillo-Gomez et al. 2023), and are commonly used in astrophysics, for example, for the classification of transients and variable stars (Jayasinghe et al. 2019; Agarwal et al. 2020; Chen et al. 2022), of quasars (Nakoneczny et al. 2019; Clarke et al. 2020; Nakoneczny et al. 2021), or of galaxies (Traven et al. 2017; Domínguez Sánchez et al. 2018; Huertas-Company & Lanusse 2023). XGBoost is also commonly used, with many applications in Galactic studies (e.g. Anders et al. 2023). To our knowledge, models like TML or UMAP are investigated for the first time here for an astrophysical problem.

For the task of determining the optimal set of features, we start with a wide range of physical stellar parameters (described in Section 2.1) and determine which combination provides the best performance

(Section 3.2.2) for our benchmark model (Section 3.2.1). The implicit assumption in our approach is that the optimal set of features would be the same for any type of ML model adopted. We then compare the classification performances of different models using the same (fixed) set of features.

In the following, we describe the metrics used for evaluating the classification performance of models (Section 3.1). The model implementations are described separately in Section 3.2, where we also provide some technical background on each of them. Otherwise specified, we will make use of common ML terminology.

The training and implementation of the ANNs is performed using the TENSORFLOW (Abadi et al. 2015) library. The XGBoost and UMAP methods are developed using the XGBOOST (Chen & Guestrin 2016) and UMAP (McInnes et al. 2018b) PYTHON packages, respectively.

3.1 Performance metrics

In the development of all models we adopt a supervised learning approach. In supervised models, the mapping between features and prediction is learned by providing a set of example-label pairs. Furthermore, the model parameters are tuned to minimize the difference between the prediction and the actual class (the label) as quantified by an objective function. The separation of the accreted versus *in-situ* stars can be thought as a binary classification problem, where the positive class is represented by the accreted stars and the negative class by the *in situ*. An ML model achieves this by applying a sequence of mathematical operations and tunable parameters to map the properties of stars to a value of either 0 or 1, representing the negative and positive classes, respectively.

For the ANN and decision tree models, the prediction for a given star is represented by the output, which is a value between 0 and 1 and measures the probability of the star belonging to the positive (accreted) class. In both cases, a star is classified as accreted if its prediction is greater than a threshold value. The optimal performance of a classifier may occur at a different threshold value for different models, especially in problems with highly imbalanced data sets, such as ours where accreted stars comprise, on average, 10 per cent of the overall stellar content of a galaxy. We therefore explore also the effect of changing threshold values on the performance of our models (see Section 4.2).

Here, we use the usual performance metrics, namely the precision (P) and recall (R) of a model. These are used to compare the performances of various models, but also to identify the optimal set of stellar features. P represents also the purity of the sample of accreted stars, while R characterizes the completeness of the sample. By definition, these two parameters correspond to the number of correctly and mis-predicted accreted stars, respectively:

$$P = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad \text{and} \quad R = \frac{\text{TP}}{\text{TP} + \text{FN}},$$

where TP, FP, and FN are the number of true positives (i.e. stars that are correctly classified as accreted), false positives (*in-situ* stars which are misclassified as accreted), and false negatives (accreted stars which are misclassified as *in situ*), respectively.

Because the precision and recall values of both ANNs and decision tree-based models vary based on the different classification thresholds, with some thresholds favouring some models over the others, we also use as metric the area under the precision–recall curve, PR-AUC. This metric is more robust across models as it accounts for the precision and recall values evaluated on a range of thresholds common to all models. A random classifier would return

a PR-AUC score equal to the fraction of accreted stars in the test data set, whereas a perfect classification algorithm would have a PR-AUC score of 1.

We have elected not to employ the accuracy metric, defined as the overall fraction of correctly classified stars, due to its shortcomings for highly imbalanced cases. For example, in galaxies with accreted fractions of ≈ 10 per cent (a typical value), a classifier which always predicts stars as being *in situ* would have an accuracy of ≈ 90 per cent even though it failed to identify any accreted stars.

Aside from the common ML metrics, we also compare the models in terms of how well they fare in terms of astrophysical diagnostics. These are not metrics per se, however, they are useful to help understand whether the models are able to grasp the ‘physics’ behind the data. We expect that a model that is able to learn (or mimic) the physical processes behind the origins of the two populations would be less precise exactly in those regions of the parameter space where the properties of the two populations are similar (e.g. stars that were born in the early phases of the galaxy formation are old, metal-poor, more α -enhanced, and move on less ordered orbits, regardless of whether they were born *in situ* or accreted). A model that cannot learn the physical patterns may still have a good performance, however its mis-classifications may be distributed more randomly in physical space. The three diagnostics used here, are:

- (i) $[\alpha/\text{Fe}] - [\text{Fe}/\text{H}]$ plane. Accreted stars tend to be located in the high $[\alpha/\text{Fe}]$, low $[\text{Fe}/\text{H}]$ region of the plane, while disc and *in-situ* halo stars generally have lower $[\alpha/\text{Fe}]$ and higher $[\text{Fe}/\text{H}]$ values.
- (ii) Toomre diagram, which is the distribution of rotational velocity, v_θ , versus the dispersion velocity σ . In this plane, the disc and the *in-situ* halo stars have high v_θ and low σ , whereas accreted stars do not have a preferred direction of motion, and generally have high σ .
- (iii) $E - L_z$ distribution. As these parameters are quasi-conserved for a given infalling satellite, stars belonging to different disrupted progenitor would appear as ‘clumps’ in this plane. These clumps are likely to be more distinct in the upper part of the plane, which is associated with late accretions. The *in-situ* halo stars, and the disc, are located on the region with positive L_z .

3.2 Supervised ML models

3.2.1 The benchmark model

For our analysis, we build a benchmark model to: (i) investigate the most informative set of stellar properties that can distinguish accreted from *in-situ* stars, that is, the *optimal* set of features; and (ii) to compare the performance of different ML algorithms trained on this optimal set.

Our benchmark model is similar to the GANN model of Tronrud et al. (2022). This is an MLP (see Section 3.2.3) comprised of an input layer, a batch normalization layer, four hidden layers of 64, 256, 64, and 32 neurons, and an output layer with one neuron, resulting in a total of 35 521 trainable parameters. We also use the same activation functions as in the GANN model in the corresponding layers. In the following, we refer to this configuration of layers and neurons as the benchmark architecture.

The model is first trained on a set of stellar features comprising ages, $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$, as in the GANN model. Following the same approach as in GANN, the chemical abundances are expressed linearly (rather than the more conventional logarithmic form) and normalized by the correspondent solar values. However, unlike in GANN, we choose not to include the hydrogen fraction with

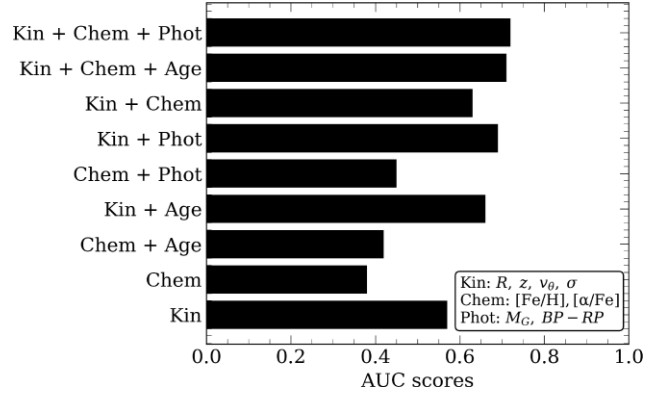


Figure 1. PR-AUC scores for the precision–recall curves for the models obtained training the benchmark architecture using different combinations of features. The feature categories are: positions and kinematics (labelled ‘kin’), $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ abundances (‘chem’), ages, and *Gaia* magnitudes and colours (‘phot’).

$[\text{Fe}/\text{H}]$, given that the information about the former can be implicitly reconstructed by the network from the latter. As for GANN, our benchmark model is trained on an equal number of accreted and *in-situ* stars. There are other differences from the GANN model also. One is that we train the model only on stars within $r \leq 50$ kpc from the centres of galaxies, as beyond this radius the distribution of stars is overwhelmingly of accreted origin (this is also valid for the other models). More significantly, we only include example stars from the main halo, whereas Tronrud et al. (2022) include also those from present-day satellites, to augment the samples of accreted stars. Moreover, they adopt a strategy of drawing equal number of stars from satellites in different mass ranges, in order to increase the number of examples of stars from objects which contribute fewer stars (i.e. the low-mass dwarfs). In choosing to train only on existing debris, our benchmark model has a slightly lower performance than that of GANN, however the training data represents a closer match to what one would expect from observations.

3.2.2 The optimal set of stellar features

With the benchmark architecture, we proceed to compare the classification performance using different combinations of stellar features as inputs, to determine the optimal set for the models. We consider the stellar properties described in Section 2.1, which we divide into four categories: positions and kinematics (R , z , v_θ , σ), chemical abundances ($[\text{Fe}/\text{H}]$, $[\alpha/\text{Fe}]$), ages (τ), and photometry (M_G , $BP - RP$). For brevity, these are also referred to as ‘kin’, ‘chem’, ‘age’, and ‘phot’, respectively.

We then train the benchmark architecture with different combinations of these four categories and evaluate the classification performance on stars from the validation data set. The results are shown in Fig. 1. On their own, positions and kinematics (kin) give a PR-AUC score of 0.57, which is significantly higher than the score for $[\text{Fe}/\text{H}]$ and $[\alpha/\text{Fe}]$ (chem), PR-AUC 0.38. This suggests that spatial distribution and kinematics are more informative than chemical abundances. In combination, kin + chem perform somewhat better than kin alone, with a PR-AUC score of 0.63, which suggests that the two categories provide complementary information. Adding ages to positions and kinematics improves the performance of the model compared to just adding chemistry (PR-AUC scores 0.66 for kin + age versus 0.63 for kin + chem). Similarly, adding ages to positions, kinematics, and chemistry improves the score

compared to not adding ages (PR-AUC of 0.71 for kin + chem + age compared to 0.63 for kin + age). Interestingly, photometry adds slightly more information than the ages, to both positions and kinematics (PR-AUC scores 0.69 for kin + phot versus 0.66 for kin + age), or chemistry categories (PR-AUC scores 0.45 for chem + phot versus 0.42 for chem + age). A possible explanation for this is that photometry is implicitly related to both the ages and metallicities of stars.

Overall, the best performance is provided by kin + chem + phot (PR-AUC score of 0.72), therefore excluding ages. We therefore assign the corresponding input parameters, namely:

$$\{R, z, v_{\theta}, \sigma, [\text{Fe}/\text{H}], [\alpha/\text{Fe}], M_G, \text{BP} - \text{RP}\},$$

as the *optimal* set features for further training and comparison of the other ML models.

We note, however, that the combinations kin + phot + age and kin + phot have classification performances which are very close to that of the optimal combination. This suggests that ML models could also be trained on reduced information, for example on just positions, kinematics, and photometry, without a significant drop in performance. In fact, this may be a preferred option for observations, given that photometry is usually more readily available than spectroscopy.

3.2.3 Multilayer perceptron

An MLP is a type of ANN consisting of multiple layers of interconnected artificial neurons, or perceptrons. The architecture typically comprises an input layer, one or more hidden layers, and an output layer. Each neuron receives input signals from neurons in the previous layer and computes a weighted sum based on internal, tunable parameters describing the importance of the single inputs. Before being forwarded to the next layer, the result is passed to a non-linear activation function to allow the learning of non-linear relations between inputs and outputs. During the training process, the weights connecting the neurons are updated to minimize the error, as estimated by an objective function, between the predicted and actual class through an optimization algorithm.

The benchmark model described earlier is also an MLP. However, because the classification between accreted and *in-situ* stars is inferred from a small number of features, we consider also shallower architectures than the one used in the benchmark model. Specifically, we explore architectures comprised of one hidden layer of 10, 50, and 100 neurons, and two hidden layers with 50 neurons in each. In total, these four MLPs have 101, 501, 1001, and 3051 trainable parameters, respectively. In comparison, the benchmark architecture contains 35 777 parameters.

We train the MLPs with various architectures on the optimal set of input features (described in Section 3.2.2), as was done for the benchmark model. We adopt the same activation functions and optimization algorithm for all cases. For the outputs of the neurons in the hidden layers, we apply a Scaled Exponential Linear Unit function (Klambauer et al. 2017), while for the neuron in the output layer we apply a sigmoid function, to ensure the prediction is in the range 0–1. The trainable parameters are updated using the Adam (Kingma & Ba 2014) optimization algorithm on the error between predictions and labels estimated by the binary cross-entropy function. The training on the optimal set of features of the four MLP plus the




Figure 2. Comparison between the precision–recall curves obtained training the benchmark architecture and shallower ANNs on the optimal set of features (see Section 3.2.2).

benchmark architectures is done for a maximum of 100 epochs,³ with an adaptive learning rate⁴ halving when the value of the objective function stops decreasing for more than five epochs.

Fig. 2 shows a comparison of the classification performances of all these architectures, based on the PR-AUC score evaluated on the test data set. All MLPs have similar performances, with PR-AUC scores ranging from 0.578 (1 hidden layer and 10 neurons model) to 0.591 (1 hidden layer and 100 neurons model). Despite the significantly larger number of trainable parameters, the benchmark architecture (0.584) outperforms only the shallowest model. The model with 1 hidden layer and 50 neurons, and the one with 2 hidden layers of 50 neurons each, have PR-AUC score of 0.589 and 0.590, respectively. As the architecture comprising 1 hidden layer with 100 neurons returns the highest PR-AUC score (0.591), we considered it for the rest of the analysis (hereafter denoted as MLP).

3.2.4 Multilayer perceptron with galaxy features

The study of Tronrud et al. (2022) suggests that MLP models may be biased towards specific assembly histories. This result is expected, especially when the training set does not contain sufficient types of accretion histories. We find a similar result when we train an MLP model on a single galaxy. As expected, the model performs better when the stars belong to the same galaxy. Table 2 shows the classification performances of these MLP models represented by the PR-AUC scores. The models are trained on the optimal set of features using accreted and *in-situ* examples from one galaxy (listed in the leftmost column) and tested on another galaxy (listed in the top row). Where the galaxy pairs are the same, we only use the stars in the validation data set. On a galaxy basis (i.e. column-by-column analysis in the table), the best classification performance is always associated to the model explicitly trained on the galaxy it is tested on.

This result is a consequence of the unique assembly history of each galaxy, where the properties of the progenitors, for example, their

³An epoch is a complete pass of the MLP through all the examples in the training data set.

⁴The learning rate is the constant of proportionality relating the gradient of the objective function and the associated change in the trainable parameters of the MLP.

Table 2. PR-AUC scores for the MLP models trained on accreted and *in-situ* examples from the galaxy in the leftmost column, and tested on the galaxies listed in the top row. Where the galaxy label is the same, stars in the validation data set were considered.

	G01	G15	G17	G18	G19	G23	G24	G25	G27	G38	G40	G44
G01	0.68	0.60	0.52	0.66	0.44	0.60	0.59	0.45	0.81	0.48	0.76	0.66
G15	0.54	0.77	0.48	0.67	0.58	0.62	0.67	0.44	0.74	0.62	0.66	0.66
G17	0.58	0.64	0.62	0.62	0.62	0.70	0.61	0.47	0.73	0.59	0.67	0.71
G18	0.55	0.66	0.45	0.76	0.46	0.50	0.65	0.46	0.78	0.48	0.74	0.57
G19	0.54	0.67	0.47	0.64	0.74	0.71	0.66	0.44	0.73	0.62	0.66	0.65
G23	0.57	0.67	0.52	0.62	0.69	0.82	0.61	0.44	0.76	0.62	0.63	0.72
G24	0.52	0.66	0.43	0.67	0.48	0.56	0.78	0.45	0.79	0.53	0.75	0.61
G25	0.47	0.43	0.44	0.58	0.44	0.46	0.39	0.61	0.62	0.32	0.46	0.58
G27	0.48	0.49	0.35	0.63	0.25	0.36	0.47	0.42	0.88	0.53	0.76	0.50
G38	0.51	0.62	0.47	0.63	0.45	0.50	0.52	0.45	0.77	0.76	0.70	0.65
G40	0.47	0.43	0.33	0.63	0.16	0.29	0.43	0.39	0.78	0.29	0.83	0.46
G44	0.57	0.64	0.55	0.63	0.64	0.65	0.63	0.46	0.78	0.58	0.73	0.83

infall times, orbits, and masses, imprint a specific characteristic in the stellar properties (features) of accreted stars. For instance, in galaxies where the MMAP was accreted at early times the accreted stars are mostly rich in α -elements and are more phase-mixed; conversely, in galaxies with a late MMAP, the accreted stars tend to be more α -poor and more spatially coherent (Font et al. 2006a). Consequently, the MLP model learns the details of the unique imprint of the assembly history of a given galaxy leaves in its distribution of accreted stars.

We note, however, that some models are able to identify a purer and more complete sample of accreted stars in other galaxies than the one they had been trained on. Again, this can be directly related to a difference in assembly histories as some galaxies (e.g. G27) have a neater distinction between the accreted and *in-situ* populations.

To account for the variability in the properties of accreted stars in different systems, we could choose to use a conglomeration of galaxies in the training set, as in the data set described in Section 2. However, this model would still underperform if applied onto a galaxy whose assembly history is not represented in the training set. In ML terminology, the performance degradation of a model trained on a source distribution ('domain') when applied to a statistically different target one is referred to as 'domain shift' (Quiñonero-Candela et al. 2009). Since our aim is to create a model capable of generalizing across different assembly histories, we retrain the MLP model by providing additional information about the galaxy from which the example stars are taken from. These additional input features are global properties of host MW-mass galaxies, measured within an aperture of 30 kpc. This approach has the advantage of providing information about the accretion history of the host galaxy, without given the model any *a priori* knowledge of which stars were accreted. The galaxy-specific input features are described in Section 2.2 and are listed in Table 1.

3.2.5 Transformational machine learning

As an alternative method to mitigate the domain shift problem, we consider the TML technique of Olier et al. (2021). In this framework, each data point is described by a vector of predictions obtained from an ensemble of *base* models. While using a common set of features, the base models are sequentially trained on different examples to perform different tasks. This representation of the data is used as an input to a new model, which combines the prior information encoded in the base models.

Before the implementation of the TML approach, we trained an MLP model (see Section 3.2.3) on each galaxy in the training data set. Because of the differences in the assembly histories of these galaxies, learning to classify the accreted stars is considered by the model to be a specific task for each galaxy. The resulting ensemble of MLPs is then used to derive a predictive description of all the stars in the training data set. These 12D vectors are then passed as inputs to a single-layer ANN with 100 neurons. During training, the internal parameters of the base MLP models are held constant, while the parameters from the neural network which combines the predictions are allowed to update. Also, to avoid any data leakage, the predictions of a base model on the stars used for the training are set to 0.

3.2.6 XGBoost

A decision tree is an ML algorithm based on a tree-like structure. It begins with a root node representing the entire data set and recursively splits the data into smaller subsets (branches) based on feature values. The endpoint of a branch is called a *leaf* and contains the model prediction. In this work, we combine many decision tree models trained on the same data set using the gradient boosting machine (GBM) method (Friedman 2001). Following this algorithm, a decision tree is created to separate accreted and *in-situ* stars by predicting their classification label as a continuous score between 0 and 1. Then, a new decision tree is added to predict the error between the predicted and actual labels (here, as for the MLP model, estimated by the binary-loss function). The prediction from the new model is then added to the initial predictions to make a more accurate classification. This continues for an arbitrary number of iterations, with each new model sequentially added and trained to minimize the error of the whole ensemble.

Here, we implement a GBM with decision tree models using the XGBoost (Chen & Guestrin 2016) method. This algorithm is particularly suitable for large data sets as the ensemble of models is built in parallel rather than serially. Moreover, it includes L1 and L2 regularization techniques to control overfitting. The number of base models in the ensemble and the number of splits in each tree are decided using the OPTUNA⁵ hyperparameter optimization framework.

⁵<https://optuna.org/>

Table 3. Fiducial classification threshold values for the models. Each value is associated to the highest F1-score calculated based on the precision and recall values on the validation data set.

Model	Threshold	F1-score
MLP	0.33	0.67
MLP + galaxy features	0.31	0.63
TML	0.24	0.62
XGBoost	0.33	0.68

3.2.7 UMAP

The UMAP method (McInnes et al. 2018a) is a dimensionality reduction technique, such as principal component analysis, with the advantage that the obtained dimensions can be non-linearly related to the starting ones. We apply this method in order to find relations, or identify possible structures within the accreted and *in-situ* examples that may exist in the training and test data sets. Given the specificity of this method, we use it only for visualizing the data, and therefore we do not include it in our classification performance comparison.

Assuming the data are uniformly distributed on a locally connected Riemannian manifold, the algorithm constructs a fuzzy topological structure of it in a 8D parameter space and maps it into a lower dimensional space with the closest equivalent structure. UMAP can be used in a supervised way by providing the labels of the classes into which the data are separated. This ensures that both the global and class-specific structures of the data are retained while maximizing their separation in the new embedding.

Here, we train a UMAP model to reduce the 8D parameter space defined by the optimal set of stellar features into a 2D plane, where the separation between the accreted and *in-situ* training examples is maximized. Because of the high-computational resources required by this method, we consider only a subset of examples from each galaxy, which consists of all the accreted examples and an equal amount of the *in-situ* ones. This results in a statistically significant number of examples from each galaxy. To investigate potential differences between the structures of the training and test data sets, we use the same UMAP model to project all the test examples into the 2D plane. A discussion of the results of this UMAP embedding is provided in Section 4.4.

4 MODEL COMPARISON

4.1 Classification performance

With the optimal set input features, we proceed to compare the classification performance of our models, MLP, MLP with galaxy features, TML, and XGBoost, in separating the two classes of stars. We evaluate all models on the test data set, and compare the classification performances using the PR-AUC scores. We also assess the purity and completeness of the retrieved accreted samples using the precision and recall metrics at the optimal threshold value. To ensure a fair comparison, for each model we consider the threshold associated to the highest value of the harmonic mean between precision and recall (i.e. the F1-score). The resulting fiducial thresholds are shown in Table 3.

Fig. 3 shows the comparison between these four models, plus the benchmark model. For the latter, we use the reduced set of features described in Section 3.2.1. This figure shows that all four models perform significantly better than the benchmark, which indicates the importance of adding more input parameters, in this case from the




Figure 3. Precision and recall values at different classification thresholds for the models. The metrics were evaluated considering all the stars in the test data set.

kinematical and photometric properties of the stars. The MLP, TML, and XGBoost models perform similarly on the test data set, with PR-AUC scores of 0.59, 0.57, and 0.59 respectively. The TML model does not provide any improvement over the MLP, which implies that there is not enough variance among the single MLPs in the ensemble.

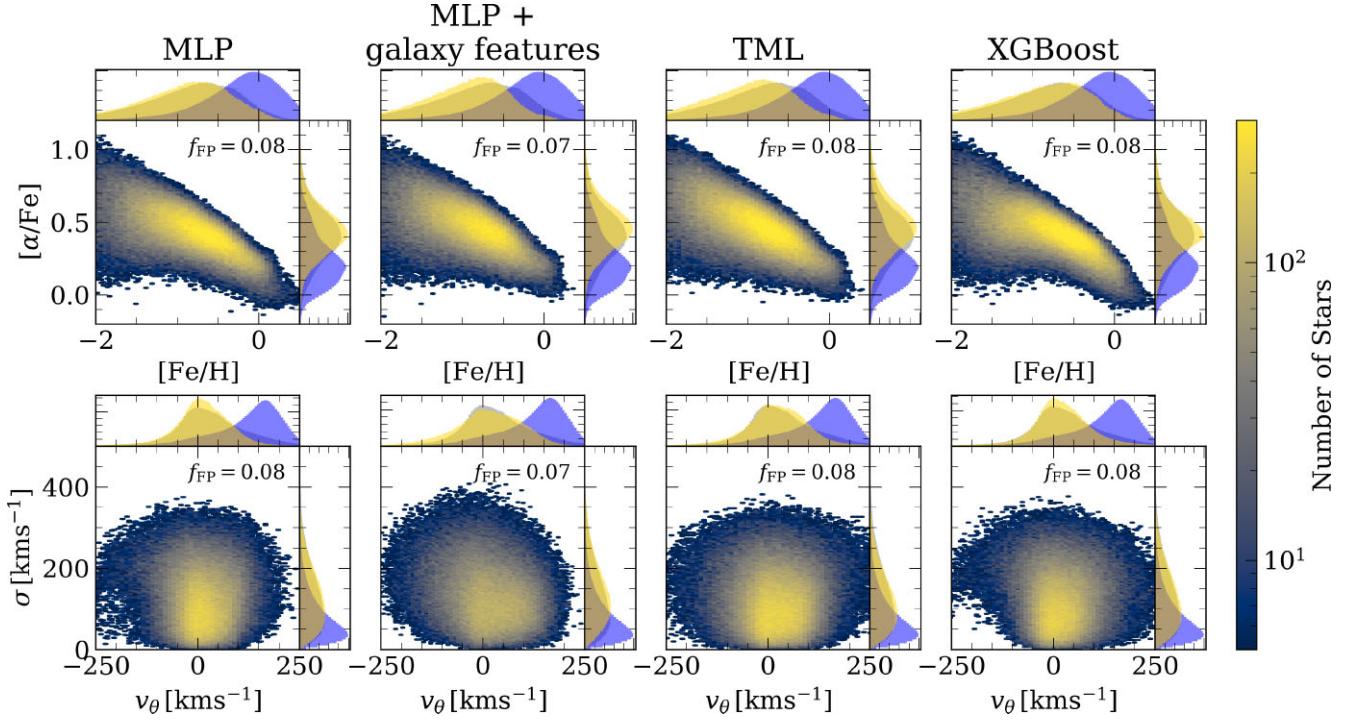
The MLP model with added galaxy-specific features gives a PR-AUC score of 0.57, which is worse than the one of the MLP model and the one calculated on validation data (0.64). This suggests that the galaxy-specific features favour the learning of specific patterns in the data, leading to overfitting, rather than learning the distinction between the accreted and *in-situ* classes in those features. This could be due to redundancies in the set of galaxy-specific features (which, unlike the stellar features, have not been optimized). For example, information gained from galaxy stellar masses may be very similar to that inferred from $\langle[\text{Fe}/\text{H}]\rangle$, as MW-mass systems follow the stellar mass–metallicity scaling relation. Also, the information gained from the overall rotation of stars (i.e. κ_{co}) may overlap with that obtained from the sizes of galaxies. In the future, we aim to investigate whether an optimal set of galaxy-specific features exists, particularly one that will increase the performance above the model without any such features. Alternatively, it could be that the galaxy properties used for this task are not representative of the specific accretion histories. In this case, one may opt, instead, to use parameters more directly related to the merger histories, for example the properties of the MMAPs. This will be investigated in a future study.

Table 4 shows a galaxy-by-galaxy comparison of all five models based on the P and R values evaluated at the fiducial classification thresholds. Confirming what was found previously, MLP, TML, and XGBoost have similar performances for every test galaxy, consistently retrieving > 50 per cent of accreted stars with similar purity. In general, the TML model retrieves a larger sample of accreted stars than the single MLP, with only a minor decrease in precision.

Overall, despite belonging to different families of ML methods, the MLP, TML, and XGBoost models show similar performances, suggesting that the same underlying relations between stellar properties and their origin are learned from the data. This is an encouraging result for further applications.

Table 4. Comparison of the precision and recall values evaluated on the test data set at the fiducial thresholds for the ML models.

Model	G29		G30		G34		G42	
	<i>P</i>	<i>R</i>	<i>P</i>	<i>R</i>	<i>P</i>	<i>R</i>	<i>P</i>	<i>R</i>
Benchmark	0.15	0.98	0.45	0.91	0.09	1.00	0.21	0.97
MLP	0.47	0.64	0.59	0.57	0.35	0.65	0.54	0.54
MLP + galaxy features	0.54	0.59	0.56	0.56	0.41	0.55	0.50	0.63
TML	0.46	0.73	0.58	0.55	0.32	0.72	0.54	0.60
XGBoost	0.49	0.70	0.58	0.58	0.36	0.68	0.54	0.58

**Figure 4.** Distribution of the FPs in the test set in the $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ plane (top row) and in the Toomre diagram (bottom row). Columns from left to right correspond to the MLP, MLP + galaxy features, TML, and XGBoost models, respectively. For each panel, the top and side subpanels show the probability density function of the FP distributions (yellow) and of the accreted (grey) and *in-situ* (blue) training examples. For each model, we also show the FP fraction (f_{FP}) of the total number of stars in the test data set. The metrics are evaluated at the fiducial threshold values listed in Table 3.

4.2 Model comparison on physical diagnostics

So far, we have shown that all ML models trained on the optimal features perform similarly well, which is remarkable given that accreted stars comprise only a small fraction of the total stellar budget. However, despite the similar overall performance, the models may still perform differently in certain regions of the physical parameter space covered by the accreted/*in-situ* stars. If the models are truly able to extract the physical properties of the two stellar populations, we expect that they will perform better in regions where the two populations are clearly distinct in physical parameters (e.g. kinematics, or metallicity) and less well in regions where these properties overlap. To investigate this possibility, we map the distribution in a chemodynamical parameter space of false positives (*in-situ* stars misclassified as accreted) and of false negatives (accreted stars misclassified as *in-situ*) in different models. In Section 3.1, we introduced several physically motivated chemodynamical diagnostics, such as the Toomre diagram, or the $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ and E – L_z planes, which we use here to evaluate how well can the models identify accreted stars in these parameter spaces.

Fig. 4 shows the distribution of the FPs in the entire test galaxy set in the $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ plane (top row) and in the Toomre diagram (bottom row), for four models: MLP, MLP + galaxy features, TML, and XGBoost, respectively. Each panel has a corresponding set of top and side subpanels, in which we compare the probability density functions of the FPs (shown in yellow) with that of the accreted stars (grey), and of the *in-situ* stars (blue) in the training set. This figure shows that the FPs tend to be located in regions of the parameter space with low rotational velocities ($v_\theta \simeq 0$) and lower metallicities ($[\text{Fe}/\text{H}] < -1$), which are regions dominated by accreted stars. This indicates that the models can learn to identify accreted stars as typically more metal-poor and characterized by more chaotic motion, which are properties expected for this category. However, models find it more difficult to identify *in-situ* stars in this chemodynamical region. This behaviour is seen across all four models, which indicates that they all learn similar physical patterns for accreted stars. This result suggests that the similar classification performances of the models, found earlier, are in fact physically motivated.

Similarly, in Fig. 5, we investigate the distribution of FNs in the same chemodynamical parameter space. This figure shows that, although the MLP, TML, and XGBoost models retrieve the majority of the accreted stars in the test galaxies (see Table 4), some accreted stars are still missed, despite them having relatively distinct motions and chemical abundance distributions from those of the *in-situ* stars (compare, again, the yellow, grey, and blue probability distribution functions). All models present a similar behaviour in this respect, as was the case for FPs. We note here, too, that the FNs represent only a small fraction of the total number of stars, with $f_{\text{FN}} \approx 5$ per cent – 7 per cent across different models.

The mis-classification in the case of FNs is likely due to the *in-situ* stars greatly outnumbering the accreted stars in the examples available to these models. To elucidate this, we also investigate the spatial location of the FNs. For the MLP, TML, and XGBoost models, we find that the majority of mis-classified accreted stars (76 per cent, 77 per cent, and 73 per cent, respectively) lie within a galactocentric radius of 5 kpc. This suggests that the models tend to identify more accurately the stars originating from late accretions, which are generally located in the outer regions of galaxies, while struggling to retrieve the stars that originate from early accretion events and which are now fully phase-mixed in the inner region. A possible solution to improving the classification of models can be provided by data augmentation techniques, which can be used to generate a higher number of accreted stars in the training sets in the inner region. Outside the 5 kpc range, the MLP, TML, and XGBoost models identify 93 per cent, 94 per cent, and 91 per cent of the accreted stars in the test data set, with a precision of 0.57, 0.56, and 0.59 at the fiducial classification threshold, respectively.

Furthermore, since the XGBoost model has a built-in degree of explainability, it can be used to determine more quantitatively the contribution of each input feature during the classification. Specifically, at each decision node, it is possible to calculate the information gain from a given feature by subtracting the impurity (i.e. a measure of the entropy in the tree) before and after the splitting. The information gain of a specific feature in the XGBoost model can be estimated by averaging over all trees in the ensemble. Fig. 6 shows that the rotational velocity (v_θ) and the distance from the centre of the galaxy (R) are the most important parameters for distinguishing between the accreted and *in-situ* stars in this model. Surprisingly, $[\alpha/\text{Fe}]$ has a significantly lower information gain. Given the high importance of the $[\text{Fe}/\text{H}]$ parameter, it is possible that the model considers the information provided by $[\alpha/\text{Fe}]$ as redundant. Therefore, although the abundance of α -elements can be used to characterize individual accreted substructures, the information from $[\text{Fe}/\text{H}]$ may be sufficient to remove the *in-situ* background.

We note that the model performances reported here depend on the chosen classification threshold. For example, for the MLP model, the average purity of the accreted stars sample at the fiducial threshold is $\simeq 50$ per cent (see Table 4). Purer samples can be obtained by increasing the classification threshold, however, this is done at the cost of completeness. This is shown in Fig. 7, where we plot the predictions of this model in the E - L_z plane, for different threshold values (0.25, 0.5, and 0.75, respectively). Here, E denotes the total energy of a star, composed of the sum of kinetic and potential energy, while colours indicate the stellar mass fraction of accreted stars, as predicted by the model. The results are shown individually, for the four galaxies in the test data set. The regions of the accreted components which are identified with more confidence by the models are those in the upper parts of the energy spectrum, which represent mostly stars originating from late accretions. These regions are

dominated by clumpy structures which correspond to tidal debris not yet fully mixed with the rest of the halo.

This result is encouraging and indicates that ML models are sensitive to physical patterns in the data, and may be used in the future to not only identify the bulk of accreted stars in the halo, but also to find individual tidal streams. Adjusting the classification threshold could provide an advantage compared with traditional methods of selecting halo stars which are fixed (see the next section), whereas ML models can be customized to provide the most appropriate samples for different types of analysis. For instance, the identification of accreted substructures in integrals-of-motion space requires a sample of high purity to avoid the identification of spurious clusters, whereas a characterization of substructures based on a large number of chemical abundances may be conducted at a lower purity, as the contamination of *in-situ* stars can be more easily identified. In this case, the model would represent a pre-processing step for reducing the number of stars to analyse.

4.3 The ML performance in separating components versus observational cuts

Observational studies of accreted substructures in the MW often focus on regions dominated by halo stars, for example away from the disc. For instance, halo stars are often identified by imposing cuts in the stellar rotational velocities, vertical distances above or below the disc plane, metallicities, or a combination thereof. These cuts are physically motivated, however they may exclude also accreted structures embedded in or near the disc. In the previous sections, we have shown that ML methods are able to identify accreted stars even in the regions which are dominated by *in-situ* stars (the fractions of FPs and FNs are low even in the disc). This suggests that it may be feasible to apply them directly on the observational data, without making recourse to specific selection cuts. To investigate this possibility, we choose a few representative examples of selection cuts from the literature, and compare the composition of accreted versus *in-situ* stellar distributions in these cases with the corresponding distributions predicted by the ML models.⁶

Specifically, we consider three examples of selection criteria for halo stars in the Galactic Solar neighbourhood: (i) a cut in velocity, $|V - V_{\text{LSR}}| > 210 \text{ km s}^{-1}$ (Helmi et al. 2018; Koppelman, Helmi & Veljanoski 2018; Lövdal et al. 2022), where the velocity of the local standard of rest is $V_{\text{LSR}} = 232 \text{ km s}^{-1}$ (McMillan 2017); (ii) a selection in the $[\text{Fe}/\text{H}]$ - v_θ plane, used by Myeong et al. (2018) to remove the disc stars in order to find accreted substructures in the halo, namely imposing $[\text{Fe}/\text{H}] < -0.5$ and $v_\theta < 150 \text{ km s}^{-1}$; and (iii) a kinematic selection used by Massari, Koppelman & Helmi (2019) to construct a sample of accreted globular clusters, by imposing $\varepsilon > 0.5$, where $\varepsilon = L_z/L_{z, \text{circ}}$ (the latter was used for larger regions of the Galaxy, however we adapt it here for the Solar neighbourhood).

For this comparison, we focus on the Solar neighbourhood regions in the simulations, which are defined as ring tori with a minor radii of 2.5 kpc. The major radii of the tori are determined on a galaxy-by-galaxy basis, by multiplying the Solar radius in the MW (assumed here to be 8 kpc) with the ratios between the disc scale lengths of the

⁶We note that, in the analysis of observations, these types selection cuts are just a first step, and further algorithms are applied to the samples to identify accreted substructure/tidal streams, for example, HDBSCAN, or k -means. As already mentioned, the motivation of this exercise is to investigate whether ML methods could discover more accreted stars in the regions that are typically overlooked by observational methods.




Figure 5. Same as in Fig. 4, but for the accreted stars mis-classified as *in situ*, that is, the false negatives.

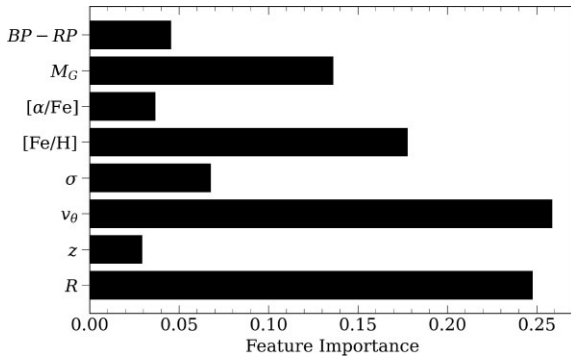


Figure 6. Information gain values describing the importance of the input features used by the XGBoost model to distinguish between accreted and *in-situ* stars.

simulated galaxies and the scale length of the thin disc of the MW, taken as 2.6 kpc (Bland-Hawthorn & Gerhard 2016). This accounts for the differences in size between the disc of the MW and the discs of the simulated galaxies.

We apply the three selection criteria above on the four galaxies from the test data set (G29, G30, G34, and G42) and label stars as accreted or *in situ* according to these cuts. In doing this, we assume that halo stars in the Solar neighbourhoods are the equivalent of stars of accreted origin, and disc stars are equivalent to those formed *in situ*. Note that, while these selection criteria are designed to exclude most of the stars of *in-situ* origin, the real compositions are a mixture of accreted and *in situ*, both in and outside the cut-out regions. The fractions of accreted stars in these two regions vary from galaxy-to-galaxy, and also on the type of selection cut that is applied. The fractions of accreted stars in the disc-like regions are very small, although the accreted stars in these regions are likely to be representative of early merger events (i.e. mostly old and metal-

poor stars). The labels assigned by these cuts are then compared with the true labels obtained from simulations, and we compute the equivalents of FPs and FNs.

We then apply the ML models on the simulated Solar neighbourhoods in the test galaxies (this time, without any selection cuts) to predict the accreted stars in these regions. As before, we use the models trained on the optimal set of features. For the sake of conciseness, we only present here the results for the MLP, but note that the XGBoost and TML models have similar classification performances.

In Fig. 8, we compare the distributions in the $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ plane of accreted and *in-situ* stars in the Solar neighbourhood regions predicted by the three selection criteria and by the MLP model. A similar comparison is shown in Fig. 9 for the distributions in the Toomre diagram. Both figures illustrate the difference in complexity between the two approaches, with the MLP model being able to provide a closer description of the true distribution of accreted stars in the chemodynamical space for all the test galaxies. This result is expected, considering that traditional observational methods are based on the assumption of a simple disc–halo dichotomy. For example, the selection cut of Myeong et al. (2018), which assumes a distinct dichotomy in terms of $[\text{Fe}/\text{H}]$, overpredicts the number of accreted stars the $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ space compared with the other methods that do not employ a $[\text{Fe}/\text{H}]$ cut (see Fig. 8). Overall, the MLP retrieves the accreted stars most accurately, compared with all the selection cut methods. This is the case not only in the overall distribution of accreted stars in the chemical abundance space, but also in terms of the number of accreted stars (compare, e.g. the accreted fractions, f_{acc} , in the corresponding panels for each galaxy in Fig. 8).

Likewise, since all selection cuts employ some type of v_θ threshold, they all underperform compared with the MLP model (see Fig. 9). The criterion of Massari et al. (2019), which allows for the inclusion of counter-rotating stars in the disc, gives a closer

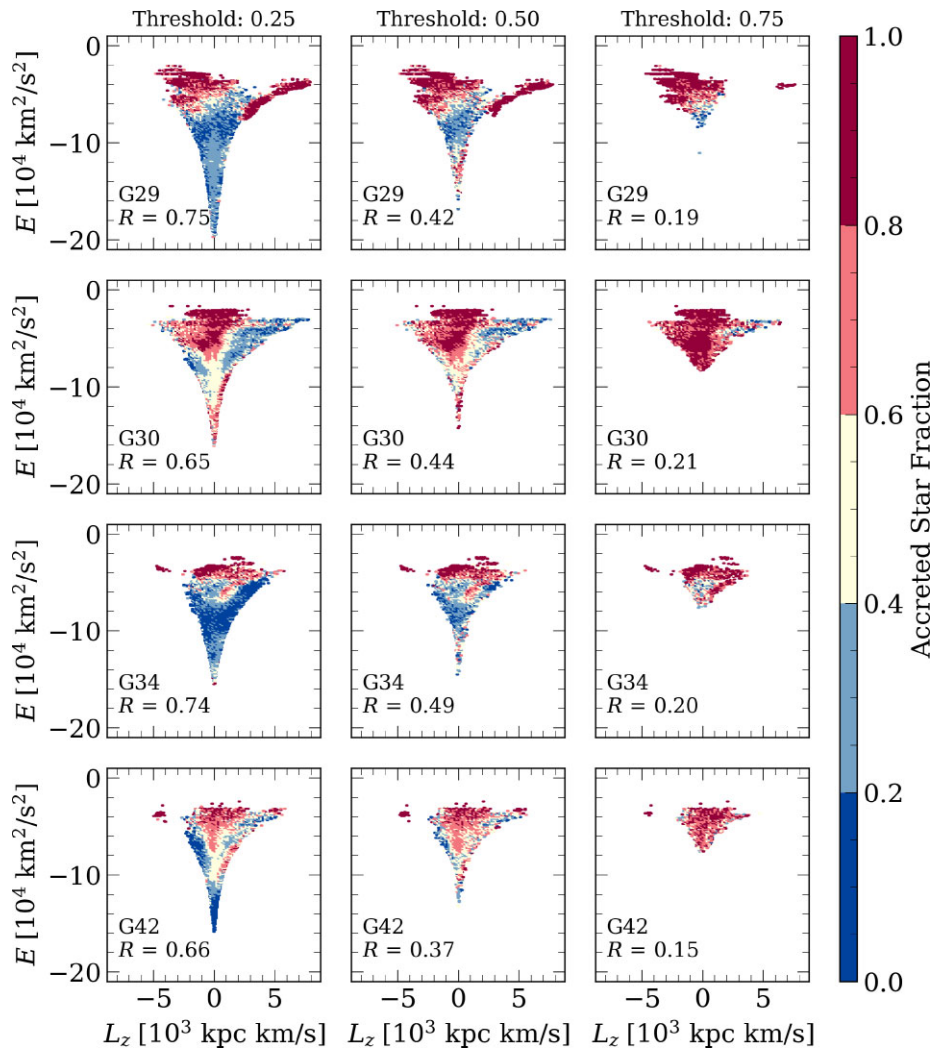


Figure 7. The distribution in the E - L_z plane of the accreted sample retrieved by the MLP model, at different classification thresholds (left to right panels), for each of the galaxy in the test data set (top to bottom rows corresponding to galaxies G29, G30, G34, and G42, respectively.). The completeness (R) of the retrieved sample is also reported. The distribution is colour-coded by the actual fraction of accreted stars as defined by the simulation label.

description of the distribution of stars in the Toomre diagram than the other two selection criteria, although it still provides a very simplified version of the true distribution of the two populations. As in the case for chemical abundances, the MLP model is able to retrieve the overall patterns in the kinematical distribution of the two stellar populations, in all four test galaxies.

Table 5 includes the purity (P) and completeness (R) of the samples of accreted stars retrieved by the MLP model, using three different classification thresholds (the fiducial value of 0.33, 0.10, and 0.75). These metrics are compared with the equivalent P and R values computed using the labels inferred from the three selection cuts versus the true labels from simulations. All values are computed for the Solar neighbourhoods in each of the four galaxies in the test data set. At the fiducial threshold, the MLP model retrieves the purest samples of accreted stars for all test galaxies. When the classification threshold is lowered to 0.10, the model identifies consistently more than 95 per cent of the accreted stars, while maintaining a precision level very similar to the purest sample retrieved by the selection criteria. As the classification threshold is increased, progressively purer samples of accreted stars are retrieved, however, at the expense of completeness. When the threshold is set to 0.75, the MLP model

is able to create samples of accreted stars, on average, twice as purer as the ones obtained through the observational selection criteria.

In addition to being more accurate in identifying the accreted stars than the selection cuts, the MLP model also retrieves fewer stars labelled accreted (on average, 20 per cent fewer than using the selection cuts). This makes it less computationally expensive, especially when applied on large observational data sets. This could be the preferred methodology for the initial processing of observational data to use for subsequent analysis, for example using clustering algorithms to identify tidal stellar streams.

4.4 Visualization of accreted and *in-situ* structures with UMAP

Further insights on the behaviour of the ML models can be obtained from the distribution of accreted and *in-situ* stars in a 2D plane constructed by the UMAP model. Through the UMAP algorithm, stars with similar properties are pulled together, revealing structures which may be hidden in the usual spatial or chemodynamical parameter spaces.

In Fig. 10, we show the distribution of the training and test examples obtained using the UMAP model. The 2D planes are




Figure 8. The $[\alpha/\text{Fe}]$ – $[\text{Fe}/\text{H}]$ distribution, colour-coded by the fraction of accreted stars, f_{acc} for stars in the simulated Solar neighbourhoods. From top to bottom, the rows correspond to galaxies G29, G30, G34, and G42, respectively. In the columns, the accreted stars are defined by: (1) the simulation label; (2)–(4) the observational selection criteria; and (5) the label predicted by the MLP model. In each panel, we show the actual (column 1) and predicted (columns 2–5) overall fractions of accreted stars in the simulated Solar neighbourhoods.

colour-coded based on the fraction of accreted stars, using the simulation labels (left and middle panels) and the predictions of the MLP model (right panel), respectively. As before, the MLP model is considered representative of all ML models. The training distribution indicates a clear separation between the accreted and *in-situ* stars. The accreted stars appear to be clustered in two main regions: the smallest cluster comprises stars located outside $R \simeq 15$ kpc, while the other includes stars with a broader range of properties, which are also more gradually changing. The stars in this second cluster typically have high velocity dispersions and low $[\text{Fe}/\text{H}]$ abundances. The *in-situ* stars appear to be mostly clustered in an L-shaped region characterized by high-rotational motions. Smaller clusters of *in-situ* stars with specific features are also present in this plane. For instance, the top cluster in this figure is comprised of young, metal-rich stars, located in the inner region of the galaxy; while the rightmost cluster is composed of stars orbiting in the plane of the disc, but at large distances from the centre of the galaxy.

In Fig. 11, we show the distribution of the training data, colour-coded by each of the stellar properties in the optimal set of input parameters. All properties are presented as normalized values. The same structures are observed in the distribution of the test examples once projected in the UMAP-defined space (central panel, Fig. 10);

however, the largest region of the accreted stars shows a significant contamination of *in-situ* examples. This seems to suggest that some *in-situ* stars in the test data set have stellar parameters that resemble those in the accreted examples from the training data set. As expected, these are the examples which the ML models struggle to associate to either the accreted or the *in-situ* classes (as shown in the right panel in Fig. 10), and which comprise the majority of the misclassified stars (in both the FP and FN cases). Moreover, the MLP predictions in the largest accreted region appear to follow a gradient as the model outputs progressively higher probabilities of being accreted, for stars occupying the higher parts of this region. In Fig. 11, a similar, although much shallower, gradient is observed in the R plot, suggesting that the accreted stars from late mergers are identified with higher probabilities by the model, as already concluded from Fig. 5.

5 TESTING THE MODELS ON THE AURIGA SIMULATIONS

We further test the performance of our models with an independent data set, drawn from the AURIGA simulations (Grand et al. 2017). This allows us to investigate whether there are any biases introduced

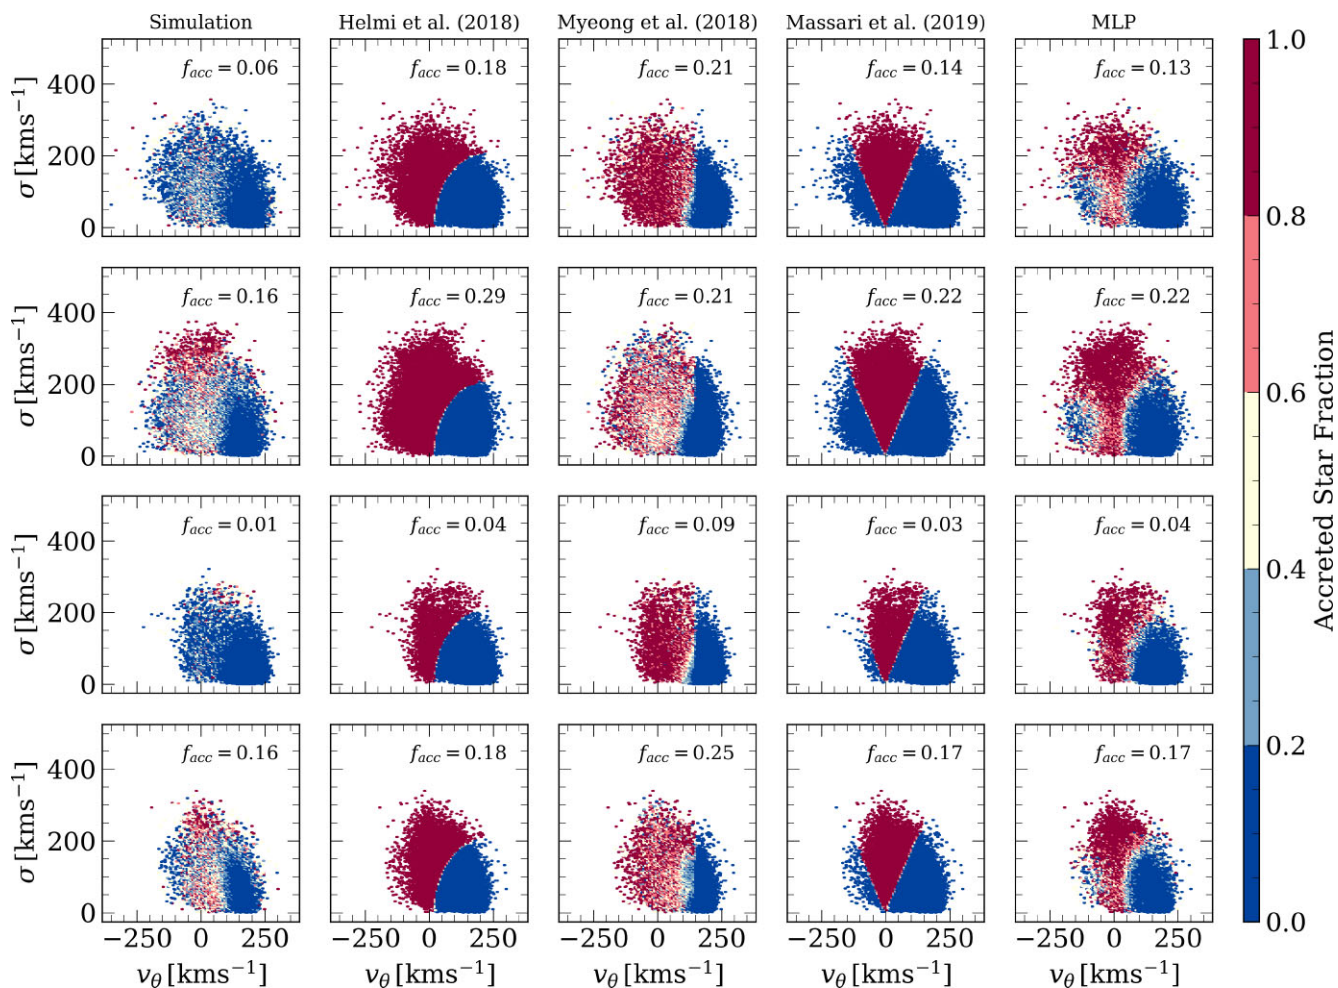


Figure 9. Same as in Fig. 8, but showing results for the distributions of accreted stars in the Toomre diagram.

Table 5. Comparison of the purity (P) and completeness (R) of the sample of accreted stars retrieved by using observational selection cuts (top three rows) and by the MLP model (bottom three rows) evaluated at different thresholds, in the simulated Solar neighbourhoods of galaxies G29, G30, G34, and G42.

	Halo (accreted) Selection criterion	G29		G30		G34		G42	
		P	R	P	R	P	R	P	R
Helmi et al. (2018)	$ \mathbf{V} - \mathbf{V}_{\text{LSR}} > 210 \text{ km s}^{-1}$	0.28	0.81	0.47	0.83	0.21	0.79	0.47	0.54
Myeong et al. (2018)	$[\text{Fe}/\text{H}] < -0.5 \wedge v_{\theta} < 150 \text{ km s}^{-1}$	0.28	0.93	0.49	0.65	0.09	0.87	0.48	0.76
Massari et al. (2019)	$\varepsilon > 0.5$	0.28	0.65	0.48	0.67	0.21	0.68	0.49	0.52
MLP (0.33, fiducial)		0.35	0.75	0.52	0.71	0.22	0.93	0.51	0.56
MLP (0.10)		0.27	0.98	0.46	0.97	0.11	1.00	0.48	0.95
MLP (0.75)		0.70	0.03	0.88	0.05	0.55	0.04	0.76	0.02

in our models due to training only on the ARTEMIS simulations, and to test the classification performance of the models outside their development environment. For this purpose, we use the ‘level 3’ set of six galaxies from the Auriga Project public data release (Grand et al. 2024): Au6, Au16, Au21, Au23, Au24, and Au27. These are also disc galaxies of MW mass, with total masses ranging between $1.04 - 1.74 \times 10^{12} M_{\odot}$ and disc-to-total ratios, D/T , ranging from 0.63 to 0.83 (Grand et al. 2017).

The simulations were run with the hydrodynamical code AREPO (Springel 2010), which includes physical subgrid prescriptions (Vogelsberger et al. 2013) that are significantly different from those implemented in the EAGLE code which was used for ARTEMIS. Furthermore, these simulations have somewhat higher numerical

resolution than ARTEMIS, with dark matter particle masses of $\sim 4 \times 10^4 M_{\odot}$ and baryonic masses of $\sim 5 \times 10^3 M_{\odot}$, respectively. However, we find that the simulated galaxies in AURIGA have similar spatial distributions of *in situ* accreted stars as in the ARTEMIS (not shown here). The level of differences in the spatial distributions between galaxies drawn from these two suites of simulations is comparable with the differences seen between galaxies from the same suite.

We apply the same ML models on these six systems and test their performance. First, we divide the sample into a training and a test data set, setting aside a fraction of stars in the training data set for validation purposes. Galaxies Au6 and Au21 experienced the most massive merging events with mass ratios of 0.54 and 0.53

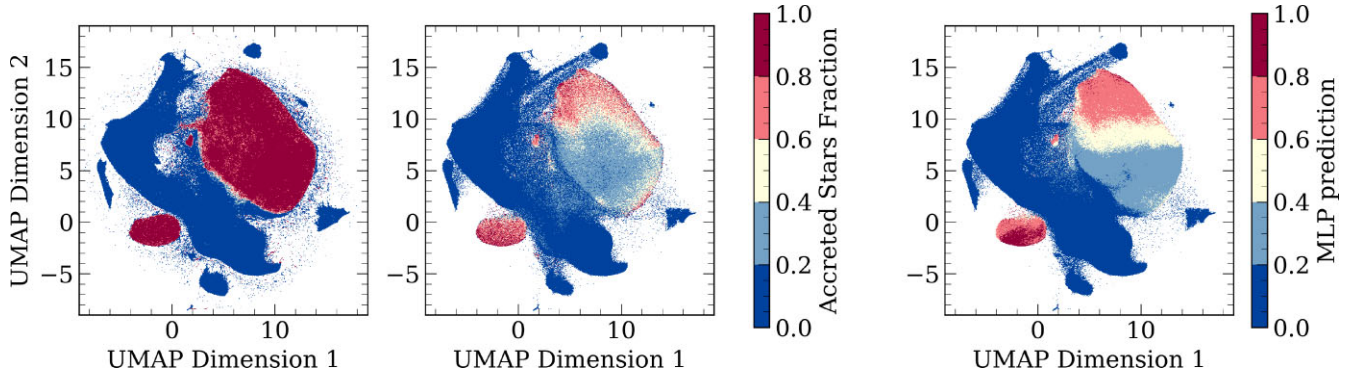


Figure 10. Distribution of stars in the training (left panel) and test (central and right panels) data sets in the parameter space defined by UMAP to maximize the separation between the accreted and *in-situ* populations. Colours represent the fraction of accreted stars in each region of the plane, as defined by the simulation labels (left and central panels) and by the predictions of the MLP model (right panel).

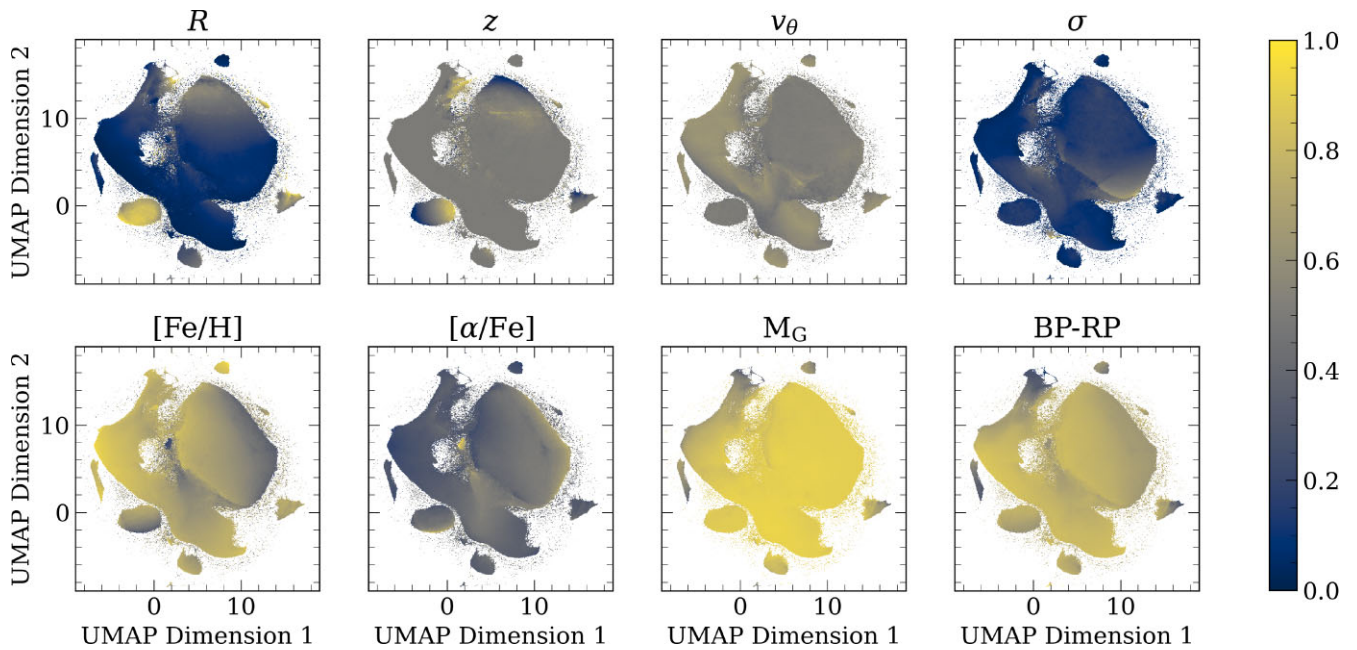


Figure 11. UMAP projections of the distribution of accreted and *in-situ* examples in the training data, weighted by the value of the corresponding physical parameters from the optimal set of features, used as input to the UMAP model. All values are normalized.

(Grand et al. 2018), respectively, and were thus selected for the test data set. The rest of the galaxies in the sample (Au16, Au23, Au24, and Au27) were used to provide training examples to the models.

Fig. 12 shows how the ANNs and decision-tree models perform when developed on the AURIGA data. The MLP, MLP + galaxy features, and the TML models have a slightly better classification performance than XGBoost. This may be caused by a missed optimization of the XGBoost hyperparameters (as we use the same hyperparameters tuned for the ARTEMIS data). Alternatively, there could be more complex non-linear relations between the features describing the accreted stars which are better modelled by the ANNs. Nevertheless, similarly to what it was found for ARTEMIS, all these four models have a similar classification performance. This is encouraging, as it suggests that the models are able to extract the relevant relationships between the accreted and *in-situ* stars, regardless of the type of simulation they were developed on. The benchmark model shows an improved performance compared

to the ARTEMIS analogue (PR-AUC score of 0.45 versus 0.36) suggesting a clearer chemical distinction between accreted and *in-situ* stars in the AURIGA galaxies. Nevertheless, the benchmark model returns samples of the lowest purity compared to the other models developed in AURIGA, showing that, as for ARTEMIS, the addition of kinematic and photometric informations improves significantly the classification.

The MLP, MLP + galaxy features, and the TML models share a similar classification performance, with PR-AUC scores of 0.59, 0.60, and 0.62, respectively, while the XGBoost model shows a drop in performance (0.53). These scores are similar to those found for ARTEMIS, which are 0.59 (MLP), 0.55 (MLP + galaxy features), 0.57 (TML), and 0.59 (XGBoost), respectively (see Section 4).

We perform an additional test, whereby the models developed on the training data set from ARTEMIS are applied on the AURIGA test dataset, and vice versa (purple lines and orange dots in Fig. 13, respectively). This allows us to further investigate the robustness of the models. We also compare the performance




Figure 12. Precision (P) and recall (R) at different classification thresholds for the benchmark model, MLP, MLP + galaxy features, TML, and XGBoost, trained on stars from six AURIGA galaxies. The metrics are evaluated considering all stars in galaxies Au6 and Au21.

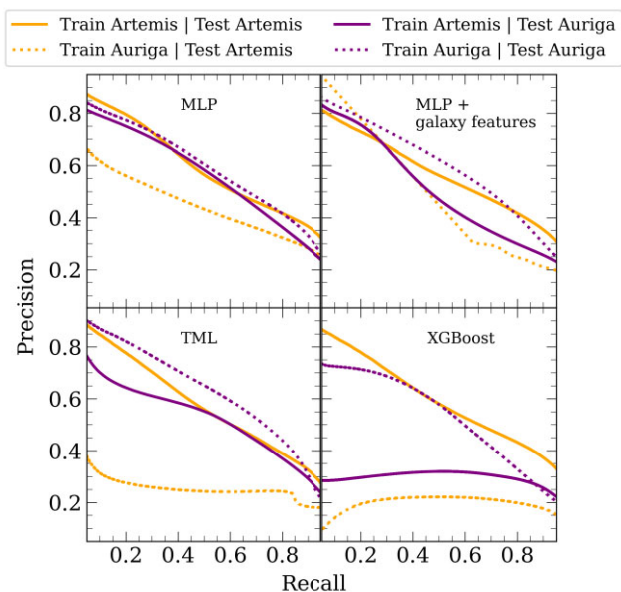


Figure 13. Comparison of the classification performance of models trained on ARTEMIS (continuous lines) and AURIGA (dotted lines). The performance of models tested on data from the same simulation is represented using the same colour, as indicated in the legend.

of criss-crossing the simulations with the performance of using sets from the same type of simulation (shown by orange lines for ARTEMIS models and purple dots for the AURIGA ones in Fig. 13).

Interestingly, when the models trained on AURIGA are applied to the test data from ARTEMIS, the classification performance drops drastically. The same is observed when the models trained on the ARTEMIS data are applied to the AURIGA test data. The lack of consistency in the classification performance of the model can be explained in terms of both data and model complexity.

Differences in the code, physical model, or the numerical resolution of the two simulations inevitably result in a domain shift between the two data sets; consequently, the models trained and tested on different simulations show a drop in classification performance. Domain adaptation techniques, such as described in Čiprijanović et al. (2020), could be explored to develop models that can maintain a consistent classification performance when applied across simulations or on observational data. Models with a high level of complexity can be affected by overfitting as they capture simulation-specific patterns while learning the distinction between accreted and *in-situ* stars. Combined with domain shift, overfitting leads to more drastic performance drops. As shown in Fig. 13, this is the case of the XGBoost model, which has sufficient complexity to extract external patterns from the data, being effectively fine-tuned on the set of simulations it is trained on. The MLP with galaxy features model also performs better on the simulation it is trained on. This can be due to the galaxy-specific properties may be affected by the different galaxy formation models used in ARTEMIS and AURIGA simulations. For the MLP and TML methods, the models developed on ARTEMIS data have a better out-of-sample classification performance than their AURIGA counterparts. Despite being similar in size, the ARTEMIS training data set comprises a larger number of assembly histories leading to a wider variety of accreted star properties learned by the models.

Fig. 13 shows that the MLP model trained on ARTEMIS data is the only model that maintains a consistent classification performance across the two simulations. This is probably due to its simpler model architecture, which makes it less sensitive to overfitting, as well as to the larger sample of assembly histories in the ARTEMIS training set (12 galaxies versus 4 in the AURIGA set). Because of its more consistent performance when applied to an entirely different set of simulations, the MLP model is perhaps better suited to be applied on entirely unseen data, such as observational data from the MW.

6 CONCLUSIONS

In this study, we have investigated the performance of different ML models in separating accreted from *in-situ* stars in MW-mass galaxies, using data from the ARTEMIS simulations. We developed and compared models based on ANN and decision trees algorithms and compared the performance of these models using the usual performance metrics and other physically motivated diagnostics.

The main conclusions of this study are as follows:

(i) We find that the optimal set of stellar features for the models includes a combination of positions and kinematics (R , z , v_θ , σ), photometry (M_G , $BP - RP$), and chemical abundances ($[\alpha/Fe]$ and $[Fe/H]$). A similar classification performance is found with a slightly reduced set of features, including positions, kinematics, and photometry (Fig. 1).

(ii) All ML models investigated here have good classification performances (Fig. 3). Moreover, MLP, TML, and XGBoost perform similarly well also in terms of the distribution of FPs and FNs in a chemodynamical parameter space (Figs 4 and 5). The majority of mis-classified accreted stars are within a 5 kpc radius from the centres of the test galaxies. However, the models perform reasonably well in the regions dominated by *in-situ* stars (e.g. the disc), and typically, the fractions of FPs and FNs are only a few per cent. In the outer regions of galaxies, the MLP, TML, and XGBoost models retrieve more than 90 per cent of the accreted stars with an accuracy close to 60 per cent, at the fiducial thresholds. These results are remarkable, given that these models belong to different families of ML methods,

suggesting that the similar underlying patterns related to accreted and *in-situ* stars are learned in all cases. Similar patterns are also retrieved by the UMAP dimensionality reduction method (Figs 10 and 11).

(iii) Of all models investigated here, MLP is less sensitive to performance drops due to domain shift related to the specificity of accretion history of any given galaxy, and could be the preferred option to use on observational data. However, developing an ML model that is able to learn the galaxy-specific properties (namely, the accretion histories of galaxies) remains a challenge. We have found that adding average galaxy properties, such as stellar masses or metallicities, to the set of input features does not improve the classification performance (specifically, the MLP + galaxy features model), and on the contrary, it leads to overfitting (Fig. 3). Also, combining the predictions of multiple models using ensemble learning (i.e. the TML model) does not provide an improvement, and more complex domain adaptation techniques may be needed to address this issue.

(iv) As shown by the importance gain in the XGBoost model, the parameters which provide the most accurate distinction between the two populations are: the rotational velocity (v_θ), the galactocentric distance in the plane of the disc (R), and the [Fe/H] abundance (Fig. 6).

(v) The purity (P) of the sample of accreted stars retrieved by the models can be increased by adjusting the classification threshold (Fig. 7), however this comes at the cost of completeness (R). The accreted stars identified more accurately by the models have chemodynamical properties associated with late accretions, and are located in the outer regions of galaxies. This suggests that adjusting the classification threshold of ML models can also be used to identify tidal streams in the outer halo.

(vi) We also compared the performance of the ML models versus imposing common observational selection cuts (either in space, kinematics or chemistry) to separate accreted stars from those formed *in situ*. We have found that ML models outperform in purity these more traditional methods (Figs 8 and 9). Therefore, ML models may be applied directly on observational data without the need of additional selection criteria. Thus, they may help in the search for accreted substructures even in the regions dominated by the disc.

(vii) Finally, we have tested the models on a different suite of cosmological simulations (namely, on AURIGA), to evaluate their performance on unseen data (Figs 12 and 13). In general, we find that the models perform similarly well on AURIGA as on ARTEMIS, which suggests that they may be also suitable to be applied on other types of previously unseen data, for example, on observations. Of all the models, XGBoost has the least performance on an unseen data set, possibly because it uses more detailed properties, which differ between the two sets of simulations. In contrast, the MLP appears to be using more broadbrush properties that are relevant to the overall trends between features. These results highlight the importance of testing not only of different ML models, but of different training sets as well.

Our study has shown that ML methods can efficiently separate accreted from *in-situ* stars in galaxies like the MW. These methods perform optimally with a combination of kinematics and chemical abundances, and can improve the detection of accreted substructures in regions of the Galaxy that are highly dominated by *in-situ* stars, and which have not been fully explored to date for identifying substructure. A wealth of high precision data are already available for millions of MW stars, from both astrometric observations, for example, with *Gaia* (Gaia Collaboration 2018, 2023), and spectro-

scopic measurements of chemical abundances, from surveys such as the SDSS Apache Point Observatory Galactic Evolution Experiment (APOGEE, Majewski et al. 2017), the Galactic Archaeology with HERMES (GALAH, De Silva et al. 2015), the Large Sky Area Multi-Object Fibre Spectroscopic Telescope (LAMOST, Zhao et al. 2012), the William Herschel Telescope Enhanced Area Velocity Explorer (WEAVE, Dalton et al. 2014), or the 4-metre Multi-Object Spectroscopic Telescope (4MOST, de Jong et al. 2019). ML methods such as the ones developed here can be directly deployed on this combined, multidimensional parameter space to help in the discovery of accreted substructures.

ACKNOWLEDGEMENTS

AS acknowledges an STFC PhD studentship at the LIV.INNO Centre for Doctoral Training ‘Innovation in Data Intensive Science’. The models have been trained and tested on LJMU’s HPC facility, Prospero. IGM has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 769130). The construction of the ARTEMIS simulations used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility. The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

All codes and models are publicly available at <https://github.com/ariasant/ML-accreted-vs-insitu>. Data from the ARTEMIS simulations may be shared on reasonable request to the corresponding author.

REFERENCES

- Abadi M. et al., 2016, TensorFlow: A System for Large Scale Machine Learning. In 12th USENIX Symposium in Operating Systems Design and Implementation (OSDI 16) 265-283. Software available at: <https://www.tensorflow.org/>
- Abadi M. G., Navarro J. F., Steinmetz M., Eke V. R., 2003, *ApJ*, 597, 21
- Agarwal D., Aggarwal K., Burke-Spolaor S., Lorimer D. R., Garver-Daniels N., 2020, *MNRAS*, 497, 1661
- Anders F. et al., 2023, *A&A*, 678, A158
- Balbinot E., Santiago B. X., da Costa L. N., Makler M., Maia M. A. G., 2011, *MNRAS*, 416, 393
- Beers T. C. et al., 2012, *ApJ*, 746, 34
- Belokurov V., Erkal D., Evans N. W., Koposov S. E., Deason A. J., 2018, *MNRAS*, 478, 611
- Belokurov V., Kravtsov A., 2024, *MNRAS*, 528, 3198
- Belokurov V., Sanders J. L., Fattahi A., Smith M. C., Deason A. J., Evans N. W., Grand R. J. J., 2020, *MNRAS*, 494, 3880
- Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition. Princeton University Press, Princeton, NJ
- Bland-Hawthorn J., Gerhard O., 2016, *ARA&A*, 54, 529
- Borsato N. W., Martell S. L., Simpson J. D., 2020, *MNRAS*, 492, 1370
- Bovy J., 2014, *ApJ*, 795, 95
- Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *MNRAS*, 427, 127
- Brook C. B., Kawata D., Gibson B. K., Gallart C., Vicente A., 2020, *MNRAS*, 495, 2645
- Bullock J. S., Johnston K. V., 2005, *ApJ*, 635, 931
- Busso G. et al., 2022, [link-to-documentation](#), Gaia DR3 documentation Chapter 5: Photometric data

- Carollo D. et al., 2007, *Nature*, 450, 1020
- Carollo D. et al., 2010, *ApJ*, 712, 692
- Chabrier G., 2003, *PASP*, 115, 763
- Chen B. H., Hashimoto T., Goto T., Kim S. J., Santos D. J. D., On A. Y. L., Lu T.-Y., Hsiao T. Y. Y., 2022, *MNRAS*, 509, 1227
- Chen T., Guestrin C., 2016, in KDD '16: The 22nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining. ACM, New York, NY, USA, p. 785
- Ćiprijanović A., Snyder G. F., Nord B., Peek J. E. G., 2020, *Astron. Comput.*, 32, 100390
- Clarke A. O., Scaife A. M. M., Greenhalgh R., Griguta V., 2020, *A&A*, 639, A84
- Cooper A. P. et al., 2010, *MNRAS*, 406, 744
- Cooper A. P., Parry O. H., Lowing B., Cole S., Frenk C., 2015, *MNRAS*, 454, 3185
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- Dalton G. et al., 2014, in Ramsay S. K., McLean I. S., Takami H. eds, Proc. SPIE Conf. Ser. Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. SPIE, Bellingham, p. 91470L
- Das P., Hawkins K., Jofré P., 2020, *MNRAS*, 493, 5195
- de Jong R. S. et al., 2019, *The Messenger*, 175, 3
- De Silva G. M. et al., 2015, *MNRAS*, 449, 2604
- Dillamore A. M., Belokurov V., Font A. S., McCarthy I. G., 2022, *MNRAS*, 513, 1867
- Domínguez Sánchez H., Huertas-Company M., Bernardi M., Tuccillo D., Fischer J. L., 2018, *MNRAS*, 476, 3661
- Ester M., Kriegel H.-P., Sander J., Xu X., 1996, in ACM, KDD '96: The 2nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining. p. 226
- Font A. S. et al., 2020, *MNRAS*, 498, 1765
- Font A. S., Johnston K. V., Bullock J. S., Robertson B. E., 2006a, *ApJ*, 638, 585
- Font A. S., Johnston K. V., Bullock J. S., Robertson B. E., 2006b, *ApJ*, 646, 886
- Font A. S., McCarthy I. G., Belokurov V., 2021, *MNRAS*, 505, 783
- 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., Le Brun A. M. C., Crain R. A., Kelvin L. S., 2017, *Publ. Astron. Soc. Aust.*, 34, e050
- Freeman K., Bland-Hawthorn J., 2002, *ARA&A*, 40, 487
- Friedman J. H., 2001, *Ann. Stat.*, 29, 1189
- Gaia Collaboration, 2018, *A&A*, 616, A1
- Gaia Collaboration, 2023, *A&A*, 674, A38
- Gallart C., Zoccali M., Aparicio A., 2005, *ARA&A*, 43, 387
- Gilbert K. M., Font A. S., Johnston K. V., Guhathakurta P., 2009, *ApJ*, 701, 776
- Gómez F. A., Helmi A., Brown A. G. A., Li Y.-S., 2010, *MNRAS*, 408, 935
- Grand R. J. J. et al., 2017, *MNRAS*, 467, 179
- Grand R. J. J. et al., 2018, *MNRAS*, 474, 3629
- Grand R. J. J., Fragkoudi F., Gómez F. A., Jenkins A., Marinacci F., Pakmor R., Springel V., 2024, preprint (arXiv:2401.08750)
- Grillmair C. J., Freeman K. C., Irwin M., Quinn P. J., 1995, *AJ*, 109, 2553
- Grimozzi S. E., Font A. S., De Rossi M. E., 2024, *MNRAS*, 530, 95
- Hawkins K. et al., 2015, *MNRAS*, 447, 2046
- Hawkins K., Jofré P., Gilmore G., Masseron T., 2014, *MNRAS*, 445, 2575
- Helmi A., 2020, *ARA&A*, 58, 205
- Helmi A., Babusiaux C., Koppelman H. H., Massari D., Veljanoski J., Brown A. G. A., 2018, *Nature*, 563, 85
- Helmi A., White S. D. M., 1999, *MNRAS*, 307, 495
- Hinshaw G. et al., 2013, *ApJS*, 208, 19
- Horta D. et al., 2023, *MNRAS*, 520, 5671
- Huertas-Company M., Lanusse F., 2023, *Publ. Astron. Soc. Aust.*, 40, e001
- Jayasinghe T. et al., 2019, *MNRAS*, 486, 1907
- 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
- Johnston K. V., Hernquist L., Bolte M., 1996, *ApJ*, 465, 278
- Johnston K. V., Spergel D. N., Haydn C., 2002, *ApJ*, 570, 656
- Kingma D. P., Ba J., 2014, preprint (arXiv:1412.6980)
- Klambauer G., Unterthiner T., Mayr A., Hochreiter S., 2017, preprint (arXiv:1706.02515)
- Koppelman H. H., Helmi A., 2021, *A&A*, 645, A69
- Koppelman H. H., Helmi A., Massari D., Price-Whelan A. M., Starkenburg T. K., 2019, *A&A*, 631, L9
- Koppelman H., Helmi A., Veljanoski J., 2018, *ApJ*, 860, L11
- 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
- Mackereth J. T. et al., 2019, *MNRAS*, 489, 176
- Majewski S. R. et al., 2017, *AJ*, 154, 94
- Malhan K. et al., 2022, *ApJ*, 926, 107
- Malhan K., Ibata R. A., 2018, *MNRAS*, 477, 4063
- Malhan K., Rix H.-W., 2024, *ApJ*, 964, 104
- Marigo P. et al., 2017, *ApJ*, 835, 77
- Massari D., Koppelman H. H., Helmi A., 2019, *A&A*, 630, L4
- McCarthy I. G., Font A. S., Crain R. A., Deason A. J., Schaye J., Theuns T., 2012, *MNRAS*, 420, 2245
- McInnes L., Healy J., Melville J., *JOSS*, 2018, 3, 29, 861
- McInnes L., Healy J., Saul N., Grossberger L., 2018b, *J. Open Source Softw.*, 3, 861
- McMillan P. J., 2017, *MNRAS*, 465, 76
- McMillan P. J., Binney J. J., 2008, *MNRAS*, 390, 429
- Monachesi A., Gómez F. A., Grand R. J. J., Kauffmann G., Marinacci F., Pakmor R., Springel V., Frenk C. S., 2016, *MNRAS*, 459, L46
- Myeong G. C., Evans N. W., Belokurov V., Amorisco N. C., Koposov S. E., 2018, *MNRAS*, 475, 1537
- Naidu R. P. et al., 2022, preprint (arXiv:2204.09057)
- Nakoneczny S. J. et al., 2021, *A&A*, 649, A81
- Nakoneczny S., Bilicki M., Solarz A., Pollo A., Maddox N., Spiniello C., Brescia M., Napolitano N. R., 2019, *A&A*, 624, A13
- Necib L., Ostdiek B., Lisanti M., Cohen T., Freytsis M., Garrison-Kimmel S., 2020, *ApJ*, 903, 25
- Olier I., Orhobor O. I., Dash T., Davis A. M., Soldatova L. N., Vanschoren J., King R. D., 2021, *Proc. Natl. Acad. Sci.*, 118, e2108013118
- Ostdiek B. et al., 2020, *A&A*, 636, A75
- Pfeffer J., Kruijssen J. M. D., Crain R. A., Bastian N., 2018, *MNRAS*, 475, 4309
- Pillepich A. et al., 2018, *MNRAS*, 475, 648
- Pillepich A., Madau P., Mayer L., 2015, *ApJ*, 799, 184
- Quiñero-Candela J., Sugiyama M., Schwaighofer A., Lawrence N. D., 2009, Dataset Shift in Machine Learning. p. MIT Press, 248, Cambridge, Massachusetts
- Robertson B., Bullock J. S., Font A. S., Johnston K. V., Hernquist L., 2005, *ApJ*, 632, 872
- Robin A. C. et al., 2012, *A&A*, 543, A100
- Rockosi C. M. et al., 2002, *AJ*, 124, 349
- Sanders J. L., Binney J., 2016, *MNRAS*, 457, 2107
- Sanderson R. E. et al., 2020, *ApJS*, 246, 6
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Searle L., Zinn R., 1978, *ApJ*, 225, 357
- Shih D., Buckley M. R., Necib L., Tamasas J., 2022, *MNRAS*, 509, 5992
- Shipp N. et al., 2019, *ApJ*, 885, 3
- Shipp N. et al., 2023, *ApJ*, 949, 44
- Springel V. et al., 2005, *Nature*, 435, 629
- Springel V., 2010, *MNRAS*, 401, 791
- Tissera P. B., Scannapieco C., Beers T. C., Carollo D., 2013, *MNRAS*, 432, 3391
- Traven G. et al., 2017, *ApJS*, 228, 24
- Tremaine S., 1999, *MNRAS*, 307, 877
- Tronrud T., Tissera P. B., Gómez F. A., Grand R. J. J., Pakmor R., Marinacci F., Simpson C. M., 2022, *MNRAS*, 515, 3818
- Trujillo-Gomez S., Kruijssen J. M. D., Pfeffer J., Reina-Campos M., Crain R. A., Bastian N., Cabrera-Ziri I., 2023, *MNRAS*, 526, 5735
- Veljanoski J., Helmi A., Breddels M., Posti L., 2019, *A&A*, 621, A13

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

Zhao G., Zhao Y.-H., Chu Y.-Q., 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.