



Universiteit
Leiden
The Netherlands

A sparse regression approach for populating dark matter haloes and subhaloes with galaxies

Icaza-Lizaola, M.; Bower, R.G.; Norberg, P.; Cole, S.; Schaller, M.

Citation

Icaza-Lizaola, M., Bower, R. G., Norberg, P., Cole, S., & Schaller, M. (2023). A sparse regression approach for populating dark matter haloes and subhaloes with galaxies. *Monthly Notices Of The Royal Astronomical Society*, 518(2), 2903-2920.
doi:10.1093/mnras/stac3265

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3566591>

Note: To cite this publication please use the final published version (if applicable).

A sparse regression approach for populating dark matter haloes and subhaloes with galaxies

M. Icaza-Lizaola^{1,2,3}★, Richard G. Bower^{1,2,4}, Peder Norberg^{1,2,4}, Shaun Cole¹ and Matthieu Schaller^{5,6}

¹*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

²*Institute for Data Science, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

³*Korea Astronomy and Space Science Institute, 776 Daedeok-daero, Yuseong-gu, Daejeon 34055, Republic of Korea*

⁴*Centre for Extragalactic Astronomy, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

⁵*Lorentz Institute for Theoretical Physics, Leiden University, PO Box 9506, NL-2300 RA Leiden, the Netherlands*

⁶*Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*

Accepted 2022 November 7. Received 2022 November 1; in original form 2021 October 1

ABSTRACT

We use sparse regression methods (SRMs) to build accurate and explainable models that predict the stellar mass of central and satellite galaxies as a function of properties of their host dark matter haloes. SRMs are machine learning algorithms that provide a framework for modelling the governing equations of a system from data. In contrast with other machine learning algorithms, the solutions of SRM methods are simple and depend on a relatively small set of adjustable parameters. We collect data from 35 459 galaxies from the EAGLE simulation using 19 redshift slices between $z = 0$ and $z = 4$ to parametrize the mass evolution of the host haloes. Using an appropriate formulation of input parameters, our methodology can model satellite and central haloes using a single predictive model that achieves the same accuracy as when predicted separately. This allows us to remove the somewhat arbitrary distinction between those two galaxy types and model them based only on their halo growth history. Our models can accurately reproduce the total galaxy stellar mass function and the stellar mass-dependent galaxy correlation functions ($\xi(r)$) of EAGLE. We show that our SRM model predictions of $\xi(r)$ is competitive with those from subhalo abundance matching and might be comparable to results from extremely randomized trees. We suggest SRM as an encouraging approach for populating the haloes of dark matter only simulations with galaxies and for generating mock catalogues that can be used to explore galaxy evolution or analyse forthcoming large-scale structure surveys.

Key words: methods: statistical – galaxies: evolution – galaxies: haloes – (cosmology:) dark matter.

1 INTRODUCTION

Within the Lambda cold dark matter (Λ CDM) paradigm (e.g. Planck Collaboration I 2014), an expanding universe filled with particles that interact only through gravity can be accurately modelled using N -body simulations (e.g. Springel et al. 2005). Because of advances in computational methods, such simulations can track the formation of galaxy-scale dark matter haloes within volumes approaching the size of the observable Universe. However, these simulations do not include the baryonic component that leads to the formation of stars and galaxies. Hydrodynamical simulations that include baryons need to deal with complicated cooling and feedback processes and are strongly influenced by events happening at scales much smaller than the size of the simulation grid. This makes them significantly more expensive to run and limits their volume to about 1 Gpc^3 (e.g. Springel et al. 2018). There is, therefore, an incentive for a hybrid approach, in which one uses hydrodynamic simulations to learn the relation between dark matter and baryonic tracers, and then uses these relations to populate N -body mock catalogues of larger volume.

In Icaza-Lizaola et al. (2021), we present a novel methodology that uses sparse regression methods (SRMs; Tibshirani 1996; Hastie, Tibshirani & Wainwright 2015) to model the relations between the stellar mass of a galaxy and its host halo in the Evolution and Assembly of Galaxies and their Environments (EAGLE; Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016) 100 Mpc hydrodynamical simulation. SRMs are a set of machine learning (ML) algorithms designed to identify the parameters that better describe a dependent variable, then discard the remaining unnecessary ones. Recently, they have been suggested as the appropriate framework to extract the equation of states of a physical system from collected data and with minimal knowledge of the physics of the system (Brunton, Proctor & Kutz 2016).

In Icaza-Lizaola et al. (2021), we were interested in developing and testing the methodology in a simple scenario without going into some of the more complicated challenges that populating a realistic N -body mock accurately would require. With that in mind, we tested our methodology on central galaxies (the main galaxy within each dark matter halo) only as they have monotonic growths with time which makes them easier to model. In this work, we extend our methodology to include satellite galaxies as well. Satellite galaxies (and their associated dark matter subhaloes) are

* E-mail: icaza@kasi.re.kr

created when a smaller dark matter halo is accreted by a larger one. This is a common process in the Λ CDM model. As they orbit within the larger halo, satellite galaxies (and their remnant dark matter subhaloes) undergo a much more diverse range of physical processes than their central galaxy counterparts.

Unlike the main dark matter halo, which undergoes monotonic mass growth, the remnants of smaller accreted haloes may decay with time (e.g. Bower & Balogh 2004; van den Bosch et al. 2018), as they lose mass due to processes such as tidal stripping and heating (Lynden-Bell 1967; Merritt 1983; Hayashi et al. 2003; Green & van den Bosch 2019). Moreover, the satellite galaxies residing inside these remnant haloes are subject to ‘environmental’ processes that remove cold gas and suppress the accretion of more material (Gunn & Gott 1972; Vollmer et al. 2001; Larson, Tinsley & Caldwell 1980; Bahé & McCarthy 2015; Correa, Schaye & Trayford 2019). As a result, star formation in satellite galaxies is significantly suppressed compared to central galaxies and we expect less stellar mass growth.

In EAGLE, the differentiation between central haloes and subhaloes is done by the SUBFIND algorithm (Springel et al. 2001). Within each halo, the algorithm identifies the self-bound overdensities and classifies them as independent subhaloes. The subhalo with the lowest potential energy is classified as the central halo and assigned any diffuse mass that has not already been associated with a subhalo. This distinction is made separately at each output time and is not a fundamental differentiation, but dependent on the details of the algorithm. In some cases, this leads to anomalous behaviour, in particular inconsistent classifications of the same subhalo at different redshift slices (e.g. Behroozi et al. 2015). It is, therefore, desirable to use a methodology that does not make a fundamental distinction between central and satellite galaxies when modelling the stellar mass, but rather to use the same approach based on the overall halo mass history.

In this paper, we use a lower threshold in the host halo mass for our central galaxy sample compared to Icaza-Lizaola et al. (2021), reducing it from $M = 10^{11.1} M_\odot$ to $M = 10^{10.6} M_\odot$. This allows us to identify low mass haloes which contain relatively large galaxies (with stellar masses greater than $10^9 M_\odot$). This is a particularly important consideration for satellite galaxies, if we are to generate a stellar-mass complete catalogue.

Other works have used ML algorithms to model the relationship between the halo and stellar properties inside a hydrodynamical simulation (e.g. Kamdar, Turk & Brunner 2016; Agarwal, Davé & Bassett 2018). Their models accurately reproduce several statistics of the original simulation. However, given that these types of models generate *black box* answers it might be complicated to modify them to reproduce statistics from observations instead. Lovell et al. (2022) trains an extremely randomized tree (ERT; Geurts, Ernst & Wehenkel 2006) model on data from the EAGLE simulation and uses it to populate the P-Millennium N -body simulation with galaxies (Baugh et al. 2018). Moster et al. (2021) uses a neural network approach that rewards the algorithm for reproducing observed statistics of a survey [like correlation functions and stellar mass functions (SMFs)] instead of properties of individual galaxies. This circumvents the problem of differences in statistics between the hydrodynamical simulation used to calibrate the model and those from an observational survey, at the cost of not requiring accuracy in the predictions of the individual values of galaxy properties. Given that our model is an equation of state with a set of input parameters fitted by the model, it is in principle possible to extract the best advantages of both approaches, extracting the important physical parameters by comparison to the simulation, but optimizing the coefficients of these terms to reproduce the statistics of an observational data set.

This paper is organized as follows. Section 2 summarizes the sparse regression methodology used in this work, with a complete discussion of the methodology presented in Icaza-Lizaola et al. (2021). Section 3 introduces the data set that we use and any enhancements to the model that we have made to handle the more complex data set. In particular, Section 3.1 explains the details of the bijective match between the hydrodynamical EAGLE simulation and the EAGLE dark matter only (EAGLE DM hereafter) simulation. Sections 3.2 and 3.3 describe the methodology used to extract our training data set from the EAGLE DM only simulation as well as the new parametrization of the model and the new weighting scheme adopted. The results from our different models are shown and analysed in Section 4. In Section 4.2, we compare the SMF and the clustering of our resulting models with the ones from the original EAGLE sample. In Section 4.3, we compare our resulting models with some available from the literature. Our conclusions and thoughts on the potential of the current methodology are discussed in Section 5.

2 METHODOLOGY

The methodology followed in this work is presented in detail in Icaza-Lizaola et al. (2021). Here, we include a summary of the key concepts, and then in Sections 3.2 and 3.3, we describe the additions and changes to the methods adopted in this specific work.

Sparse regression methods (SRM; Tibshirani 1996; Hastie et al. 2015; Tibshirani & Friedman 2017) are a set of ML algorithms designed to develop a fitting function by selecting linear combinations from a large library of candidate functional forms. The method selects only a minimal subset of functions from the library such that the combination describes the input data well but does not overfit and hence avoids poor interpolation between input points. One key advantage of SRM methods over other ML techniques is that the resulting model is in the form of an equation with nominally a small subset of terms, making it more likely to have a clear physical interpretation.

In Icaza-Lizaola et al. (2021), we used SRM to model the relation between the stellar mass (M^*) of central galaxies and a set of properties of their host haloes. We found that a good, but simple description could be obtained based on the final mass of the host halo and its parametrized formation history. In this paper, we aim to provide a similar relationship that describes all galaxies in the simulation, whether they are the dominant galaxy within the halo (which we refer to as *central*) or a galaxy that was formed in a separate sub-halo that has been subsequently been accreted (we refer to such galaxy as a *satellite*). Although we follow very similar methodologies to Icaza-Lizaola et al. (2021), the halo and subhalo properties used as input parameters here have been adapted so that we can model both satellites and central galaxies consistently. The details are described in Section 3.2.

Let us call M the number of host halo properties, and N the total number of galaxies in our data set. For each halo property we define the vector $\vec{x}'_i = [x'_{i1}, \dots, x'_{iN}]$ that contains the observed value of the i th property of each host halo ($i \leq M$).

The input halo properties need to be standardized so that they all vary within a consistent range. This is done using the following transformation:

$$\vec{x}_i = \frac{\vec{x}'_i - \mu(\vec{x}'_i)}{\sigma(\vec{x}'_i)}, \quad (1)$$

where μ and σ are the mean and standard deviation operators, respectively. We use the $\vec{x}_i$ vectors to build a set of D polynomial

functions $F_l(\vec{x})$ ($l < D$), where the function F_l can be either a linear, a quadratic or cubic combination of the dependent variables, i.e. $F_{l\alpha} = x_{i\alpha}$ or $F_{l\alpha} = x_{i\alpha} \times x_{j\alpha}$ or $F_{l\alpha} = x_{i\alpha} \times x_{j\alpha} \times x_{k\alpha}$, where $1 \leq i \leq j \leq k \leq M$ and $\alpha < N$. We use all possible linear quadratic and cubic combinations of the input properties and so $D = 1 + M + M(M + 1)/2 + M(M + 1)(M + 2)/6$.

The value of the stellar mass predicted by our model for galaxy α ($M_{p\alpha}^*$)¹ is expressed as the linear combination of functions $F_{l\alpha}$:

$$M_{p\alpha}^* = \sum_{l=0}^D C_l F_{l\alpha}, \quad (2)$$

where $\vec{C} = [C_0, \dots, C_D]$ are a set of coefficients. The optimal values of these coefficients are the quantities determined by our methodology.

Following the SRM approach, most coefficients are discarded (i.e. we set $C_l = 0$) and only a small subset of the possible coefficients are retained. This is achieved by minimizing the LASSO function defined as

$$L(\vec{C}) = \chi^2(\vec{C}) + \lambda P(\vec{C}), \quad (3)$$

where $\chi^2(\vec{C})$ is a statistic that determines the goodness of the fit, $P(\vec{C})$ is a penalty term that incentivizes the minimization to discard unnecessary input parameters and λ is a hyperparameter of our methodology that regulates the relative magnitude of $P(\vec{C})$. We define $\chi^2(\vec{C})$ as

$$\chi^2(\vec{C}) = \sum_{\alpha=1}^N \frac{(M_{\alpha}^* - M_{p\alpha}^*(\vec{C}))^2}{\sigma^2}, \quad (4)$$

where σ is an estimate of the uncertainty of the measurement of M_{α}^* (as defined by equation 11 of Icaza-Lizaola et al. 2021).

The penalty term $P(\vec{C})$ is defined in such a way that its value increases significantly with the number of coefficients C_j that are non-zero. The shape of $P(\vec{C})$ is given by the following equation:

$$P(\vec{C}) = \sum_{l=1}^D \left[\sum_{m \neq l} |C_m| e^{-(\epsilon/C_m)^2} \right] |C_l| e^{-(\epsilon/C_l)^2}, \quad (5)$$

where ϵ is a small constant that determines how close to zero a coefficient needs to be before its contribution to the penalty is negligible.

In this work, all coefficients below 10^{-3} are discarded. We refer the reader to Icaza-Lizaola et al. (2021) for a discussion of the choice of this specific value and of the optimal ϵ value for third-order polynomials.

Equation (3) is designed to avoid overfitting the input data, which is a necessity in any model with a large space of input parameters. This is achieved by the balancing between the goodness of fit and the penalty term. An overfitted model would have a small χ^2 by including many non-zero parameters, which, in turn, would make the penalty term very large. Therefore, the minimum of $L(\vec{C})$ should correspond to a model that is as simple as possible (small $P(\vec{C})$), while still being a good fit (small χ^2). The equilibrium between the need to fit the data well and to keep the number of non-zero coefficients small is set by the choice of the λ parameter: a large value strongly reduces the number of coefficients selected, while a small value does not

penalize the goodness of fit enough. We determine the optimal value using the k -fold methodology (Hastie et al. 2015), where the data are separated into a training set and a test set k -times. The optimal value of λ , and its associated uncertainty, can then be determined by examining how well a model fitted to the training set can predict the data in the test set. The full details of this process are described in Icaza-Lizaola et al. (2021).

We use 85 per cent of our data to train our model, with the remaining 15 per cent labelled as the Holdout data set. The latter is used in Section 4.1 to test the accuracy of the method, while the full data set is used in Sections 4.2 and beyond.

3 DATA

Our data set comes from the EAGLE (Crain et al. 2015; Schaye et al. 2015; McAlpine et al. 2016) simulations, which are a suite of hydrodynamical simulations built using the Planck 2014 cosmology (Planck Collaboration I 2014). During the rest of this work, we define the stellar mass of a galaxy in EAGLE as the sum of all stellar particles inside a sphere with an aperture of 30 kpc centred at the centre of the potential of the galaxy.

We use the simulations built in a 100 comoving Mpc box, which is the largest box available. Haloes in the simulation are identified using a Friends-of-Friends (FoF) algorithm (e.g. Davis et al. 1985) with a linking length of $b = 0.2$. Subsequently, the SUBFIND algorithm (Springel et al. 2001) finds the subhaloes within each halo and selects one of them as the central halo. The simulation outputs are saved in 29 snapshots going from $z = 20$ to $z = 0$. The snapshots are used to build merger trees (Qu et al. 2017) by identifying haloes with their progenitors at the previous redshift slice. Main progenitors are defined as the progenitors with the larger branch mass (De Lucia & Blaizot 2007), defined as the sum of the progenitors mass at all previous snapshots. During this work, we use the main progenitor branch to track the mass evolution of a halo.

3.1 Matching

The goal of this work is to develop a fitting function that allows the mass of a galaxy to be estimated from knowledge of its DM halo formation history only. Since DM haloes in hydrodynamical simulations are affected by baryonic processes that might alter their density profile (Navarro, Eke & Frenk 1996; Martizzi et al. 2012; Schaller et al. 2015b), or other properties like the shape of the halo (Katz & Gunn 1991; Bryan et al. 2013), it is important that we match the haloes in the hydrodynamical simulation with the same haloes in a dark matter only simulation (with identical cosmology and initial conditions). By making a one-to-one matching between the DM only simulation and the hydrodynamical one, the properties of the DM only simulations can be used as the input variables of the model (the vectors $\vec{x}_j$ of Section 2), while the stellar mass is measured in the full-physics hydrodynamical simulation. The matching is done by following the procedure of Schaller et al. (2015a). To summarize, we look at the 50 most bound DM particles of each halo or subhalo in the hydrodynamical simulation: if a halo or subhalo of the DM only simulation contains at least half of these particles, then they are matched. The matching is done for all haloes above $M_{\text{total}} > 2 \times 10^9 M_{\odot}$ and both haloes need to be above this value to be matched, where M_{total} is the summed mass of all particles assigned to the halo or subhalo.

Fig. 1 shows the SMF of the full EAGLE hydrodynamical simulation and compares it to the SMF of the galaxies living in haloes that were successfully matched. The bottom panel of Fig. 1

¹ As mentioned later, our code actually models $\log_{10}(M^*/M_{\odot})$. We opt against including the full logarithmic expression in the main text of the paper and associated equations to simplify the notation, while we show the explicit dependencies in the figure labels.




Figure 1. Comparison of the SMF of the full EAGLE simulation (solid lines), with the SMF from the galaxies living in haloes that were successfully matched (dashed lines). The plot shows results for both central haloes (red) and satellites (green), and the combined sample of central and satellites haloes (blue). The bottom panel shows the ratios of the SMF of comparable galaxy types (while keeping the colour coding the same as in the top panel) and quantifies the fraction of matching failures per galaxy type.

shows that the fraction of matching failures for central galaxies is around 1 per cent for all stellar mass scales of interests. This explains why it was not necessary to consider the effect of unmatched haloes in Icaza-Lizaola et al. (2021). However, the number of unmatched satellite galaxies is significantly larger, with a matching success rate around 80 per cent for galaxies with $\log_{10}(M^*/M_\odot) > 10$ (green line in the bottom panel of Fig. 1).

With this in mind, all statistics presented from Section 4.2 onwards result from applying the model to all haloes in the EAGLE DM only simulation (matched and unmatched) and compares them to statistics from all galaxies in the hydrodynamical simulation. This comparison assumes that the distribution of unmatched haloes in both simulations is similar. We explore the validity of this assumption in Appendix A.

3.2 Halo selection and input parametrization

We begin our selection of haloes by tracing the evolution of the halo mass in the DM only simulation at 19 redshift slices between $z = 0$ and $z = 4$. This initial selection is based on $M_{\text{total}}(z)$, the total mass of the particles associated to the halo or subhalo by the SUBFIND algorithm. These trajectories summarize the evolution of the galaxies host halo mass as a function of redshift and give us a relation between halo mass and time for each galaxy. In order to ensure that the trajectory is not overly affected by the algorithm used in the selection process, we use a Gaussian kernel with a σ of one redshift slice to smooth this evolution history. Since halo masses can increase as well as decrease (for satellite galaxies in particular), we base our halo selection on the maximum value of $M_{\text{total}}(z)$ in the smoothed trajectory. The success rate of the matching is dependent on the halo mass, with more massive haloes being more likely to be matched. We find that $\text{Max}(M_{\text{total}}(z)) = 10^{10.66} M_\odot$ corresponds to the threshold at which more than 90 per cent of haloes




Figure 2. Distribution of the central galaxies (blue dots) and satellite galaxies (red dots) in our sample in the $M_{\text{max}}-M^*$ space, where M_{max} is the largest halo mass the halo's main progenitor reached (see equation 6). The solid lines show the median value of the distributions. The plot shows that at a fixed M_{max} the median galaxy mass of a satellite galaxy is larger than that of a central galaxy.

are successfully matched. We define this threshold as the halo mass cut-off of our sample. In order to avoid missing data, we discard those that do not have a well-defined main progenitor in all redshift slices up to $z = 4$. For $\text{Max}(M_{\text{total}}(z)) > 10^{10.66} M_\odot$, this cut is unimportant, with 99.6 per cent of the sample being kept. Our final sample consists of a total of 35 456 galaxies, of which 9967 live inside subhaloes, and 25 489 inside central haloes.

As a pre-processing step, we use the interpolation scheme developed in Icaza-Lizaola et al. (2021) to ensure the halo masses of central galaxies are not affected by inconsistent classification between snapshots. Nominally haloes in our models have their evolution tracked with $M_{\text{total}}(z)$ at all redshifts. We have compared models with different halo mass definitions for centrals, like M_{200}^c ² and found negligible differences on the accuracy of the stellar mass predictions.

Since the satellite halo mass cannot be expected to grow monotonically with decreasing redshift, a more important parameter for each galaxy is instead its maximum halo mass. In the rest of the paper, we refer to this as M_{max} :

$$M_{\text{max}} = \text{Max}(M_{\text{total}}(z)). \quad (6)$$

Central galaxies tend to grow monotonically with time, and M_{max} is correlated with the stellar mass through the $z = 0$ stellar mass–halo mass (SMHM) relation. In satellite galaxies, however, M_{max} corresponds to the redshift at which their host halo merges and becomes the subhalo of a larger system. Once a halo merges the mass of the halo declines due to tidal processes. We can expect, therefore, that the galaxy mass at $z = 0$ will be well correlated with the mass of the host halo before merging. Fig. 2 shows the distribution of galaxies in the $M_{\text{max}}-M^*$ space.

²The mass within a radius for which the density is 200 times larger than the critical density of the Universe. We note that M_{200}^c is only defined for central galaxies in EAGLE.

We note that the median stellar mass of satellite galaxies is larger than that of centrals at fixed $M_{\max}$, i.e. for a fixed $M_{\max}$ satellite galaxies are more massive. The offset in the SMHM relation for satellites and centrals is driven by two competing processes. On the one hand, satellites may undergo a strong suppression of their star formation as they orbit within the main halo due to the combined effects of ram-pressure stripping (the removal of the interstellar medium of the galaxy by ram pressure) and *strangulation* (the absence of gas infall on to the satellite). On the other hand, while the halo mass of the central continues to grow with cosmic time, the satellite reaches its peak mass and $M_{\max}$ becomes frozen thereafter. The net offset is determined by whether the halo mass or the stellar mass grow fastest in the central galaxies, and by whether satellite galaxies are able to continue to grow in stellar mass after they are accreted (Behroozi et al. 2019). Because the effect on the stellar mass growth tends to be delayed compared to the effect on the halo, satellite galaxies tend to have larger stellar mass than their central counterparts.

We now describe the input parameters used in this work, which are the values of the vectors $\vec{x}'_j$ of Section 2.

In Icaza-Lizaola et al. (2021), we tested different parametrizations and concluded that parameters that measure the SMHM relation and the halo growth trajectory are the most useful for modelling the stellar mass at $z = 0$. We also found no improvement in our models when adding parameters correlated with the angular momentum evolution of the halo. The best model that we found used $\log_{10}(M_{200}^c(z = 0)/M_{\odot})$ as the input parameter that traced the SMHM relation, as well as a set of formation criteria parameters $\mathbf{FC}_p$ that model the assembly history, where $\mathbf{FC}_p$ is the redshift by which a central galaxy has assembled $p = [20, 30, 50, 70, 90]$ per cent of its current mass. In order to accommodate satellite galaxies, we substitute the input parameter $M_{200}^c(z = 0)$ with $M_{\max}$ and we define the dimensionless parameter

$$\lg M_{\max} = \log_{10}(M_{\max}/M_{\odot}) \quad (7)$$

and redefine the formation criteria parameters $\mathbf{FC}_p$ as follows. First we find the redshift z_i at which a halo or subhalo reaches $M_{\max}$. Then we look at the evolutionary history of the halo from $z = 4$ up until z_i , and find the redshift ($z_i \leq \mathbf{FC}_p \leq z = 4$) at which the halo has assembled a percentage p of $M_{\max}$.

Note that if z is such that $M(z) = M_{\max}$, then $z < \mathbf{FC}_{90} < \mathbf{FC}_{70}$. This parametrization is almost equivalent to the one used in Icaza-Lizaola et al. (2021) when only considering central galaxies as in this case $M_{200}^c(z = 0) \sim M_{\max}$. As a check, we ran our methodology on the data set of Icaza-Lizaola et al. (2021) with the new parametrization. The resulting model is comparable to the original one in accuracy and simplicity. In total we use six independent variables in our methodology [$\lg M_{\max}$, $\mathbf{FC}_{20}$, $\mathbf{FC}_{30}$, $\mathbf{FC}_{50}$, $\mathbf{FC}_{70}$, $\mathbf{FC}_{90}$]. Each of these parameters is transformed to the standardized space defined by equation 1. Since we consider cubic combinations of these parameters, this leads to a model with up to $D = 84$ parameters.

Many methodologies have found that parameters related to the circular velocity of haloes, like the maximum of the radial circular velocity profile at $z = 0$ ($V_{\max}$) or even the maximum value of $V_{\max}$ among all redshifts (V_{peak}), are more accurate than the halo mass when modelling the stellar mass of their host galaxy (e.g. Conroy, Wechsler & Kravtsov 2006; Chaves-Montero et al. 2016; Matthee et al. 2017; Kamdar et al. 2016; Lovell et al. 2022). In our current implementation, strongly correlated parameters that serve a similar function in the modelling of the stellar mass, like $V_{\max}$, V_{peak} and M_{total} , are not easily distinguished by our algorithm. This leads to subtle variations in the surviving parameters of a given model that

can depend on configuration parameters, like the starting point of the minimization and the specific training set selection. Degeneracies due to correlated model parameters are further discussed in Icaza-Lizaola et al. (2021) and at the end of Section 4.

We have run a model where we use both $M_{\max}$ and V_{peak} as free parameters simultaneously, and compare it to the model with only $M_{\max}$ that we present in the next section. We found no difference in accuracy or simplicity between the two models. However, the fact that one model is a function of both parameters made its interpretation less straightforward. For example, when running our algorithm using only $M_{\max}$, the SMHM relation is modelled as a third-order polynomial of $M_{\max}$ (as we show in Section 4.1), which makes intuitive sense when looking at Fig. 2. However, when using both $M_{\max}$ and V_{peak} , the SMHM function is now modelled by a more complicated function of both parameters. Therefore, by adding parameters that are strongly correlated with $M_{\max}$, we lose explainability without gaining accuracy, and hence we decide to keep only one of the two correlated parameters. In Appendix B, we discuss why we did select $M_{\max}$ instead of V_{peak} . A possibility to work with correlated parameters without the need of doing this sort of correlation analysis beforehand would be to use some principal component analysis (e.g. Jolliffe 2005).

To test the differences between modelling satellite and central galaxies separately and modelling them together with a single model, we run three models independently of each other:

- (i) A model that only contains central galaxies, with $N = 25\,489$ data points.
- (ii) A model that only contains satellite galaxies, with $N = 9967$ data points.
- (iii) A model that combines central and satellite galaxies and fits them all at the same time, with $N = 35,456$ data points.

3.3 Weighting the cost function

In Icaza-Lizaola et al. (2021), we used a simple χ^2 measure to assess the quality of the model's prediction of the data (i.e. χ^2 is the cost function). In the CDM paradigm, however, smaller haloes are always much more numerous than massive ones. As a consequence, such methodology would have a stronger incentive to fit numerous smaller haloes more accurately at the expense of a less accurate fit to less numerous massive ones. In Icaza-Lizaola et al. (2021), we concluded that our methodology became more inaccurate for galaxies larger than $\log_{10}(M^*/M_{\odot}) > 11.0$ (see discussion of fig. 14) due to a relatively small fraction of galaxies above the threshold (90 out of ~ 9500). Given that in this iteration of the work we reduced the cut-off value of galaxies even further, we now have a larger number of smaller galaxies making the issue even more problematic. A good solution to this problem is to assign a weight w'_i to each halo. This weight determines how much of an incentive the code will have to fit a particular halo mass correctly. If the weight w'_i is larger for galaxies in larger haloes, then by modifying equation (4) to include a normalized weight w_i as below, we will give a larger importance to the rarer larger haloes:

$$\chi_w^2 = \sum_{\alpha=1}^N \frac{w_{\alpha} (M_{\alpha}^* - M_{p\alpha}^*(C))^2}{N^2}. \quad (8)$$

To compute the weight of a halo, we first look at the halo mass function (HMF) as a function of $\lg M_{\max}$. To avoid noisy weights from having a small number of objects in the more massive bins, we make use of a linear fit to the HMFs. Referring to the linear fits as

$\text{fl}(\text{lgM}_{\text{max}})$, the weight of a halo is defined as

$$w'_\alpha = \sqrt{\frac{10^{\text{fl}(\mu)}}{10^{\text{fl}(\text{lgM}_{\text{max},\alpha})}}}, \quad (9)$$

where μ is the median value of lgM_{max} . As a final step, we normalize the weights of a sample as follows:

$$w_\alpha = \frac{N \times w'_\alpha}{\sum_{\alpha=1}^N (w'_\alpha)}. \quad (10)$$

We emphasize that in the combined model, the weighting scheme does not distinguish between central and satellite galaxies.

4 RESULTS

We start in Section 4.1 by comparing input and predicted stellar masses, using the holdout data only. As mentioned in Section 2, haloes in the holdout set were not used to train the model. Therefore, comparisons with the holdout data enable the accuracy of our method to be tested by making model predictions on EAGLE data that the model has not seen before. In Section 4.2, we present model predictions using the full data set for the galaxy SMF and galaxy clustering split by stellar mass. In Section 4.3, we compare our EAGLE SRM predictions with an SHAM model (Chaves-Montero et al. 2016) and an ML method (Lovell et al. 2022) applied to EAGLE as well. In Section 4.4, we consider whether some of the additional parameters identified by the two aforementioned papers could improve our model.

4.1 Comparing input and predicted stellar masses

We now present the results of each of our three models. The surviving coefficients and their respective values are shown in Table 1. In order to extract a fitting function that can be applied directly to the input variables, one first needs to transform the input data using equation (1), which requires the mean and standard deviation values of the dependent variables. The values of these parameters for our combined model are given in Table 2.³

Fig. 3 shows a comparison between the stellar masses predicted by the models for haloes in the holdout set and their actual values in EAGLE. This choice of sample enables the accuracy on the model to be assessed by considering data that was not used in training the model. The left-hand and right-hand panels show the results for central and satellite haloes, respectively. The figure shows that the mean closely follows the one-to-one relation (black dashed line) for all models above $\log_{10}(M^*/M_\odot) \sim 8$. The bottom panels highlight how accurate the models are, with the shaded area corresponding to an estimate of the error on the mean. The latter is computed using the central 68 per cent range of the stellar mass distribution divided by the square root of the number of galaxies in a given stellar mass bin. The mean model stellar mass is predicted to percent level accuracy for all stellar masses of interest and always within our estimate of the error on the mean.

Overall, the plot is encouraging and shows that the properties of satellites, as well as centrals, can be accurately predicted by the SRM approach. This is an important prerequisite for constructing accurate mock catalogues from dark matter simulations. We will explore the performance of the models in more detail below.

³Note that the resulting stellar mass also needs to be converted from standardized units, and we have therefore included the stellar mass parameters in Table 2 as well.

Table 1. Parameters and their respective values for the surviving coefficients of the three models. Note that the parameters presented here are in the standardized space defined by equation (1). Parameters are shown to three significant figures, sufficient to make the RMSE accurate to four significant figures.

Coefficient	Centrals	Satellites	Combined
Constant	0.122	0.172	0.171
lgM_{max}	1.20	1.12	1.17
$(\text{lgM}_{\text{max}})^2$	−0.144	−0.154	−0.146
$(\text{lgM}_{\text{max}})^3$	0.00527	0.00633	0.00509
FC_{20}	0.0435	—	0.0136
FC_{30}	—	—	0.0223
FC_{50}	−0.0732	0.0603	0.0560
FC_{70}	0.0803	0.110	0.0953
FC_{90}	0.0262	0.100	0.190
$(\text{FC}_{30})^2$	—	—	0.0107
$\text{lgM}_{\text{max}} \times \text{FC}_{20}$	−0.0392	—	−0.0224
$\text{lgM}_{\text{max}} \times \text{FC}_{30}$	—	—	−0.00508
$\text{lgM}_{\text{max}} \times \text{FC}_{50}$	—	−0.0595	−0.0263
$\text{FC}_{20} \times \text{FC}_{90}$	—	—	−0.0220
$\text{FC}_{30} \times \text{FC}_{90}$	—	—	−0.0192
$\text{FC}_{50} \times \text{FC}_{90}$	—	−0.0450	−0.0636
$\text{FC}_{70} \times \text{FC}_{90}$	—	−0.0121	—
$(\text{FC}_{20})^3$	0.0106	—	—
$(\text{FC}_{30})^3$	0.00521	—	—
$(\text{lgM}_{\text{max}})^2 \times \text{FC}_{30}$	—	−0.00217	—
$(\text{lgM}_{\text{max}})^2 \times \text{FC}_{90}$	—	−0.00521	−0.0124
$\text{lgM}_{\text{max}} \times (\text{FC}_{20})^2$	—	0.00197	—
$\text{lgM}_{\text{max}} \times (\text{FC}_{90})^2$	—	—	−0.00433
$(\text{FC}_{20})^2 \times \text{FC}_{70}$	—	0.00567	0.00875
$(\text{FC}_{30})^2 \times \text{FC}_{20}$	—	—	−0.00186
$(\text{FC}_{50})^2 \times \text{FC}_{20}$	−0.00158	—	0.0243

A subsidiary aim, however, is to determine whether it was necessary to explicitly distinguish between central and satellite galaxies in constructing the model. We test this by comparing the model in which central and satellite galaxies are fitted separately with one that combines all galaxies into one single model and relies on the methodology to distinguish between satellite and central galaxies only on the basis of their different formation histories. The dashed coloured lines in Fig. 3 show the mean stellar mass of the central (left-hand panel) and satellite (right-hand panel) galaxies in the holdout set when the combined model was used, i.e. a model that is trained on all galaxies simultaneously with no binary distinction between satellites and centrals. Those dashed coloured lines are virtually identical to the models inferred using central and satellite information alone (solid lines).

Removing this binary condition should result in an algorithm that is less dependent on the details of the SUBFIND algorithm, making results simpler to interpret.

In order to compare the accuracy of the models, we use the root mean square error (RMSE) statistic defined as follows:

$$\text{RMSE} = \sqrt{\frac{\sum_{\alpha=1}^N (M_{p\alpha}^*(C) - M_\alpha^*)^2}{N}}, \quad (11)$$

We find the same RMSE of 0.203 for the central galaxies in the holdout set when we predict their stellar mass with either our combined model or the model run with central galaxies only. Similarly, satellite galaxies in the holdout set have an RMSE of 0.236 in the individual model, and an RMSE of 0.243 in the combined model. This shows that a binary distinction between central and

Table 2. Normalization parameters used for the stellar mass and the DM halo variables. These parameters are for the model that mixes central and satellite galaxies. The μ and σ rows correspond to the mean and standard deviation of the variables, respectively, and are used in equation (1) to standardize the range of the variables considered.

	$\log_{10} M^*/M_\odot$	$\lg M_{\max}$	FC'_{20}	FC'_{30}	FC'_{50}	FC'_{70}	FC'_{90}
μ	8.760	11.13	3.054	2.481	1.644	1.034	0.531
σ	0.8002	0.4566	0.8311	0.8786	0.7666	0.6291	0.5206




Figure 3. Comparison between the stellar masses of galaxies in EAGLE and those predicted by the models for all haloes within the holdout set. The coloured shaded areas on the top panels show the boundary encompassing 68 per cent of this holdout galaxies within bins of fixed model SM, and the solid lines are their mean values. The black dashed line corresponds to the one-to-one line. The black horizontal lines show the resolution limit of galaxies within the EAGLE simulation (Schaye et al. 2015). Below this line galaxies are defined by fewer particles and numerical noise starts to become an issue. The left-hand panel shows the result for the central haloes: the solid blue line and light blue shading corresponding to the model trained on centrals alone, while the green dashed line and light green shading to the combined model, trained using centrals and satellites. The right-hand panel is equivalent to the left-hand panel but for satellite galaxies. The bottom panel shows the relative difference between our model prediction and EAGLE data, defined as $(y - x)/x = [\text{EAGLE}(\log_{10}(M^*/M_\odot)) - \text{Model}(\log_{10}(M^*/M_\odot))]/\text{Model}(\log_{10}(M^*/M_\odot))$. It represents the relative difference between the coloured lines and shades and the one-to-one line (black dashed line) shown in the top panel.

satellite galaxies does not improve significantly the accuracy of the models.

We can also look at all centrals and satellites of the individual models used together which have a RMSE of 0.215.

This is very comparable to the RMSE of the combined model which is 0.216.

This indicates that the individual models and the combined model have comparable accuracies. Note that the combined model ends up with 21 terms, while modelling satellites and centrals individually requires 14 and 12 terms, respectively (hence 26 terms in total).

We would like to highlight that none of the three models shows a significant difference between the RMSE of the holdout and training sets at the third significant figure. This suggests that our methodology is robust against overfitting, as overfitting would result in a difference between the RMSE of the holdout and training set. Hence our method of selecting the hyperparameter λ in equation (3), designed to avoid overfitting, works as intended.

In the rest of this work, we present our statistics using the whole data set. This is justified as we have shown that the accuracy of the models is similar for galaxies in the training set and in the

holdout set. The holdout set alone is rather small (about five thousand galaxies typically), and therefore statistics like SMFs or galaxy correlation functions would result with comparatively large statistical uncertainties, if the models are applied to the holdout data only.

A significant appeal of the SRM approach is that the surviving terms in Table 1 have a physical interpretation. Following the discussion in Icaza-Lizaola et al. (2021), we note that there are four types of surviving parameters:

- (i) A constant, or normalization, term.
- (ii) Terms that only include $\lg M_{\max}$ and no formation criteria parameter: these terms model the underlying relation between $M_{\max}$ and M^* . For central galaxies they should correspond to a model of the SMHM relation.
- (iii) Terms that only include formation criteria parameters (e.g. FC_{50} and higher order combinations): these terms quantify the growth history of the halo, capturing scatter in the relation.
- (iv) Terms that are a product of halo mass, $\lg M_{\max}$, and formation criteria parameters: these terms model the dependence of the assembly history on the final halo mass.

Comparing the models, we see that the constant term and the coefficients that depend only on the $\lg M_{\max}$ coefficients are similar between all three models. This reflects the similar underlying shape of the $M_{\max}$ and M^* relation.

In the combined model, central and satellite galaxies are treated on an equal footing and their offset is captured by the more complex dependence on formation time parameters. The combined model needs 21 parameters, which are less free parameters than the combination of the two separate models, which each require 12 and 14 parameters to model centrals and satellites, respectively. One noticeable difference is that the combined model relies on terms of the shape $\mathbf{FC}_i \times \mathbf{FC}_{90}$ which measure the time it takes a halo to evolve into their maximum mass with respect to the time it took them to reach a smaller percentage of that mass.

It is interesting to compare the central galaxy model with the one presented in Icaza-Lizaola et al. (2021). It is important to stress that we do not expect identical models, since we have broadened the range of masses considered and weighted the cost function to emphasize the importance of predicting stellar masses well over the full halo mass range. These changes resulted in a slightly simpler model.

The number of free parameters selected by the algorithm has decreased from 17 to 12. However, a close inspection of the surviving parameters of both models reveals a lot of striking similarities between the two. Many of the surviving terms are similar despite the differences in the definition of the halo mass term and, to some extent, the formation criteria definition (see Section 3.2). Here, we use $\lg M_{\max}$, while it was $\log(M_{200}^c)(z=0)$ in Icaza-Lizaola et al. (2021). Both models have surviving coefficients of similar amplitudes for the *constant*, the $\log(M)^x$ and the $\mathbf{FC}_j^x$ terms (with $x < 3$), with a difference now that $\mathbf{FC}_{20}$ is selected instead of $\mathbf{FC}_{30}$. In summary, the main difference between both models is that the model in Icaza-Lizaola et al. (2021) required more cross terms between the mass and the formation criteria parameters while now we only require one ($\lg M_{\max} \times \mathbf{FC}_{20}$).

One difficulty becomes apparent when comparing the models in greater detail, however. Because of the significant correlation between parameters, models of almost equivalent accuracy and complexity can vary in the final parameters chosen if these parameters are correlated. For example, the current central model includes strong dependencies on terms in $\mathbf{FC}_{20}$, while the model of Icaza-Lizaola et al. (2021) had most terms as function of $\mathbf{FC}_{30}$.

It is difficult to decide on the significance of these differences because of the underlying correlations of $\mathbf{FC}_{20}$ and $\mathbf{FC}_{30}$. As mentioned in Section 3.2, future investigations could consider methods like principal component analysis to transform our input functions into a parameter space where they are uncorrelated. However, this would lose the benefit of having a simple physical interpretation of the input parameters and the resulting model.

4.2 Predicting clustering and the SMF

In this section, we explore the SMF and the clustering of the stellar population generated by applying our model to the haloes in our DM only simulation. We compare our resulting statistics to the ones we get from the stellar population of the full EAGLE hydrodynamical simulation. All of the statistics presented here include all haloes in the DM only simulation, even those that were not matched in Section 3.1. Fig. 4 shows how the SMF of our models, split by galaxy type (total in green, centrals in blue and satellites in red) compares to those from the EAGLE hydrodynamical simulation. The plot shows the SMF of the combined model (solid lines), of the individual models (dotted lines) and of the EAGLE data (shaded area), with the shading




Figure 4. The galaxy SMF of EAGLE, represented as the shaded areas, compared to the galaxy SMF of our models, shown as solid (combined model) and dotted (individual models) lines. The green line corresponds to combined samples of all galaxies, and the red and blue lines to the satellite and central subsets respectively. The shaded region shows the bootstrap error on the EAGLE SMF estimate. The bottom panel shows the relative difference of the model predicted SMFs compared to the EAGLE SMF, with the same line styles and colours as in the top panel. The SMFs are shown to $\log_{10}(M^*/M_{\odot}) = 8.8$, the threshold below which the EAGLE galaxy sample starts to be incomplete due to our halo mass cut.

indicating a bootstrap error estimate to account for sampling effects (Efron 1979). The different model SMFs are all comparable, as they seem to agree all similarly well with the EAGLE SMFs, with the agreement worsening somewhat for masses around $\log_{10}(M^*/M_{\odot}) = 10.5$, as identified already in Icaza-Lizaola et al. (2021). As we suggested in that work, one possible reason behind this disagreement is the stochasticity of certain baryonic processes which might affect the stellar mass, for example the feedback from supermassive black holes (Martizzi et al. 2012; Bower et al. 2017). While this would be a challenging phenomenon to predict using input parameters from a DM only simulation, it should be possible to develop, in a future work, SRM models that estimate both a central value and a stochastic scatter in the predicted quantities.

In what follows we show different predictions of galaxy correlation functions and analyse how do they compare with the original statistics from EAGLE. We emphasize that our model is not tuned to reproduce the clustering of EAGLE. Therefore, any success that we may find is a consequence of correlation functions being preserved when populating the correct haloes with galaxies of a given stellar mass.

Fig. 5 shows the galaxy correlation functions of our models, split by the predicted galaxy stellar mass. The figure also includes the correlation function of galaxies when split by their stellar mass in the EAGLE simulation. As with Fig. 4 we have included an estimate of the error due to sampling effects using the bootstrap method. The correlation function of both models with central and satellites galaxies (green lines) agrees within the errors with the EAGLE correlation function. The same is true for central galaxies (blue lines). On the other hand, satellite galaxies (red lines) are slightly more strongly clustered compared to EAGLE in the lowest stellar mass bin. There seem to be no discrepancies in the correlation functions when satellites and central galaxies are modelled together or separately.




Figure 5. Correlation function of EAGLE galaxies split into different stellar mass bins (as indicated in the title of each panel). The solid (dotted) lines show the correlation function of all galaxies in our combined (individual) models. Like in Fig. 4, the colour coding refers to the galaxy sample type: all, central, and satellite galaxies are in green, blue, and red, respectively. The shaded area corresponds to the correlation function of the corresponding EAGLE galaxies including bootstrap errors. The bottom panels show the relative difference of the model predicted correlation function compared to the EAGLE one, with the same line styles and colours as in the top panels.

This is encouraging as it implies that the binary distinction between central and satellite galaxies becomes unnecessary to model the overall correlation function using our prescription.

One of the advantages of our methodology over standard ML techniques is the fact that our solution is expressed as a simple equation of state with 21 free parameters fitted by the algorithm. This is important as the model coefficients can be modified so that other data sets (different from EAGLE) can be fit. This would be needed when for example one wants to populate DM only simulations with EAGLE informed physical processes to create mocks that mimic observational data set. This could not be achievable by a more complex *black box* model.

4.3 Comparison with other models

4.3.1 Comparison with SHAM

We have stated that we are interested in using our methodology as an alternative for populating haloes in DM only simulations with galaxies. To test if our methodology is adequate, we first need to compare our accuracy to that obtained from standard methods like subhalo abundance matching (SHAM; e.g. Vale & Ostriker 2004; Conroy et al. 2006) that makes a one-to-one matching between haloes and galaxies, based on a property that correlates with the stellar mass. More recent implementations of SHAM add some




Figure 6. Correlation functions of the combined model produced with our SRM method (green solid lines), of the SHAM results presented in Chaves-Montero et al. (2016) (blue solid lines), and of the EAGLE hydrodynamical simulation (red dashed line and shading). The correlation functions are computed for galaxies in the stellar mass bins indicated in the title of each panel. The shading correspond to bootstrap errors. The bottom panels show the relative difference of the model predicted correlation functions compared to the EAGLE ones, with the same line styles and colours as in the top panels.

stochasticity to the methodology to account for the scatter in the correlation (e.g. Behroozi, Conroy & Wechsler 2010; Zentner, Hearin & van den Bosch 2014). Therefore, regular SHAM implementations produce models that depend on only one free parameter and one subhalo property, which makes them simpler than our SRM models that consider six halo properties and fit several free parameters.

In what follows, we compare the correlation function from our combined model to the one presented by Chaves-Montero et al. (2016). They used a SHAM methodology to populate galaxies in the EAGLE simulation by studying the relation between the stellar mass of a galaxy and the maximum circular velocity of a halo once it reaches equilibrium after a merger.

Fig. 6 shows how our correlation functions (blue lines) compare to the ones from Chaves-Montero et al. (2016). The right-hand panel shows that for larger stellar masses, both methods agree with EAGLE within the bootstrap errors, while they provide reasonable accuracy in recovering the correlation function for smaller stellar masses (as shown by the left-hand panel). However, Chaves-Montero et al. (2016) seems to struggle to recover the EAGLE correlation function on the smaller scales. They report differences of 20 per cent to 30 per cent, as confirmed in the bottom left-hand panel of Fig. 6. Our SRM model shows a slight improvement on these smaller scales and agrees better with the EAGLE correlation function. For stellar masses larger than the ones shown in Fig. 6, we continue to agree with the EAGLE simulation within errors.

4.3.2 Comparison with ML tree methods

We have stated that our goal is to develop an explainable ML methodology. However, for this to be of use we need to make sure that the accuracy of our model is comparable to that of more established ML methods. With this in mind, in what follows we compare our model with the ML model presented in Lovell et al. (2022), which

uses ERTs (Geurts et al. 2006) to model galaxy properties from EAGLE halo information. ERT methods are emerging as a popular and highly accurate ML method to model the relations between galaxies and host haloes (e.g. Kamdar, Turk & Brunner 2015; Jo & Kim 2019).

The model of Lovell et al. (2022) is trained using data from the EAGLE and the C-EAGLE simulations (Barnes et al. 2017; Bahé et al. 2017). The latter is a set of zoom-in hydrodynamical simulations of massive galaxy clusters. The calibration of C-EAGLE is slightly different from the standard EAGLE one, with changes in the values of the parameters determining the AGN feedback and the black hole accretion rates. This new parametrization is usually referred to as AGNdt9 (Schaye et al. 2015). The EAGLE data used in Lovell et al. (2022) comes from a smaller box of 50 comoving Mpc that has the same resolution and cosmology as the standard 100 Mpc box, but uses the AGNdt9 parametrization of C-EAGLE.

We decided to compare with the model of Lovell et al. (2022), as it was also constructed using the EAGLE simulation and therefore it shares the same cosmology and resolution and was built using the same algorithm that our data, which makes a direct comparison of the models more straightforward. ML methods that have been trained on other simulations might have differences in the accuracy of the models that could be a consequence of the training data and not of the methodology itself.

Lovell et al. (2022) uses either eight or twelve properties of the host DM haloes to model the stellar properties of galaxies (depending on the specific model). Hence, the number of input parameters they consider is comparable to our work, as we use six halo properties. Their properties include information that parametrize the host halo mass at $z = 0$, like the total mass of the halo M_{FoF} , and properties that are more correlated with the assembly history, like V_{max} or the radius at which V_{max} is reached. On the other hand, our formation criteria parameters contain a more direct parametrization of the assembly history.

The top left-hand panel of Fig. 7 shows how the SMFs from our model and from Lovell et al. (2022) compare to the one from the EAGLE hydrodynamical simulation. The bottom left-hand panel of Fig. 7 shows the relative difference of the model SMFs w.r.t. to the EAGLE 100 Mpc. We note that both models have comparable accuracy, with our SRM model being slightly more accurate for stellar masses between $\sim 10^9$ and $\sim 10^{10} M_\odot$. However, we should emphasize that given the Lovell et al. (2022) model is trained on a combination of C-EAGLE and AGNdt9 data, it is less likely to reproduce the SMF of EAGLE as accurately as a model that is trained solely on EAGLE, like ours.

McAlpine et al. (2016) show that the SMF of the AGNdt9 simulation agrees well with the one from the larger EAGLE 100 Mpc simulation, with both SMFs being identical in all but the larger stellar mass bins where AGNdt9 lacks volume to be representative, which is precisely what the C-EAGLE data used by Lovell et al. (2022) compensates for. Therefore, it is reasonable to compare both models with the SMF of the EAGLE 100 Mpc box, bearing in mind those limitations.

The right-hand panel of Fig. 7 shows the projected correlation function ($w_p(r_p)$)⁴ for a stellar mass selected sample as defined by the panel title. The clustering of the 50 Mpc box built with the AGNdt9 parametrization is slightly different from the one built with the standard 100 Mpc box, as shown by the two dashed lines in the right-hand panel of Fig. 7. The AGNdt9 simulation (along with C-EAGLE data) was used to build Lovell et al. (2022) model and therefore the correlation function of the model applied to the DMO should be compared with the correlation function of the AGNdt9, which is why the ratio of the bottom panel of the right plot is done w.r.t. the EAGLE AGNdt9 50 Mpc box.

The two solid lines in the right-hand panel of Fig. 7 correspond to two different models, as indicated by the key. We note that the projected correlation functions of our SRM model agrees well with the one from EAGLE: on all scales considered, the line from our SRM model is within the bootstrap errors of the EAGLE sample.

Similarly, the clustering of Lovell’s model applied to the DMO simulation agrees well with that of the EAGLE AGNdt9 simulation used to build the model. The accuracy with which this model reproduces the projected correlation function is similar to the one from our model in all but the smallest scales. This is clear from the bottom right panel of Fig. 7 where the relative clustering difference of the Lovell et al. (2022) model with respect to that of the EAGLE AGNdt9 sample is shown by the blue solid line. As the Lovell et al. (2022) model was tuned to a combination of C-EAGLE and AGNdt9 data, it is not straightforward to make a direct comparison with the clustering of their training data, a comparison to the clustering of the AGNdt9 simulation is therefore the best alternative.

We have shown that the projected correlation function and the SMF resulting from our SRM methods are comparable to the ones obtained by Lovell et al. (2022) using ERT methods. As we have stated, the comparison between Lovell et al. (2022) and our model cannot be done fully accurately, as they use data from the C-EAGLE simulation to build their models. Nevertheless, we consider the fact that the models seem to have a similar level of accuracy as an encouraging result, especially as ERT methods are designed to be accurate and cost-efficient (Geurts et al. 2006). Unlike our SRM method, explainability is not an aim within the design philosophy of ERT models.

⁴The projected correlation function (Davis & Peebles 1983) is defined as $w_p(r_p) = 2 \int_{-\infty}^{\infty} \xi(r_p, \pi) d\pi$, where r_p and π are the components of r perpendicular and parallel to the line of sight, respectively.

4.4 Models with additional halo properties

The parametrization of halo properties presented in Section 3.2 is different from the parameters selected by other ML methods. For example, Lovell et al. (2022) uses exclusively properties at $z = 0$ to build a model with an accuracy comparable to ours. Lovell et al. (2022) finds that the maximum circular velocity ($V_{\max}$), the half mass ratio ($R_{1/2}$), the mass of the halo at $z = 0$ (M_0), and the potential energy of the halo (E_p) are the parameters that have significant contributions to their stellar mass model (see fig. 11 of Lovell et al. 2022).

In this section, we explore whether some of these parameters could improve our baseline model, as presented in Section 4. This is not a trivial question, as some of these parameters, like $V_{\max}$ and $R_{1/2}$, might be useful in other ML models as they are a better tracer of the inner part of the halo than $M_{\max}$. However, the halo evolution is already well tracked in our model by our parametrization of the halo evolution with the FC_i parameters.

Another issue when including additional parameters is that some might be strongly correlated with each other. In Appendix C, we show that parameters like M_0 and E_p will provide essentially the same information to our models. Including highly correlated parameters in our current implementation could reduce the explainability of the model, as coefficients corresponding to different polynomial terms of correlated parameters can have different physical interpretations while modeling the same underlying behaviour.

In SRM, the standard approach for dealing with extra variables that one does not know if they could improve a model or not is to add them as free parameters and to see if the algorithm discards them by itself. This is one of the original design philosophies behind these methodologies, as discussed in Brunton et al. (2016). Hence, we run our methodology using our regular six halo parameters to which we add five extra parameters: four free parameters defined at $z = 0$ and suggested by Lovell et al. (2022) [$V_{\max}$, $R_{1/2}$, $\log_{10}(M_0)$ and $\log_{10}(E_p)$] and a fifth parameter V_{peak} , defined in Section 3.2. Throughout the rest of this work, we use the following unitless parameters:

$$\begin{aligned} \lg V_{\text{peak}} &= \log_{10}(V_{\text{peak}}/(\text{km s}^{-1})) \\ \lg V_{\max} &= \log_{10}(V_{\max}/(\text{km s}^{-1})) \\ \lg E_p &= \log_{10}(E_p/(M_\odot(\text{km s}^{-1})^2)) \\ \lg M_0 &= \log_{10}(M_0/M_\odot). \end{aligned} \quad (12)$$

Several of these new halo properties are correlated with each other, as shown in Appendix C. As discussed in section 4.3 of Icaza-Lizaola et al. (2021), correlated parameters have the effect of generating multiple local minima, resulting in a highly non-convex configuration space to be explored. In such spaces, our implementation of the minimization algorithm struggles to find the global minimum, as it spends time exploring unstable local minima. This, in turn, has the net effect of building models with slight variations in the surviving coefficients, which depend on the starting point of the minimization and on the specific selection of galaxies in the holdout set. To address this limitation, we run our methodology five times, using the same initial set of galaxies, but modifying the random seed so that the subset of galaxies selected for the holdout set and the starting points for the minimization algorithm change. These five runs provide an idea of the average model that can be built with this new configuration. Finally, by adding five new parameters to the model, the number of coefficients to minimize over goes from 84 to 364, which increases significantly the dimensionality of the problem.

These three observations, i.e. correlated parameters, the need for statistically equivalent runs and the larger dimensionality of the problem, have the net effect of increasing significantly the




Figure 7. Left-hand panel: the SMF predicted by the ERT method of Lovell et al. (2022) (blue line) and our SRM model (green line) when applied to haloes within the EAGLE DMO simulation box. Right-hand panel: The projected correlation functions predicted by Lovell et al. (2022) (blue solid line) and by our SRM model (green solid line) when applied to haloes within the EAGLE 100 Mpc DMO simulation, using the stellar mass bin of Lovell et al. (2022). In both top panels, the red dashed lines (and shading) show the corresponding statistics (and bootstrap errors) measured directly from the EAGLE hydrodynamical simulation. The blue dashed line of the right-hand panel shows the projected correlation function of the EAGLE simulation built with the AGNdt9 parametrization. The bottom left-hand panel shows the relative difference of the model predicted SMF w.r.t. the EAGLE 100 Mpc box, with the same line styles and colours as in the corresponding top panels. The green line in the bottom right-hand panel shows the relative difference of the projected correlation functions of our SRM model w.r.t. the one from the EAGLE 100 Mpc box, while the blue line is the relative difference of the projected correlation function from Lovell et al. (2022) model w.r.t. the one from the EAGLE AGNdt9 50 Mpc box. We note that Lovell et al. (2022) model is trained on data from both the C-EAGLE and EAGLE AGNdt9 simulations, and therefore, it is less likely to reproduce the SMF of EAGLE as accurately as our model that is trained solely on EAGLE. See the main text for a more detailed discussion.

computational cost of running our algorithm. As a consequence, we make the compromise of using only 4000 randomly selected galaxies to run our models on (as opposed to the nominal 35 456 galaxies), as this keeps the overall computational running costs manageable. In parallel, we run another set of five models that uses our standard configuration of 6 parameters from Section 4, but built with the same 4000 galaxies and random seeds as these new models. These five models correspond to our baseline models throughout this subsection, and we refer to them as our 6 parameter models. In the rest of this section these models are contrasted with their equivalent models built with extra parameters but the same random seeds, to which we refer to as our 6 + 5 parameter models.

Before analysing the subset of parameters selected by the algorithm for the new 6+5 parameters models, we show that these models are as accurate as the models run with the nominal 6 parameter configuration. All five 6 + 5 parameters models have a RMSE between 0.22 and 0.23, which is comparable to within the uncertainty of the model fitting to the RMSE of the corresponding nominal 6 parameter models, which is between 0.21 and 0.22. These values are also comparable with our final model from Section 4, that has an RMSE of 0.22 when estimated with the set of 4000 galaxies used in this section. We note that all five runs of the 6 + 5 parameter models choose a similar number of surviving coefficients, with two runs selecting 13 and 15 coefficients each, while the other three runs all selecting 10. This is in agreement with the variance on the methodology due to variations in the holdout set selection found in Icaza-Lizaola et al. (2021). We find no correlation

between the number of surviving coefficients and the RMSE of the models.

Fig. 8 shows the values of each of the selected coefficients for our five new models using both our new configuration with 6 + 5 parameters (left) and the standard configuration with 6 parameters (right).

Via the SRM methodology, most models will have a subset of their allowed parameters discarded (when none of the coefficients associated with these parameters are chosen by the algorithm). In our case, the five runs of the 6 + 5 parameter model end up keeping between 5 and 9 parameters. The differences in the number of surviving parameters between the runs show how correlated these parameters are with each other, making them somewhat interchangeable. As shown in the left-hand panel of Fig. 8, the only new parameter selected by all five runs is $\lg V_{\text{peak}}$, while four out of the 6 standard model parameters ($\lg M_{\text{max}}$, FC_{30} , FC_{70} , and FC_{90}) are kept in each resulting new model.

We note that all of the 6 parameter models keep all input parameters without discarding any. This suggests that the information contained inside the parameters used in our standard configuration is more unique than the one from 6 + 5 parameter models used in this section. This is further discussed in Appendix C, where we show how most of the new parameters included are strongly correlated with each other and with $\lg M_{\text{max}}$, while the correlations between the formation criteria parameters are comparably weaker and hence contain more specific information. The fact that the models that start with 11 free parameters require in some cases a large number of parameters to




Figure 8. The left-hand panel shows the absolute values of all selected coefficients for each of the five statistically equivalent 6 + 5 parameter models (RS-1 to RS-5) trained and validated on a sample of 4000 galaxies. The right-hand panel shows similar information using the same data set, but for our standard 6 parameter model. The coloured labels at the right of each plot correspond to parameters that were used in at least three of the five models, while the labels inside the plot correspond to parameters used only in one or two of the models. The grey shading shows the threshold below which coefficients are discarded by a given model.

reach an accuracy similar to the model in Section 4 suggests that the new parameters did not contain much additional (if any) new information that was not already present in our initial model.

The five runs with our standard 6 parameter model selected between 13 and 18 coefficients, less than the 21 coefficients of our final model from Section 4. The differences in selected coefficients is due to these new models being built with less data. Given that the SRM method is very strict about avoiding over fitting, it becomes harder to justify a larger set of coefficients. As a test of this, we ran another set of five models using our standard 6 parameters but using 12 000 galaxies (3 times more than the data in this section and 3 times less than the nominal set) and we found that the selected models use between 16 and 19 coefficients.

From the left-hand panel of Fig. 8, we note that the five 6 + 5 parameter models use between 10 and 15 coefficients, which is slightly less than the 13–18 coefficients of the 6 parameter models. This suggests that, while the new parameters might not necessarily contain much new information required to model the SMHM relation, they might be more efficient at compressing the relevant information.

While the 6 + 5 parameter models select on average less coefficients, it is significantly less consistent in the subset of coefficients selected by any particular model. This is shown by the fact that only 9 coefficients are selected in at least three models, while another 14 coefficients are selected once or twice only. This is in contrast with the 6 parameters models, where most coefficients are present in at least three models and only 6 out of 21 coefficients are selected once or twice. In fact, we note that even if the 6 + 5 parameter model uses fewer coefficients per model, the number of coefficients selected by at least one model (23 coefficients) is comparable to that from the 6 parameter model (21 coefficients). This shows how the inclusion of correlated parameters increases the stochasticity of the method, which in turn complicates the interpretation of the resulting models.

Out of the 11 parameters, all 6 + 5 parameter models select linear contributions from $\lg M_{\max}$, $\lg V_{\text{peak}}$, FC_{70} and FC_{90} and cubic

contributions from $\lg V_{\text{peak}}$ and FC_{30} , in order of decreasing linear coefficient value. Almost all models select some contribution from $R_{1/2}$ and FC_{50} as well, except for one model, RS-4 in the left-hand panel of Fig. 8. That latter model has a comparable RMSE to the other models and requires 10 coefficients, the same number as two other models.

This highlights the difficulties of using our current SRM implementation on spaces with highly correlated parameters: RS-4 has an accuracy (RMSE) and simplicity (number of parameters) that are equivalent to those of the other four models (within the variations of the methodology due to different holdout sets). Therefore, it is neither better nor worse than the other models within the standards that we designed our models to meet. However, due to parameters being strongly correlated and sharing similar information, we see that this model requires two less free parameters and therefore would have a simpler physical interpretation than the others.

As shown in Section C, out of all new input parameters, $R_{1/2}$ is the one that is the least correlated with the rest of the new parameters. This suggests that the $R_{1/2}$ information provided to the model is possibly more unique than that from some of the other new parameters, which might explain why four new models had some contributions from the $R_{1/2}$ parameter.

We note that all models discard contributions from $\lg V_{\max}$ and $\lg E_p$ up to order 3, and that only one model (RS-2 in the left-hand panel of Fig. 8) includes a very minor contribution from $\lg M_0$ in the form of the coefficient $\lg M_0^2 R_{1/2}$. This suggests that the contribution to the accuracy of the model after including any of these three parameters is negligible and that none of these parameters contributed additional information that was not already provided without them.

FC_{20} is only selected by two models, RS-1 and RS-2 in the left-hand panel of Fig. 8. Those models are the two that have the largest number of coefficients, which suggest that the information that was previously provided by FC_{20} to our model is also contained in some of the new set of parameters.




Figure 9. The SMF as predicted by the five new models (blue lines) that are built by adding five new parameters to our method. The purple lines show the SMF of their corresponding SRM models built with our standard configuration of 6 parameters. We also include the SMF predicted by our final model from Section 4 (green line) when applied to the same subset of 4000 galaxies. The EAGLE SMF of this subset is shown as the red line. The bottom panel shows the ratio of each model predicted SMF to the EAGLE SMF, with all models predicting the SMF to a similar level of accuracy.

In summary and as stated already, all five runs of the 6 + 5 parameter model selected contributions from $\lg M_{\max}$, $\lg V_{\text{peak}}$, FC_{70} , FC_{30} , FC_{90} . Of these parameters, $\lg V_{\text{peak}}$ is the only one that is not within our original set of parameters. Given that the five new models do not seem to be more accurate than the model presented in Section 4, the contribution provided by $\lg V_{\text{peak}}$ could also be obtained by a combination of the FC_i parameters within our original model, as shown in Appendix B. However, the fact that these new models require in general less coefficients seems to indicate that including $\lg V_{\text{peak}}$ is an efficient way of compressing some of the information contained in our FC_i parameters.

Fig. 9 shows the SMFs of the five 6 + 5 parameter models (blue) and the complementary 6 parameter models (purple). These SMFs are built using a subset of 4000 galaxies and we account for this sampling in the SMF estimates. The bottom panel indicates that all models have a similar accuracy (to within 25 per cent typically) when predicting the SMF of EAGLE. This suggests that including the extra parameters does not improve our ability to reproduce the stellar mass distribution. In addition, while the corresponding 6 parameter models do not have better accuracy than their 11 parameter counterparts, they seem to be far more consistent with each other. This can be seen by the purple lines being more similar to each other than the blue ones in the bottom panel of Fig. 9. This is due to the 6 parameter models being more consistent in their selection of surviving coefficients, as shown in the right-hand panel of Fig. 8.

As mentioned already, these models are trained with a small subset of the full data (4000 galaxies as opposed to 35 456 galaxies). Given that our model from Section 4 is trained using our full data set, we could expect its stellar mass predictions to be less accurate for this smaller subset of data, as it was constrained to model a larger data

set. However, we see that both the RMSE and the SMF of the new models are comparable to the one from our final model in Section 4. The fact that our original model seems to do as well as these new ones suggest that our method is robust against sample size variations, and that it is able to deal effectively with overfitting.

Given that we see no improvement in accuracy using these new models, and given that they were not trained on our full data, we do not quote these new models as our final result, but keep the model of Section 4.

5 CONCLUSIONS

In Icaza-Lizaola et al. (2021), we used a sparse regression methodology to fit the stellar mass of central galaxies as a function of properties of their host halo. In this paper, we expand our study to cover a wider halo mass range and to model the properties of satellite galaxies. The distinction between central and satellite galaxies relies on identifying subhaloes as self-bound substructures within larger haloes, for example by using the SUBFIND algorithm. This classification is uncertain and may be inconsistent for the same subhaloes in adjacent snapshots outputs. We therefore explored whether we need to make a fundamental distinction between haloes and subhaloes. With this in mind, we use the maximum mass that a halo has ever reached during its evolution, denoted $\text{Max}(M_{\text{total}}(z))$ and use this in place of the final (sub)halo mass at $z = 0$. Given that central galaxies grow monotonically then $\text{Max}(M_{\text{total}}(z)) \sim M(z = 0)$ and this results in little change. In subhaloes, however, it corresponds to the mass of their main progenitor before merging with their central halo. In order to quantify the prior growth history of the halo, we define a set of formation criteria parameters, that measure the redshift at which a halo has formed a given percentage of its maximal mass and before it reaches $\text{Max}(M_{\text{total}}(z))$.

Our data is taken from the EAGLE hydrodynamical simulation. In order to avoid selection biases when predicting stellar mass, we use a bijective matching between the EAGLE hydrodynamical simulation and a DM only simulation with the same cosmology and initial conditions. We select all galaxies that have a halo mass larger than $\text{Max}(M_{\text{total}}(z)/M_{\odot}) > 10^{10.66}$, this value corresponds to the threshold at which our matching methodology successfully matches more than 90 per cent of all galaxies. We use a total of 35 456 galaxies, 9967 of them live inside subhaloes and 25 489 inside central haloes. Because our sample has significantly increased the fraction of low-mass galaxies considered compared to our previous work (Icaza-Lizaola et al. 2021), we weight residuals according to stellar mass, giving a larger incentive to the model to accurately fit less well represented galaxy masses.

We build our models only using information on the accretion history of the halo or subhalo and its maximum mass. Using these parameters our methodology seem to predict the stellar mass of galaxies in haloes and subhaloes with a singular model and without needing to distinguish between the two. We note that there are other parameters that we have not tested for in our analysis that might break this symmetry, for example, the infall angle of subhaloes, which is not defined for central haloes, might improve our modelling of subhaloes.

The SMF of our models agrees well with that of EAGLE at all stellar masses except at $\log_{10}(M^*/M_{\odot}) = 10.5$ where our models tend to slightly underpredict the amount of galaxies when compared with the EAGLE simulation. This could be related to the stochasticity of baryonic processes that might alter the stellar mass of a galaxy, which could be hard to predict using parameters from a DM only simulation. We also calculate the correlation functions of our models

split by their predicted stellar mass, and find that they also agree well with the EAGLE correlation functions. The model that combines central and satellite galaxies has comparable accuracy to the models in which central and satellites are treated independently, while using an overall smaller number of model parameters. This suggests that a binary classification is unnecessary and the stellar mass of both galaxy types can be predicted by suitable measurement of their halo mass history.

The SRM approach can be viewed as an ML algorithm. It can accurately model the stellar masses of EAGLE from the data itself and without requiring previous knowledge of physics behind the system. At the same time, the approach results in a prediction algorithm that is explicit and simple (compared with the solutions of other ML techniques), and the terms that are retained give physical insight into the important processes at work.

We have seen that the correlation function and the SMF of our models agree well with the EAGLE data set. This is encouraging as both of these EAGLE statistics have been positively compared with observational data. For example, Furlong et al. (2015) has shown that the EAGLE SMF at $z = 0$ agrees reasonably well with the ones observed by the SDSS (Li & White 2009) and GAMA (Baldry et al. 2012) surveys. Similarly, Artale et al. (2017) shows that the EAGLE correlation function reproduces observations accurately between $1 h^{-1}$ Mpc and $6 h^{-1}$ Mpc. Additional statistics, like Counts-in-Cells and multipoles of the correlation function, were successfully reproduced by the models, but we leave to future work a more in depth discussion of their successes and limitations.

Our method compares favourably with the SHAM methodology from Chaves-Montero et al. (2016), with both models being able to reproduce well the correlation function of EAGLE at larger stellar masses with our SRM models being slightly more accurate on smaller scales.

We also compare our model with the one presented in Lovell et al. (2022), using ERT which is a highly accurate ML methodology. ERT makes accurate models but the resulting models are less explainable than our SRM models. Both methods reach comparable accuracy on the SMF predictions, with our model being slightly more accurate at smaller stellar masses. We find similar predictions for the projected correlation function of a stellar mass selected sample between both models. We note that Lovell et al. (2022) data were trained using C-EAGLE zoom-in simulation data that is not identical to the EAGLE data used in training our model, which might explain some of the small differences seen in the accuracy of the predictions of both models.

Finally, we analyse the inclusion of additional halo properties into our methodology. This is done by building new models with some of the halo parameters used in other successful ML models. We run five new models to account for differences due to variance in our methodology, which increases due to the correlations between the new parameters. We find no improvement in accuracy which suggests that any information provided by the new parameters was already present in our standard parametrization. We find a slight reduction in the number of surviving coefficients, which suggests that some parameters, like $\lg V_{\text{peak}}$ and $R_{1/2}$, are possibly more efficient at summarizing some of the relevant information required to described the SMHM relation. However, the number of free parameters varies between five and nine depending on the model realization, which complicates significantly the model interpretation, one of the underlying aims of this SRM methodology. Due to this fact along with the reduced stability of the model as evidenced by the increase in scatter on the predicted SMFs (Fig. 9), we do not quote these new models as final result.

All of this suggests that our methodology could be a promising approach to populate N -body simulations with galaxies of the correct stellar mass and spatial distribution.

However, several complications will make this an interesting challenge. First, EAGLE is run in a comparatively small volume with respect to other DM simulations that means that the number of massive haloes is comparatively small and it will be necessary to test the accuracy of the resulting SMF at the larger stellar masses. Secondly, larger simulations normally produce large amounts of output data, which generates challenges in storing the necessary halo history to build merger trees, some simulations either save only a small number of redshifts or no halo evolution information at all. Finally, the distribution of our required input halo parameters such as $\lg M_{\text{max}}$ or FC_i might differ from simulation to simulation. All of these reasons make populating larger simulations with galaxies using our methodology a challenging endeavour that we will explore in more depth in future papers.

Our ultimate goal is to generate mock catalogues that provide an accurate representation of the observed Universe. An attractive idea is to iterate on the coefficients of the terms selected by comparison to EAGLE (or another hydrodynamic simulation), creating an even closer match to target observations. This would retain the same physical processes, but accept that their relative importance might differ between the true Universe and the simulation used for the training.

ACKNOWLEDGEMENTS

We thank the anonymous referee for their very insightful comments that improved the paper. We thank Arnau Quera-Bofarull for his help in making our code faster and more efficient. We thank Christopher Lovell and Jonás Chaves-Montero for sharing their data with us, enabling the detailed comparisons presented. MIL is supported by a PhD Studentship from the Durham Centre for Doctoral Training in Data Intensive Science, funded by the UK Science and Technology Facilities Council (STFC, ST/P006744/1) and Durham University. MIL, RGB, PN, and SMC acknowledge support from the Science and Technology Facilities Council (ST/P000541/1 and ST/T000244/1). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, and ST/R002371/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure. Some of the numerical calculations in this work were done in the high performance computer cluster of the Korea Astronomy and Space Science Institute.

DATA AVAILABILITY

The data used in this work can be shared if requested from the authors. The data from the EAGLE simulations has been publicly released (see McAlpine et al. 2016).

REFERENCES

- Agarwal S., Davé R., Bassett B. A., 2018, *MNRAS*, 478, 3410
- Artale M. C. et al., 2017, *MNRAS*, 470, 1771
- Bahé Y. M. et al., 2017, *MNRAS*, 470, 4186
- Bahé Y. M., McCarthy I. G., 2015, *MNRAS*, 447, 969
- Baldry I. K. et al., 2012, *MNRAS*, 421, 621
- Barnes D. J. et al., 2017, *MNRAS*, 471, 1088

Baugh C. M. et al., 2018, *MNRAS*, 483, 4922
 Behroozi P. et al., 2015, *MNRAS*, 454, 3020
 Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
 Behroozi P. S., Conroy C., Wechsler R. H., 2010, *ApJ*, 717, 379
 Bower R. G., Balogh M. L., 2004, in Mulchaey J. S., Dressler A., Oemler A., eds, *Clusters of Galaxies: Probes of Cosmological Structure and Galaxy Evolution*. Cambridge University Press, Cambridge, U.K., p. 325
 Bower R. G., Schaye J., Frenk C. S., Theuns T., Schaller M., Crain R. A., McAlpine S., 2017, *MNRAS*, 465, 32
 Brunton S. L., Proctor J. L., Kutz J. N., 2016, *Proc. Natl. Acad. Sci.*, 113, 3932
 Bryan S. E., Kay S. T., Duffy A. R., Schaye J., Dalla Vecchia C., Booth C. M., 2013, *MNRAS*, 429, 3316
 Chaves-Montero J., Angulo R. E., Schaye J., Schaller M., Crain R. A., Furlong M., Theuns T., 2016, *MNRAS*, 460, 3100
 Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
 Correa C. A., Schaye J., Trayford J. W., 2019, *MNRAS*, 484, 4401
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Davis M., Peebles P. J. E., 1983, *Machine Learning*, 267, 465
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
 Efron B., 1979, *Ann. Statist.*, 7, 1
 Furlong M. et al., 2015, *MNRAS*, 450, 4486
 Geurts P., Ernst D., Wehenkel L., 2006, *Machine Learning*, 63, 3
 Green S. B., van den Bosch F. C., 2019, *MNRAS*, 490, 2091
 Gunn J. E., Gott J. R. I., 1972, *ApJ*, 176, 1
 Hastie T., Tibshirani R., Wainwright M., 2015, *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman & Hall/CRC, New York
 Hayashi E., Navarro J. F., Taylor J. E., Stadel J., Quinn T., 2003, *ApJ*, 584, 541
 Icaza-Lizaola M., Bower R. G., Norberg P., Cole S., Schaller M., Egan S., 2021, *MNRAS*, 507, 4584
 Jo Y., Kim J.-h., 2019, *MNRAS*, 489, 3565
 Jolliffe I., 2005, *Principal Component Analysis*. American Cancer Society (<https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470013192.bsa501>)
 Kamdar H. M., Turk M. J., Brunner R. J., 2015, *MNRAS*, 455, 642
 Kamdar H. M., Turk M. J., Brunner R. J., 2016, *MNRAS*, 457, 1162
 Katz N., Gunn J. E., 1991, *ApJ*, 377, 365
 Larson R. B., Tinsley B. M., Caldwell C. N., 1980, *ApJ*, 237, 692
 Li C., White S. D. M., 2009, *MNRAS*, 398, 2177
 Lovell C. C., Wilkins S. M., Thomas P. A., Schaller M., Baugh C. M., Fabbian G., Bahé Y., 2022, *MNRAS*, 509, 5046
 Lynden-Bell D., 1967, *MNRAS*, 136, 101
 Martizzi D., Teyssier R., Moore B., Wentz T., 2012, *MNRAS*, 422, 3081
 Matthee J., Schaye J., Crain R. A., Schaller M., Bower R., Theuns T., 2017, *MNRAS*, 465, 2381
 McAlpine S. et al., 2016, *Astron. Comput.*, 15, 72
 Merritt D., 1983, *ApJ*, 264, 24
 Moster B. P., Naab T., Lindström M., O'Leary J. A., 2021, *MNRAS*, 507, 2115
 Navarro J. F., Eke V. R., Frenk C. S., 1996, *MNRAS*, 283, L72
 Planck Collaboration I, 2014, *A&A*, 571, A1
 Qu Y. et al., 2017, *MNRAS*, 464, 1659
 Schaller M. et al., 2015a, *MNRAS*, 451, 1247
 Schaller M. et al., 2015b, *MNRAS*, 452, 343
 Schaye J. et al., 2015, *MNRAS*, 446, 521
 Springel V. et al., 2005, *Nature*, 435, 629
 Springel V. et al., 2018, *MNRAS*, 475, 676
 Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
 Tibshirani R., 1996, *J. R. Stat. Soc. B (Methodol.)*, 58, 267
 Tibshirani R., Friedman J., 2020, *Journal of Computational and Graphical Statistics*, 29, 215
 Vale A., Ostriker J. P., 2004, *MNRAS*, 353, 189
 van den Bosch F. C., Ogiya G., Hahn O., Burkert A., 2018, *MNRAS*, 474, 3043

Vollmer B., Cayatte V., Balkowski C., Duschl W. J., 2001, *ApJ*, 561, 708
 Zentner A. R., Hearin A. P., van den Bosch F. C., 2014, *MNRAS*, 443, 3044

APPENDIX A: MATCHING FAILURES

As mentioned in Section 3.1, the matching success rate of satellite galaxies is around 80 per cent. With this in mind, we decided to apply our model to all haloes in the DM only simulation (match and unmatched) and compare the resulting statistics to the ones obtained from all galaxies in the EAGLE hydrodynamical simulation.

Fig. A1 shows the success rate of the matching algorithm as a function of the halo mass M_{total} at $z = 0$ for both the EAGLE-DMO simulation (red line) and the EAGLE hydrodynamical simulation (green line). For haloes larger than $\log_{10}(M_{\text{total}}/M_{\odot}) = 11$ the percentage of unmatched haloes is small (< 2 per cent) and similar across both simulations. This suggests that most of the haloes that were unmatched in the hydrodynamical simulation did have an equivalent halo in the DM only simulation, but the algorithm had trouble matching them. This justifies the decision made in Section 3.1 to compare models applied to all haloes (match and unmatched) in the DM only simulation to all haloes (match and unmatched) in the hydrodynamical simulation.

The matching algorithm runs at $z = 0$ and therefore the lines at this redshift are adequate to show the success rate of the algorithm. However, we select our halo sample using M_{max} , which is the maximum mass reached by the halo at any redshift. The matching success rate as a function of M_{max} is shown as the blue line in Fig. A1. The matching success rate as function of M_{max} is smaller than when considering M_{total} at $z = 0$. The success rate is around 88 per cent at $\log_{10}(M_{\text{max}}/M_{\odot}) = 10.66$ (our mas cut), and grows to around 94 per cent at $\log_{10}(M_{\text{max}}/M_{\odot}) = 11.5$. These differences in success rate are due to a significant fraction of haloes being disrupted after a merger. These disrupted haloes would be smaller at $z = 0$ than at the redshift of their maximum mass, and therefore their probability of being matched decreases. This is consistent with other works that




Figure A1. The success rate of the matching methodology for haloes in the hydrodynamical simulation (green line) and in the DM only simulation (red and blue lines) as a function of halo mass. For the green and red lines, the halo mass is M_{total} at $z = 0$, while for the blue line it is M_{max} , the mass parameter used by our SRM model. The coloured shadings show the error on the matching rate assuming binomial statistics.




Figure B1. RMSE reached by our algorithm at different values of the hyperparameter Λ (equation 3) for the test and training sets of the k -fold method. The RMSE from the combined model of Table 1 with the free parameter $\lg M_{\max}$ are shown with green and purple lines, while those from a new model with $\lg V_{\text{peak}}$ as a free parameter are shown with blue and red lines. The thin lines represent the RMSE for each of the $k = 10$ individual data sets of the k -fold method and the thick lines show their mean value. The vertical black dashed line correspond to the Λ value for which the model accuracy, described by the RMSE, is the same for both set of models.

have found that an accurate measurement of $M_{\max}$ requires higher numerical resolution.

APPENDIX B: COMPARING WITH V_{PEAK}

In this section, we explore which one of the two halo properties V_{peak} and $M_{\max}$ would be a better input for our models. As stated in Section 3.2, the consensus is that stellar mass models that use properties correlated with the circular velocity profile of haloes, like $V_{\max}$ and V_{peak} , tend to outperform those based on the mass of the halo. This is due to $V_{\max}$ being a good representation of the inner part of the halo, which affects galaxies more directly and is less sensitive to mass stripping. However, we note that the evolutionary history of the halo is well tracked in our SRM model due to our definition of $M_{\max}$ and the inclusion of formation criteria parameters. Therefore, it is not trivial to know which of the two properties will perform better in our model.

We run our combined model from Table 1, but substituting $\lg M_{\max}$ for the unitless parameter $\lg V_{\text{peak}}$ defined in equation (12).

As mentioned in Section 2 the optimal value of the hyperparameter Λ from equation (3) is found using a k -fold method, where the data are separated into a training set and a test set k -times. We examine how well a model fitted to the training sets at different values of Λ predicts the test sets. We refer the reader to Icaza-Lizaola et al. (2021) for an in-depth discussion of this process. Fig. B1 shows the RMSE resulting from the exploration of the Λ space for both models and the training and test sets of all k -folds. The figure shows how models that use $\lg M_{\max}$ as a parameter are more accurate than those using $\lg V_{\text{peak}}$ for both training and test sets. We note that in this comparison the same set of formation criteria parameters were considered by both

set of models and it is within this specific modelling context that we draw our conclusions.

The models that use $\lg V_{\text{peak}}$ are less accurate, but they are simpler than the one with $\lg M_{\max}$, as the former require only 6 parameters. We can build a simpler $\lg M_{\max}$ model by increasing the magnitude of Λ beyond its nominal optimal value. The black dashed line in Fig. B1 shows the value of Λ at which a model built with $\lg M_{\max}$ reaches the same accuracy as the one with $\lg V_{\text{peak}}$. The resulting $\lg M_{\max}$ model built with this Λ contains seven free parameters, which is very comparable with the six of the $\lg V_{\text{peak}}$ model. With this in mind, we conclude that models built with $\lg M_{\max}$ are more accurate and can be as simple as models built with $\lg V_{\text{peak}}$. This justifies our selection of $\lg M_{\max}$ as the mass parameter used in this work.

APPENDIX C: CORRELATED PARAMETERS

Fig. C1 shows the correlations of most halo properties used throughout this work, built from the 4000 haloes considered in Section 4.4. For clarity, the parameters FC_{30} and FC_{70} have been omitted, as they show similar correlation trends to the other three formation criteria parameters already included. Each panel includes P_r , the value of the Pearson correlation coefficient⁵ for each pair of halo properties. The closer the absolute value of this coefficient is to unity, the more linearly correlated those two parameters are.

Fig. C1 shows that parameters can be divided into two subgroups of correlated halo properties:

- (i) The first subgroup includes $\lg M_0$, $\lg M_{\max}$, $\lg V_{\text{peak}}$, $\lg V_{\max}$, $R_{1/2}$, and $\lg E_p$. They are all strongly correlated with each other, with $|P_r|$ around than 0.9 typically.
- (ii) The formation criteria parameters FC_i form the second subgroup. Their correlations, as measured by the Pearson coefficient, are weaker than those within the first group, with $|P_r|$ less than 0.7 typically.

Out of all of the parameters in the first group, $R_{1/2}$ is the least correlated with the rest, with Pearson coefficients between 0.61 and 0.88 with respect to the rest of the halo properties of this subgroup. This might explains why most of the five models of Section 4.4 select a small but noticeable contribution from $R_{1/2}$ after already having strong contributions from $\lg M_{\max}$ and $\lg V_{\text{peak}}$. It is also noticeable how correlated $\lg M_0$ and $\lg E_p$ are with each other, with a correlation coefficient of 0.99. This suggests that the information that they could provide to a model is almost identical. The strong correlation between both parameters comes from the similarities in the way they are defined and computed in EAGLE (see appendix D in McAlpine et al. 2016).

The fact that the parameters of this first subgroup are so correlated with each other might explain why the model of Section 4.1, which is built using only $\lg M_{\max}$ has comparable accuracy to the models of Section 4.4, which are made using all of the 6 parameters of this subgroup.

The weaker correlations observed between parameters in the second group, i.e. the formation criteria parameters FC_i , suggest that the information held by those halo properties is somewhat more unique, especially when compared to the halo properties of the first subgroup. This might explain why all of the models

⁵The Pearson correlation coefficient is defined as the ratio of the covariance of the parameters with the product of their standard deviations.




Figure C1. Correlations of most halo properties used in this work, as indicated by the axis labels of each panel. The Pearson correlation coefficient is indicated in each panel, as a measure of how correlated two halo properties are.

presented in this work select several parameters of this subgroup simultaneously.

As we discuss in detail in Section 4.4, the inclusion of correlated parameters adds stochasticity to our resulting models. This can be seen in models selecting very different collections of surviving coefficients when built with different subsets of training data. As mentioned this is due to correlated parameters making the parameter

space non-convex, with several local minima. Dealing with correlated parameters is something that would need to be implemented into our methodology in future work, if uniqueness of the solution and maximal parameter reduction is a priority.

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