

# Dark from light (DfL): Inferring halo properties from luminous tracers with machine learning trained on cosmological simulations

## I. Method, proof of concept, and preliminary testing

Asa F. L. Bluck<sup>1,\*</sup>, Joanna M. Piotrowska<sup>2</sup>, Paul Goubert<sup>1</sup>, Roberto Maiolino<sup>3,4,5</sup>, Camilo Casimiro<sup>1</sup>, Thomas Pinto Franco<sup>1</sup>, and Nicolas Cea<sup>1</sup>

<sup>1</sup> Stocker AstroScience Center, Department of Physics, Florida International University, 11200 SW 8th Street, Miami, FL, USA

<sup>2</sup> Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA, USA

<sup>3</sup> Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

<sup>4</sup> Cavendish Laboratory Astrophysics Group, University of Cambridge, 19 JJ Thomson Avenue, Cambridge, CB3 0HE, UK

<sup>5</sup> Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

Received 22 March 2025 / Accepted 28 June 2025

### ABSTRACT

We present Dark from Light (DfL), a novel method for inferring the dark sector in wide-field galaxy surveys that leverages a machine learning approach trained on contemporary cosmological simulations. The aim of this algorithm is to provide a fast, straightforward, and accurate route to estimating dark matter halo masses and group membership in wide-field spectroscopic galaxy surveys. This approach requires a highly limited number of input parameters (RA, Dec,  $z$ , and stellar mass plus uncertainties) and yields full probability distribution functions for the output halo masses of galaxies, groups, and clusters. To achieve this, we trained a series of random forest (RF) regression models on the IllustrisTNG and EAGLE simulations at  $z = 0-3$ , which provided model-dependent mappings from luminous tracers to dark matter halo properties. We incorporated the individual regression models into a virial group-finding algorithm (DfL), which outputs halo properties for observational input data. We tested the method at  $z = 0-2$  for both the EAGLE and IllustrisTNG models, as well as in a cross-validation mode (where one simulation is used to train the model and the other to test). We demonstrate that known halo masses can be recovered with a mean systematic bias of  $\langle b \rangle = \pm 0.10$  dex (resulting from simulation choice), a mean statistical uncertainty of  $\langle \sigma \rangle = 0.12$  dex across epochs, and a central – (core) satellite classification accuracy of 96%. We establish that this approach yields halo mass recovery that is superior to standard abundance matching applied to groups identified through a friends-of-friends algorithm. Additionally, we compare the outputs of DfL to observational constraints on the  $M_* - M_{\text{Halo}}$  relation from strong gravitational lensing at  $z \sim 0$ , demonstrating the promise of this novel approach. Finally, we systematically quantify how DfL performs on observational-like input data with varying stellar mass uncertainty and spectroscopic incompleteness, enabling robust error calibration in applications with observational galaxy surveys.

**Key words.** galaxies: abundances – galaxies: evolution – galaxies: formation – galaxies: statistics – dark matter

## 1. Introduction

Dark matter remains a mystery in large part because, whatever its ultimate physical nature turns out to be, it does not interact with the electromagnetic field (e.g., Frenk et al. 1988; Efstathiou et al. 1990; Navarro et al. 1996). Even with the advent of gravitational wave and neutrino astronomy (e.g., Mannheim et al. 2000; Abbott et al. 2016, 2017; Ansoldi et al. 2018), the electromagnetic spectrum provides the vast majority of information we have on the Universe. Therefore, looking for ways to infer the presence and properties of dark matter using luminous tracers has a long history (e.g., Zwicky 1933; Refsdal 1964; Sunyaev & Zeldovich 1972; Rubin et al. 1978; Walsh 1979; Hu & Sugiyama 1995; Warren et al. 2006; Xue et al. 2008; Cappellari et al. 2006, 2013).

Today, the existence of dark matter is almost beyond serious scientific dispute, due primarily to precision constraints from the power spectrum of the cosmic microwave background (CMB; see Spergel et al. 2003, 2007; Komatsu et al. 2011; Planck Collaboration XVI 2014; Planck Collaboration XIII

2016; Planck Collaboration VI 2020) and strong gravitational lensing measurements in galaxy clusters (e.g., Turner et al. 1984; Schneider et al. 1992; Bolton et al. 2006). Despite its elusive nature, dark matter provides the backdrop to modern cosmology and galaxy formation theory, driving the growth of cosmic structure and dictating the underlying environment in which baryons coalesce to become galaxies (see, e.g., Cole et al. 2000; Springel et al. 2005; Bower et al. 2006, 2008; Henriques et al. 2015; Vogelsberger et al. 2014a,b; Schaye et al. 2015; Nelson et al. 2018; Pillepich et al. 2018; Davé et al. 2019).

The most accurate way to measure the total mass content in high-mass groups and clusters is through strong gravitational lensing (e.g., Turner et al. 1984; Schneider et al. 1992; Bolton et al. 2006; Hoekstra et al. 2013; Vegetti et al. 2014; Atek et al. 2018). This approach uses the least assumptions, primarily just general relativity (which is now tested to extreme precision, see Shapiro 1964; Bertotti et al. 2003; Everitt et al. 2011; Abbott et al. 2016). Additionally, by leveraging the assumption of virialization of galaxies and/or groups and clusters, in conjunction with Newtonian gravity (which is an excellent

\* Corresponding author: [abluck@fiu.edu](mailto:abluck@fiu.edu)

approximation of general relativity on these scales), a host of alternative dynamical methods have been developed.

In disk galaxies the rotation curve can be used to infer the mass within a given radius (e.g., Rubin et al. 1978; Begeman 1989; Persic et al. 1996; de Blok et al. 2008). Similar techniques can also be used to leverage velocity dispersion in spheroids (e.g., Erb et al. 2006; Auger et al. 2010; Wolf et al. 2010; Cappellari et al. 2013). In groups and clusters, the satellite velocity dispersion can be used to trace the total mass of the system (e.g., Hickson et al. 1992; Carlberg et al. 1996). Additionally, measurements of the temperature of the circumgalactic medium (CGM) or the intracluster medium (ICM) from X-ray observations can be used to infer mass (e.g., Evrard et al. 1996; Giodini et al. 2009; Pratt et al. 2010). Furthermore, there are also methods to infer cluster masses from observations of the cosmic microwave background (CMB), leveraging the Sunyaev – Zeldovich effect (SZ effect: Sunyaev & Zeldovich 1972; Carlstrom et al. 2002; Bocquet et al. 2019)<sup>1</sup>.

Given the list of dynamical techniques outlined above, one might wonder if there is any need for additional approaches. However, the need for an alternative approach is a simple consequence of all of the above techniques being limited in their applications. For many science goals in extragalactic astrophysics one requires complete halo mass estimates for large galaxy surveys, which rules out the use of the above techniques (even in tandem). Ultimately, all of the above techniques are highly expensive, in terms of observational time required on premier facilities and in terms of funding. These issues are further exacerbated by attempting to constrain dark matter masses at early cosmic times. At best, one can hope to obtain dynamical constraints on dark matter halo masses for a small subset of objects within a wide-field galaxy survey, and frequently none at all for high- $z$  surveys.

As a consequence of the above limitations, an entire industry of techniques has emerged over the past two decades to approximate the halo masses (and more generally the dark sector) in wide-field galaxy surveys. The two main variants are (i) halo occupation distributions (HOD; see Benson et al. 2000; Peacock & Smith 2000; Berlind & Weinberg 2002; Zheng et al. 2007; Croton et al. 2007); and (ii) halo abundance matching (HAM; see Vale & Ostriker 2004; Kravtsov et al. 2004; Conroy et al. 2006; Moster et al. 2010; Guo et al. 2010; Behroozi et al. 2010; Trujillo-Gomez et al. 2011; Klypin et al. 2011, 2015).

The first of these techniques, HOD, functions by constructing the two-point correlation function around each galaxy in the survey as a function of physical (or comoving) distance, and matching this to the dark matter two-point correlation function from a dark matter-only simulation, analytical approximation, or gravitational lensing survey. This process is complicated by the need to assume (or infer) a bias that allows luminous and dark matter structures to have different clustering properties (see Benson et al. 2000; Croton et al. 2007 for discussions). The bias must be learned either from simulations, which historically have been primarily semi-analytic models, or else from high-quality gravitational lensing surveys matching the relevant parameter space (e.g., Mandelbaum et al. 2006, 2016; Mandelbaum 2018). The difficulty with the latter approach is in obtaining strong constraints on the low-mass end of the relationship between luminous and dark matter, as well as challenges with extending to high redshifts. Alternatively, the primary limitation with the for-

mer approach is the inevitability of incurring model bias in the halo estimation when using simulations.

The second of these techniques, HAM, essentially just matches the number density of galaxy stellar mass to the number density of halo mass, the latter historically taken from a dark matter-only N-body simulation. This approach relies on the simple assumption that the most massive galaxies reside in the most massive haloes (see Vale & Ostriker 2004; Conroy et al. 2006; Moster et al. 2010). The basic HAM technique has been improved with sub-halo abundance matching (SHAM; e.g., Hearin et al. 2013; Chaves-Montero et al. 2016; Campbell et al. 2018). This approach leverages information on the satellite population (in addition to centrals), using group stellar masses to constrain dark matter content through a similar number density matching assumption. In recent years, a further enhancement has been constructed, referred to as conditional abundance matching (CAM; see Hearin et al. 2014; Lehmann et al. 2017). This approach utilizes parameters in addition to stellar mass to further constrain dark matter properties, usually in conjunction with semi-analytic models (SAMs).

In addition to the rise of halo abundance matching and halo occupation distribution techniques, some authors have sought to directly obtain a simple mapping from stellar mass of galaxies (and/or groups and clusters) to dark matter halo masses leveraging the  $M_{\text{Halo}} - M_*$  relationship in SAMs (see, e.g., Yang et al. 2007, 2009). The approach of this paper is qualitatively similar to this technique, but applied here to cosmological hydrodynamical simulations, incorporating machine learning to enable multi-parameter dark matter constraints within an iterative group finding algorithm.

Perhaps the most important theoretical contribution to galaxy formation over the past decade has been the emergence of cosmological hydrodynamical simulations (e.g., Vogelsberger et al. 2014a,b; Schaye et al. 2015; Nelson et al. 2018; Pillepich et al. 2018; Davé et al. 2019). These simulations simultaneously model the gravitational and hydrodynamical physics of galaxy formation and evolution within a given background cosmology. It is natural, therefore, to improve both the HOD and (C/SH/H)AM techniques by leveraging these simulations, and several examples of this now exist in the literature (see, e.g., Khandai et al. 2015; Artale et al. 2018; Cochrane et al. 2018).

Nonetheless, the basic assumptions underpinning both halo occupation distributions and abundance matching still exist in these approaches. In the case of abundance matching the essential assumption is that there is a one-to-one mapping between the number densities of either galaxies, groups, or a subset thereof and dark matter haloes. In the case of halo occupation distributions the essential assumption is that there is a simple bias which can be utilized to bring the halo and galaxy density distributions into accord. Yet, important open questions remain: Are these assumptions correct? And, to what extent are the halo properties derived from these approaches biased by their modeling choices?

These questions have been brought into sharp focus by Hadzhiyska et al. (2020), with numerous solutions proposed in Hadzhiyska et al. (2021, 2022, 2023a,b). In summary, different galaxy clustering is expected for emission line and red galaxies in cosmological simulations, which brings into question the standard HOD approach. One solution is to use different biases for different classes of galaxies. In some ways this approach is similar to CAM in that the mapping between luminous tracers and the dark sector can be made more accurate by adding dependencies upon galaxy properties or classes, which can be carefully controlled for in the methodology. However, one potential

<sup>1</sup> This is by no means a complete list of methods to infer dynamical masses.

issue with these approaches is that they impede the possibility of establishing independent connections between these properties and dark haloes in observations, since they are assumed in advance. One of the first steps in the analysis of this work is to investigate precisely which parameters are effective at estimating halo masses in cosmological simulations and limit our mapping to only those which are essential. Ultimately, this is a trade-off between accuracy and independence in the output halo catalog.

In parallel with the theoretical advances leading to cosmological hydrodynamical simulations, a revolution in machine learning is now underway (e.g., Pedregosa et al. 2011; Brink et al. 2013; Baron 2019), with many applications to extragalactic astrophysics (e.g., Banerji et al. 2010; Mucesh et al. 2021; Bluck et al. 2022; Piotrowska et al. 2022; Remy et al. 2023; Huertas-Company et al. 2024; Ho et al. 2024). Consequently, it is now possible to learn arbitrarily complex mappings from luminous tracers to dark matter properties, utilizing multiple input parameters, even among inter-correlated variables. Hence, one need not restrict the mapping to individual number densities of a given parameter, most frequently stellar mass (as in AM), or to the identification of one or multiple biases (as in HOD).

Moreover, the machine learning tools may be trained on state-of-the-art cosmological hydrodynamical simulations, enabling one to fully account for the baryon – dark matter interaction in the relationship between dark matter and luminous tracers. In conjunction, these advances remove the need for simplifying assumptions in constructing a mapping from luminous tracers to dark matter properties. One interesting example of this approach is to use the morphologies of galaxies to constrain their dark matter properties in imaging surveys (see Hahn et al. 2024). Our approach should be seen as complementary to this work in that it is geared towards spectroscopic surveys without the need for high-resolution imaging. One advantage of our approach is that morphology remains independent of the estimated halo masses and hence one can look for connections between halo mass and galaxy morphology without prejudicing the result.

Ultimately, the overarching goal of this work is to leverage cosmological hydrodynamical simulations and modern machine learning techniques to construct a novel approach for solving the problem of which galaxies reside in which dark matter structures. One of the main advantages of this methodology is to rigorously determine the uncertainties of the estimated halo masses, and group properties. We also aim for scalability in the halo finder algorithm, enabling straightforward application to the vast galaxy surveys of the future. Our primary motivation is to provide a fast, efficient, and accurate halo finder and halo mass estimator which can be used at cosmic noon ( $z \sim 1-2$ ) in the upcoming VLT-MOONRISE spectroscopic galaxy survey (see Maiolino et al. 2020; Cirasuolo et al. 2020).

To this end, we present *Dark from Light: Dfl*<sup>2</sup>. DfL is a publicly available halo finder and halo mass estimator intended for application to wide-field spectroscopic galaxy surveys. DfL utilizes random forest regression trained on contemporary cosmological simulations (particularly, EAGLE: Schaye et al. 2015; Crain et al. 2015; McAlpine et al. 2016; and IllustrisTNG: Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018; Springel et al. 2018). The machine learning modules are incorporated into a physically motivated virial group finder, which simultaneously identifies clustering and creates a mapping from luminous tracers to the dark sector. Special care is made to ensure that this method can be applied to a wide

range of redshifts. Moreover, DfL also characterizes both the random and systematic error associated with the method (including an estimate of uncertainty from model choice). This enables the robust application to a host of contemporary cosmological and astrophysical problems.

The paper is structured as follows. In Section 2 we present the simulated data used in this work from EAGLE and IllustrisTNG. In Section 3 we present a thorough description of the halo finder pipeline. Initially, we explore which luminous tracers are most effective at constraining halo mass. We then detail the training of the machine learning regression models. Finally, we present a step-by-step explanation of the full pipeline. In Section 4 we test the performance of the halo finder pipeline at  $z = 0-2$  in EAGLE and IllustrisTNG. We also test the performance in cross-validation mode (where one model is used in training and another for testing). Moreover, we compare the results from DfL to a more conventional method (abundance matching applied to total stellar masses output from a friends-of-friends, FOF, group finding algorithm), and compare the output  $z = 0$  halo mass – stellar mass relationship to observational constraints utilizing strong lensing. Finally, we systematically explore how stellar mass uncertainty and survey incompleteness impact halo recovery in DfL. In Section 5 we discuss the potential of DfL applied to VLT-MOONRISE data for a host of scientific applications. The main contributions of this paper are summarized in Section 6.

We also provide three appendices to this work. Appendix A presents a detailed users guide for the public DfL code. Appendix B presents the full details of the FOF-AM approach used for comparison to DfL in Section 4. Appendix C lists the hyper-parameters used in the random forest modules incorporated within DfL, and in the ANN modules used to ensure the RF achieves an optimal regression solution.

Unless otherwise specified, we assume a spatially flat  $\Lambda$ CDM background cosmology with  $\Omega_M = 0.3$ ,  $\Omega_\Lambda = 0.7$ , and  $h \equiv H_0/(100 \text{ km/s/Mpc}) = 0.7$ .

## 2. Simulations data

In this work we employ two cosmological simulation suites to learn redshift and model-dependent mappings from luminous tracers to the dark sector. Specifically, we utilize: (i) IllustrisTNG<sup>3</sup> (hereafter TNG; Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018; Springel et al. 2018); and (ii) EAGLE<sup>4</sup> (Crain et al. 2015; Schaye et al. 2015; McAlpine et al. 2016). From each simulation we extract data from snapshots at  $z = 0-3$  for all sub-haloes with stellar mass,  $M_* > 10^{9.5} M_\odot$ . Full information on the two simulations are provided in the above references. Additionally, we have prepared a docker<sup>5</sup> to aid in data access and the application of these catalogs.

Both simulation suites model a  $\Lambda$ CDM background cosmology tracing both gravitational and fluid dynamical physics directly. Subgrid models are utilized to implement gas cooling, star formation, stellar evolution, supernova feedback, supermassive black hole seeding and growth, active galactic nucleus (AGN) feedback, and the formation of metals, among several other properties. The principal differences between the models are due to the choice of code used to solve the hydrodynamical fluid equations (explicitly, adaptive moving mesh in

<sup>3</sup> IllustrisTNG data access: [www.tng-project.org/](http://www.tng-project.org/)

<sup>4</sup> EAGLE data access: <http://icc.dur.ac.uk/Eagle/>

<sup>5</sup> Simulations docker: <https://hub.docker.com/u/jpiotrowska>

<sup>2</sup> <https://github.com/abluck/Dark-from-Light>

IllustrisTNG vs. smoothed particle hydrodynamics in EAGLE), and most importantly, the details of the subgrid recipes used to model feedback, especially from supermassive black holes and AGN.

### 2.1. IllustrisTNG

We utilize the IllustrisTNG-100-1 simulation (Nelson et al. 2018; Pillepich et al. 2018), which has a box size of  $\sim(100\text{cMpc})^3$  and is run using the AREPO unstructured moving mesh code (Springel 2010). This code has been updated to include magnetic fields in addition to gravity and hydrodynamics. The cosmological parameters are based on Planck Collaboration XIII (2016). Dark matter and stars are treated as particles, with particle masses of  $M_{DM} = 7.5 \times 10^6 M_\odot$  and  $\langle M_{\text{star}} \rangle = 1.4 \times 10^6 M_\odot$ , respectively. Baryons in gas are modeled using the adaptive moving mesh Voronoi binning approach (see Springel 2010).

Baryons transform into stars following a simple density threshold prescription (see Pillepich et al. 2018), based on the empirical Kennicutt (1998) law, assuming a Chabrier (2003) initial mass function (IMF). Dark matter haloes and sub-haloes (including galaxies) are identified through the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009). This approach uses a linking-length friends-of-friends algorithm to associate dark matter, stars, and gas within gravitationally connected systems.

In this work, the principal parameters of interest are dark matter halo mass (specifically,  $M_{200}$ ) in addition to the stellar mass of galaxies, groups, and clusters. Additionally, we also consider the number of galaxies above a stellar mass threshold associated with each group. Hence, we are primarily interested in the relationship of stellar and halo properties. The main subgrid physics which can impact the stellar mass – halo mass relationship at the relatively high masses probed in this work is AGN feedback, which we discuss below.

IllustrisTNG-100-1 seeds supermassive black holes in low-mass dark matter haloes and models their evolution through Eddington-limited Bondi-Hoyle accretion (Hoyle & Lyttleton 1939; Bondi & Hoyle 1944). This simulation incorporates three distinct modes of AGN feedback (see Weinberger et al. 2017, 2018): (i) quasar mode (which operates at high Eddington ratios); (ii) kinetic mode (which operates at low Eddington ratios); and (iii) radiative ionization (operational at all Eddington ratios). Among these, only the kinetic mode has a significant impact on quenching within the simulation (see Weinberger et al. 2017; Zinger et al. 2020; Piotrowska et al. 2022), and hence leaves a major imprint on the stellar – halo mass relationship, which is critical for this work.

In the kinetic mode, which is only operational at  $M_{\text{BH}} \gtrsim 10^8 M_\odot$ , a fraction of the accreted rest energy is converted into the kinetic energy of gas cells surrounding the black hole in a stochastic fashion (see Weinberger et al. 2017). Although the direction of the momentum kick is random, this is constrained to be isotropic over long time scales, preserving conservation of momentum in the simulation. This kinetic feedback induces turbulence and outflows in the interstellar medium, and additionally provides long-term heating of the circum-galactic medium through gas percolation and shocks. In conjunction, kinetic mode AGN feedback shuts down star formation in the highest mass galaxies, leading to curvature in the fundamental  $M_* - M_{\text{Halo}}$  relationship.

### 2.2. EAGLE

We utilize the EAGLE-RefL0100N1504 simulation (Schaye et al. 2015; McAlpine et al. 2016), which has a box size of  $(100\text{cMpc})^3$  and is run using an updated version of the GADGET-3 smoothed particle hydrodynamics (SPH) code (Springel et al. 2005); Dalla Vecchia et al. in prep.). This simulation adopts a cosmological model based on Planck Collaboration XVI (2014), assuming a spatially flat  $\Lambda$ CDM background universe, dominated by mass at early cosmic times and dark energy at late cosmic times. EAGLE models dark matter and baryons (in all phases) as particles with  $M_{\text{DM}} = 9.7 \times 10^6 M_\odot$  and an initial baryonic mass of  $M_b = 1.81 \times 10^6 M_\odot$ , respectively.

Baryons transform into stars following a metallicity-based criterion with a pressure dependent star formation rate (see Schaye et al. 2015), which is ultimately based on the empirical Kennicutt (1998) law. Dark matter haloes and sub-haloes (including galaxies) are identified through the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009), as in TNG. As with TNG, the main subgrid physics which can impact the stellar mass – halo mass relationship at high masses is AGN feedback. Stellar and supernova feedback have a big impact on star formation for low-mass galaxies, but by  $M_* \sim 10^{9.5} M_\odot$  these effects become marginal, due to the gravitational potential of galaxies far exceeding the binding energy required to hold on to supernova ejecta. Therefore, we focus on the AGN feedback prescription here.

In EAGLE, AGN feedback serves as the primary mechanism for quenching massive galaxies, and hence the major driver of the  $M_* - M_{\text{Halo}}$  relation at high masses. Supermassive black holes are seeded in low-mass dark matter haloes and grow through Eddington-limited Bondi-Hoyle accretion (Hoyle & Lyttleton 1939; Bondi & Hoyle 1944), as in TNG. However, AGN feedback is modeled in EAGLE using a single mode, where a fraction of the accreted rest mass energy is converted into thermal energy of neighboring gas particles. Energy exchange between the black hole accretion disk and the galaxy is implemented via thermal injections in a stochastic, burst-like manner (see Crain et al. 2015).

AGN feedback in EAGLE initiates quenching by driving ISM outflows, and helps maintain quiescence in massive galaxies by heating the circum-galactic medium (CGM) from shocks induced by the outflows. However, we have previously noted that the long-term quiescence of massive galaxies in EAGLE is less effective than in TNG (see Piotrowska et al. 2022). Nonetheless, EAGLE does provide a reasonable approximation of the multi-epoch stellar mass functions (see Schaye et al. 2015).

### 2.3. Why two simulations?

In principle one can train a machine learning algorithm to learn how to estimate dark matter properties from luminous tracers in a single simulation (as in, e.g., Yang et al. 2009; Hahn et al. 2024). We opt not to do this due to remaining uncertainty in the correct form of the baryonic physics of galaxy formation and evolution.

Gravitational physics is extremely well understood on the scales of galaxies, groups, and clusters. Additionally, the cosmological parameters governing the background expansion of the Universe are also very well constrained (e.g., Planck Collaboration XVI 2014; Planck Collaboration XIII 2016; Planck Collaboration VI 2020). Hence, a single model would likely be sufficient for investigating the formation of dark matter haloes in isolation. However, in this work we seek to

estimate dark matter properties from luminous tracers, which are all baryonic in nature. More precisely, all of the luminous tracers in this work are of stellar origin. Hence, the baryonic physics – including gas cooling, star formation, and, crucially, feedback – is critical to the required mapping of luminous tracers to dark matter properties.

Unlike the background cosmology and gravitational physics, the physics of baryons in galaxy evolution is much less well constrained (see, e.g., [Somerville & Davé 2015](#) for a review). To combat this issue, we elect to incorporate two successful<sup>6</sup> simulations of galaxy evolution with very different subgrid recipes and hydrodynamical solvers. This allows us to begin to quantify the impact of model choice on the inferred dark matter halo properties. Ultimately, this will yield an initial estimate of the bias induced on halo mass recovery from the choice of the (relatively) poorly constrained baryonic physics. In future work we plan to extend this framework to incorporate essentially all public cosmological simulations, enabling users to have complete flexibility with the application of DfL. In this first work in the series, we aim more modestly for a first proof of concept of this approach with two popular simulations.

### 3. Methods

#### 3.1. Defining the problem

In essence we seek a mapping from observable luminous tracers, which are relatively straightforward to obtain in extant and upcoming wide-field galaxy surveys, to dark matter halo mass. To be useful, this mapping must yield accurate halo masses within some well defined random uncertainties, and have constrainable biases. Hence, schematically, we look for a mathematical relation of the form:

$$\{\text{Luminous Tracers}\} \rightarrow \{\text{Dark Matter Properties}\}. \quad (1)$$

As a starting point, one must employ known dark matter halo masses for a representative sample of haloes, and hence galaxy, group, and cluster masses. At cosmic noon this is only possible in simulations at present. Direct measurements of halo masses (e.g., from strong gravitational lensing or density – temperature profiles) are limited to very high-mass haloes at these epochs, and hence cannot be used to learn a general mapping across the full diversity of halo types. This poses an immediate problem in that any mapping will be model-dependent. This furthermore implies that the method will most probably be biased relative to the correct galaxy formation model of the Universe. To combat these issues we employ two very different cosmological simulations and rigorously compare the predicted dark matter halo properties from each to ascertain how impactful the model choice is in the mapping from luminous tracers to the dark sector.

The next issue is how to optimally learn the mapping from light to dark in a given simulation. In the simplest case of one luminous tracer, a straightforward nonlinear fit could be made for each appropriate epoch and galaxy type. However, in general this is not necessarily optimal. To form a more general mapping one needs to consider multiple input parameters, which are in general inter-correlated. This leads to a highly complex mathematical problem, which is not easily tractable with conventional

fitting techniques. Conversely, this problem is ideally suited to machine learning.

#### 3.2. A machine learning solution

Machine learning describes the capacity of a certain class of algorithms to adapt to inputs, learning to solve complex problems without explicit instructions. In supervised learning, known solutions are made available to the machine learning system to aid the algorithm in assigning specific weights (or hyperparameters) which minimize the inaccuracy of the task at hand. There are a number of machine learning algorithms which are potentially suitable to solving the mapping from luminous tracers to dark matter properties, described above. These range from deep learning with artificial neural networks to more straightforward approaches like decision trees and random forests.

Additionally, one must specify whether classes or numbers are desired as the output. Since dark matter halo masses form a continuum, it is clear in this case that we require a numerical output and hence we must utilize some form of supervised machine learning with regression (see, e.g., [Predgosa et al. 2011](#)).

One possibility is deep learning with multi-layered, fully connected artificial neural networks (ANN, e.g., [Ilbert et al. 2006](#); [Dieleman et al. 2015](#); [Teimoorinia et al. 2016](#)). This type of machine learning is highly effective at learning nonlinear mappings between input data and a desired output. However, the key weakness of these techniques is interpretability. Due to multiple inputs with nonlinear activation functions between layers in an ANN, it is extremely hard to reverse engineer the trained network to understand why it works. For many applications this is not especially problematic. However, if one wishes to gain insight from the machine learning technique, artificial neural networks are frequently highly limited.

In the present application, our first goal is to learn precisely which parameters are most effective at making the mapping from light-to-dark, and, crucially, limit their number as much as possible. The latter is a result of our ultimate goal to have the final dark matter halo catalogs be as independent as possible from other measurable parameters. This goal is not well suited to deep learning. Fortunately, in preliminary testing we have found that ANN is no more effective than the more transparent Random Forest technique for this particular problem, and so we continue with this more simplistic approach.

Moreover, the RF technique has numerous advantages over the ANN approach. First, the RF enables the extraction of feature importance, which we utilize in Section 3.3 to investigate the best parameters to incorporate into DfL. Furthermore, the RF offers a natural way to extract uncertainties (via the distribution of individual tree predictions) and propagate these through the DfL pipeline efficiently (which is discussed in Section 3.5.7). While there are methods that enable full PDF error propagation through ANN with Bayesian neural networks (see, e.g., [Chua & Vallisneri 2020](#); [Hortúa et al. 2023](#); [Jones et al. 2024](#)), this is much more complex, slower to train and apply, and in this case not necessary to yield optimal performance.

##### 3.2.1. The random forest method

Random forests are formed from a collection of decision trees with differences between them enforced through bootstrapped random sampling of the input data, and (optionally) through partial sampling of the input features. A single decision tree is formed in a deterministic manner from a set of input data, a desired output parameter, and a performance indicator to minimize.

<sup>6</sup> This is meant in the sense that both simulations accurately reproduce the observed stellar mass functions of galaxies across a wide range in cosmic time, making them highly suitable to learn a light-to-dark mapping from (see, e.g., [Schaye et al. 2015](#); [Nelson et al. 2018](#); [Pillepich et al. 2018](#)).

In our application, the input data will be luminous tracers that are known, or expected, to reveal information on the dark sector, and the output parameter will be halo mass. The performance indicator, or loss-function, is chosen to be the mean square error, explicitly

$$\text{MSE} \equiv \frac{1}{N} \sum_{i=1}^N (Y_i - \hat{Y}_i)^2, \quad (2)$$

where  $N$  is the total number of predictions made,  $Y_i$  is the true value of each prediction, and  $\hat{Y}_i$  is the predicted value from the random forest, or partial step in each decision tree in training. The MSE quantifies how accurate the predictions of a given quantity are, treating over- and under-estimates as equally erroneous.

Each tree in the random forest has branches structured as follows. The single variable (referred to as a feature), and boolean operation using that variable, which minimizes the MSE in the combined daughter nodes is chosen at each parent node. For instance, this may be stellar mass with a cut at  $M_* > 10^{10} M_\odot$ . After splitting the data with this criterion and this feature, the mean halo mass for each subset is computed, using the known halo masses in the training data, and the MSE performance is computed. For each branch the process continues until either every final (leaf) node contains just one training datum, or else a predetermined limit is reached, employed to prevent over-fitting the data. This is continued for each tree in the forest and the final predictions are taken as the average over all of the individual tree predictions.

Due to the relative transparency of the random forest architecture, it is possible to extract precisely how much value each input parameter has to solving the regression problem. This may be quantified as the relative importance of each variable as

$$I_R(k) = \frac{1}{N_{\text{trees}}} \sum_{\text{trees}} \left( \frac{\sum_{nk} N(nk) \Delta\text{MSE}(nk)}{\sum_n N(n) \Delta\text{MSE}(n)} \right), \quad (3)$$

where in the numerator the sum is made over all nodes in a given tree which utilize variable,  $k$ , and the denominator sums over all nodes in the tree, regardless of which features are used to make the cuts. The change in MSE ( $\Delta\text{MSE}$ ) is weighted by the number of data which reach that individual node ( $N$ ). This gives the importance of any given variable,  $k$ , to that tree. Finally, the overall importance is taken as the average importance across all of the trees within the forest.

In this work we generate our Random Forest models using `RandomForestRegressor` from the `SciKit-Learn` PYTHON package (see [Pedregosa et al. 2011](#))<sup>7</sup>. All hyper-parameters are listed in [Appendix C](#) to allow the reproducibility of our results.

To prevent over-fitting, we systematically vary the `min-samples-leaf` (MSL) hyper-parameter to ensure an optimal regression performance in unseen validation data, but, crucially, not the training data. Following [Bluck et al. \(2022\)](#), we also impose a limit on the similarity of performance between training and testing data sets. Here we require that  $\Delta\sigma < 0.03$  dex (where  $\sigma \equiv \sqrt{\text{MSE}}$ , and MSE is defined in Eq. (2)).

We also allow only partial access to features in any given tree, using the `sqr` mode for the `max-features` hyper-parameter. While this is not the optimal route to establish causality (see [Bluck et al. 2022](#)), it is the most accurate way to train the random forest architecture (see, e.g., [Breiman 2001](#)), limiting the impact of erroneous values in any one feature in the overall estimation of halo masses.

<sup>7</sup> SciKit-Learn: <https://scikit-learn.org/stable/>

### 3.3. What are the best luminous tracers of halo mass?

We start by investigating which parameters are most useful for predicting dark matter halo masses in the two simulations utilized in this work, EAGLE & TNG. This will inform us as to the optimal parameters to incorporate into our halo finder pipeline.

In [Fig. 1](#), we show the results from numerous random forest regression runs in the TNG (left column) and EAGLE (right column) simulations. Each random forest run predicts halo masses (explicitly,  $M_{200}$ ) from a wide set of luminous tracers. The decision trees are trained with: stellar mass ( $M_*$ ); group total stellar mass ( $GM_*$ , evaluated within  $R_{200}$ ); the total number of galaxies in the group above  $M_* > 10^{9.5}$  ( $N_g$ , evaluated within  $R_{200}$ ); the galaxy over-density at the 3rd, 5th and 10th nearest neighbor (evaluated in 2D for an observational-like parameter); and two star formation tracers: `sSFR` ( $=\text{SFR}/M_*$ ), and the star forming class (discussed below).

The local galaxy over-densities are computed as

$$\delta_N \equiv \log(\Sigma_N) - \langle \log(\Sigma_N) \rangle_z, \quad (4)$$

where

$$\Sigma_N = \frac{N}{\pi D_N^2} \quad (5)$$

and where  $D_N$  is the physical distance to the  $N$ th nearest galaxy neighbor, evaluated in 2D observer-like space. We also require a velocity difference along the line of sight (i.e., perpendicular to the observer-like plane) of  $|\Delta V| < 1000$  km/s. The over-density (Eq. (4) above) is then taken as the difference between each galaxy's individual local surface density and the geometric mean of all surface densities in the simulation at the same redshift snapshot as the galaxy in question (as in [Bluck et al. 2020a](#); [Goubert et al. 2024](#)).

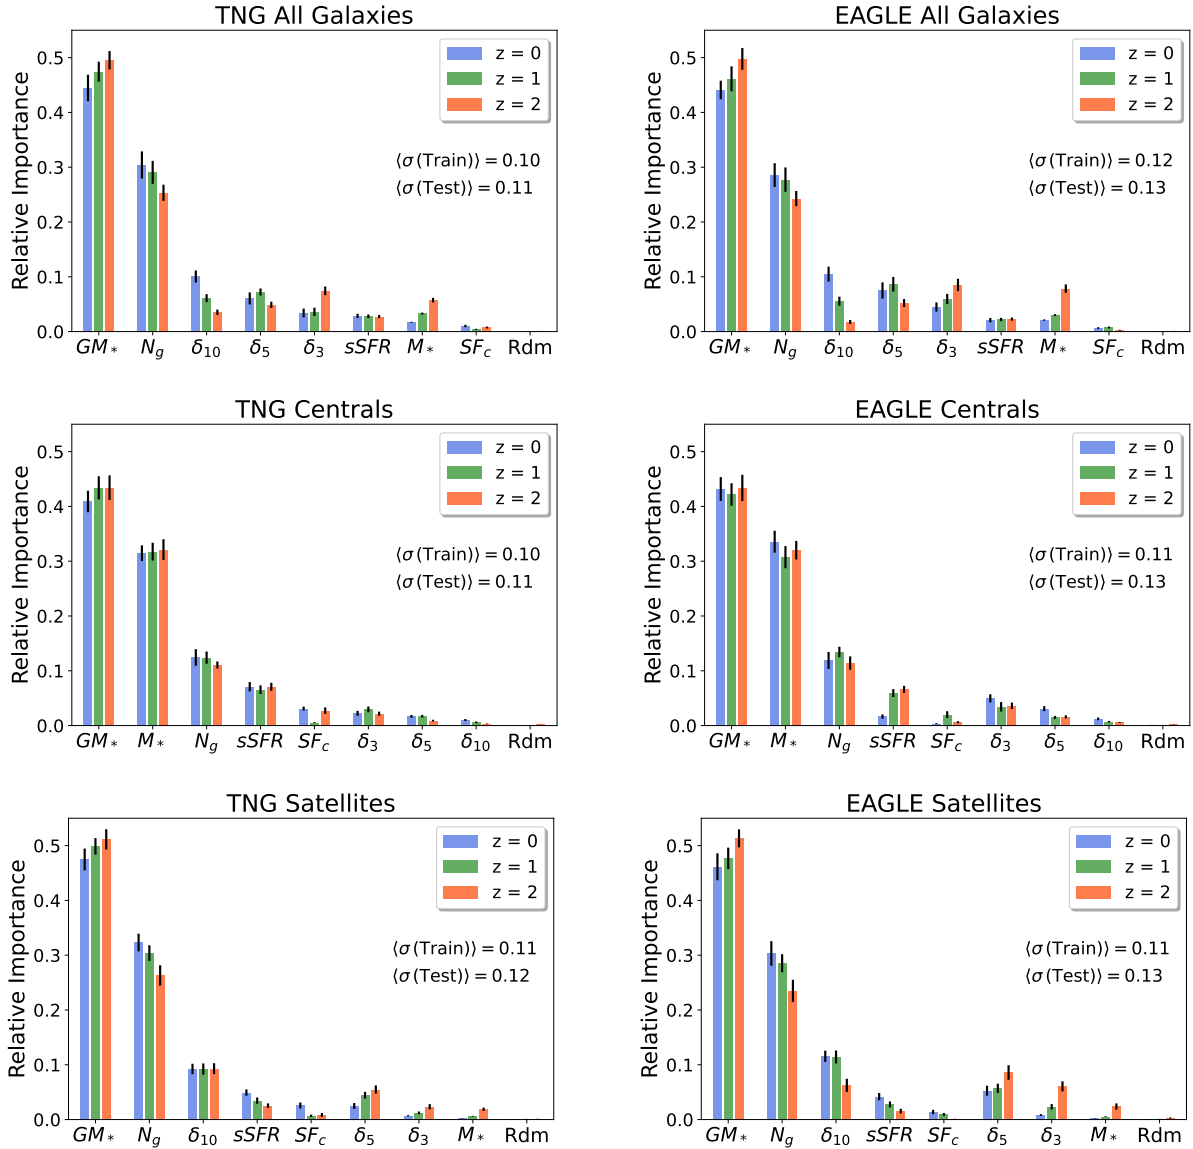
Additionally, we add information on the star forming state of galaxies. We include the specific star formation rate (defined above) and the star forming class (`SFc`), defined by identifying galaxies which are forming stars with `sSFR` lower than one order of magnitude of the peak of the star forming distribution at that epoch (as in [Bluck et al. 2023](#)). Specifically, star forming galaxies are assigned a number of zero, and quenched galaxies a number of one, with the criterion for being quenched being

$$\text{sSFR}(z) < \text{sSFR}_{\text{peak}}(z) - 1 \text{ dex}. \quad (6)$$

Finally, we also include a random number (`Rdm`) to aid in interpreting whether any importance may be spurious.

In the top row of [Fig. 1](#) we present results for all galaxies treated together, in the middle row we show results for central galaxies, and in the bottom row we show results for satellite galaxies. Within each panel, we present the results separately for three redshift snapshots in the simulations,  $z = 0, 1, 2$ . The parameters are listed along the  $x$ -axis of each plot, ordered by how important they are for predicting halo mass in the regression analysis for that class of galaxy in TNG at  $z = 0$ . The  $y$ -axes in [Fig. 1](#) display the relative importance of each feature, as defined in Eq. (3) above. This quantifies which parameters are most effective for predicting dark matter halo masses in simulations.

For every galaxy class, redshift snapshot, and in both simulations, the group total stellar mass is consistently revealed to be the most important (and hence predictive) luminous tracer of dark matter halo mass. Therefore, this is clearly a parameter we want to include in our halo finder pipeline. However, it is not



**Fig. 1.** Random forest regression analyses showing the predictive importance of various observable galaxy and environmental parameters for inferring dark matter halo mass, explicitly  $M_{200}$  (as listed along the  $x$ -axis; see Section 3.3 for definitions). The left column shows results from the TNG simulation; the right column shows results for the EAGLE simulation. Results for all galaxies are shown in the top row, results for central galaxies are shown in the middle row, and results for satellite galaxies are shown in the bottom row. In each panel, results are shown separately for three redshift snapshots ( $z = 0, 1, 2$ ). The uncertainties on the relative importance are computed as the standard deviation across 100 independent runs. The overall average performance of the random forest regressor ( $\sigma \equiv \sqrt{\text{MSE}}$ ) is shown in each panel for both the training and independent testing samples (each containing 50% of the full data set). For all populations of galaxies, and at all epochs studied, the total group stellar mass is seen to be the most predictive parameter for inferring halo masses in both EAGLE & TNG. The second most important parameter depends on galaxy class, with the number of galaxies in the group being the second most important parameter for all galaxies and satellites, and stellar mass being the second most important parameter for centrals. All other investigated parameters (e.g., local densities, star formation tracers) are substantially less predictive than the first three parameters. We note that there is a high degree of similarity between the simulations and redshift snapshots, which implies that these results are not strongly model-dependent and do not vary significantly with cosmic time.

immediately obvious how to compute this directly from survey data. Obviously, one must first have a group defined. Of course, one could use a group finding algorithm such as a friends-of-friends (FOF, see, e.g., Press & Davis 1982; Davis et al. 1985; Springel et al. 2001; Knebe et al. 2011) in order to identify the groups first. But our present goal is to simultaneously find halo masses and group information, enabling the interrelated information on clustering and mass to interact. To help with this problem, we next look at the second most predictive parameter in various classes.

For galaxies as a whole, and for satellites, the number of galaxies in the group is the second most predictive parameter of halo mass. While potentially useful once groups are defined, this still does not help us to identify groups, or motivate their expected size. For central galaxies, however, the stellar mass of the galaxy is the second best parameter, with a very respectable importance score. This is interesting because if we can identify central galaxies we can immediately get a reasonable estimate of halo mass and, moreover, use this to motivate the expected group size. Via iteration one can then ascertain the group parameters

**Table 1.** Performance comparison between the random forest (RF) and artificial neural network (ANN) methods.

| Simulation          | ANN All      | ANN Group    | ANN Central  | RF All       | RF Group     | RF Central   |
|---------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| TNG (All)           | (0.00, 0.12) | (0.00, 0.12) | (0.00, 0.14) | (0.00, 0.11) | (0.00, 0.12) | (0.00, 0.15) |
| EAGLE (All)         | (0.00, 0.12) | (0.00, 0.13) | (0.00, 0.14) | (0.00, 0.12) | (0.00, 0.13) | (0.00, 0.15) |
| TNG ( $N_g > 1$ )   | (0.00, 0.12) | (0.00, 0.11) | (0.13, 0.16) | (0.00, 0.11) | (0.00, 0.11) | (0.12, 0.18) |
| EAGLE ( $N_g > 1$ ) | (0.01, 0.12) | (0.00, 0.12) | (0.13, 0.16) | (0.00, 0.12) | (0.00, 0.12) | (0.15, 0.17) |

**Notes.** The performance is given in each cell position as the bias followed by the dispersion (defined as in Fig. 2): ( $b$ ,  $\sigma$ ) [dex]. We note the similar performance between the ANN and RF methods. Moreover, we note the significant improvement when moving from central to group parameters (especially for multi-parameter groups, as seen in Fig. 2), but the lack of additional improvement moving to the “all” parameter set (as seen in Fig. 1). This serves to justify our choice in parameters and machine learning method.

( $GM_*$  and  $N_g$ ) and improve the halo mass estimate further. This is, of course, expected from the abundance matching approach (see Moster et al. 2010; Behroozi et al. 2010). However, as we will see, the stellar mass of central galaxies alone is not an optimal predictor of group halo mass.

All other parameters considered in Fig. 1 are of much less predictive power than the three discussed above. Furthermore, we have determined that the improvement in overall performance (total MSE) is only negligibly affected by removing these parameters (see Table 1). Thus, even though they do technically add something to the regression analysis, their value is extremely limited. Including parameters with very little predictive power only makes the pipeline outputs harder to use in practice. Moreover, using the minimum number of parameters to achieve the desired mapping results in having the final halo mass estimates independent of more galaxy properties, which is especially valuable for many applications with star formation tracers.

Hence, as a result of the above discussion on Fig. 1, we opt to utilize the following luminous tracers to infer dark matter halo masses in our pipeline: (i) total group stellar mass ( $GM_*$ ), (ii) total number of group members ( $N_g$ ), (iii) central stellar mass ( $CM_*$ ), and (iv) redshift ( $z$ ). The final parameter allows for the possibility of significant evolution in the relationships between luminous and dark matter throughout an observational survey.

### 3.4. Training the RF regression models

From each simulation we construct three subsamples. First, we make a simple cut to the simulation boxes at  $X = 50$  cMpc. This approximately splits each simulation in half. The first half ( $X < 50$  cMpc) is used to train and validate the RF models, as outlined in this section. Within this half, training operates on a randomly selected 70% of the available simulated groups. The remaining 30% of groups are used to validate the RF models to ensure that we do not over-fit the data (as discussed below). The second half of each simulation ( $X > 50$  cMpc) is reserved for testing the DfL algorithm in Section 4. We emphasize that no individual groups or galaxies used in either training or validation of the RF models are used to test the performance later on in the paper. Additionally, for the testing data we restrict to an observational-like 2D +  $z$  space from the full 6D phase space of each simulation (see Section 3.6).

We adopt a two stage approach for estimating halo masses, and more generally halo properties, from luminous tracers. First, we identify systems which are most likely centrals (discussed below in Section 3.5) and leverage the strong relationship between central stellar mass and halo mass in simulations to gain a first estimate of the halo mass, and hence the virial radius and velocity of the group. This part is very similar to the

standard abundance matching procedure, with a redshift dependence added. We then use this information to identify likely group members (discussed at length in Section 3.5).

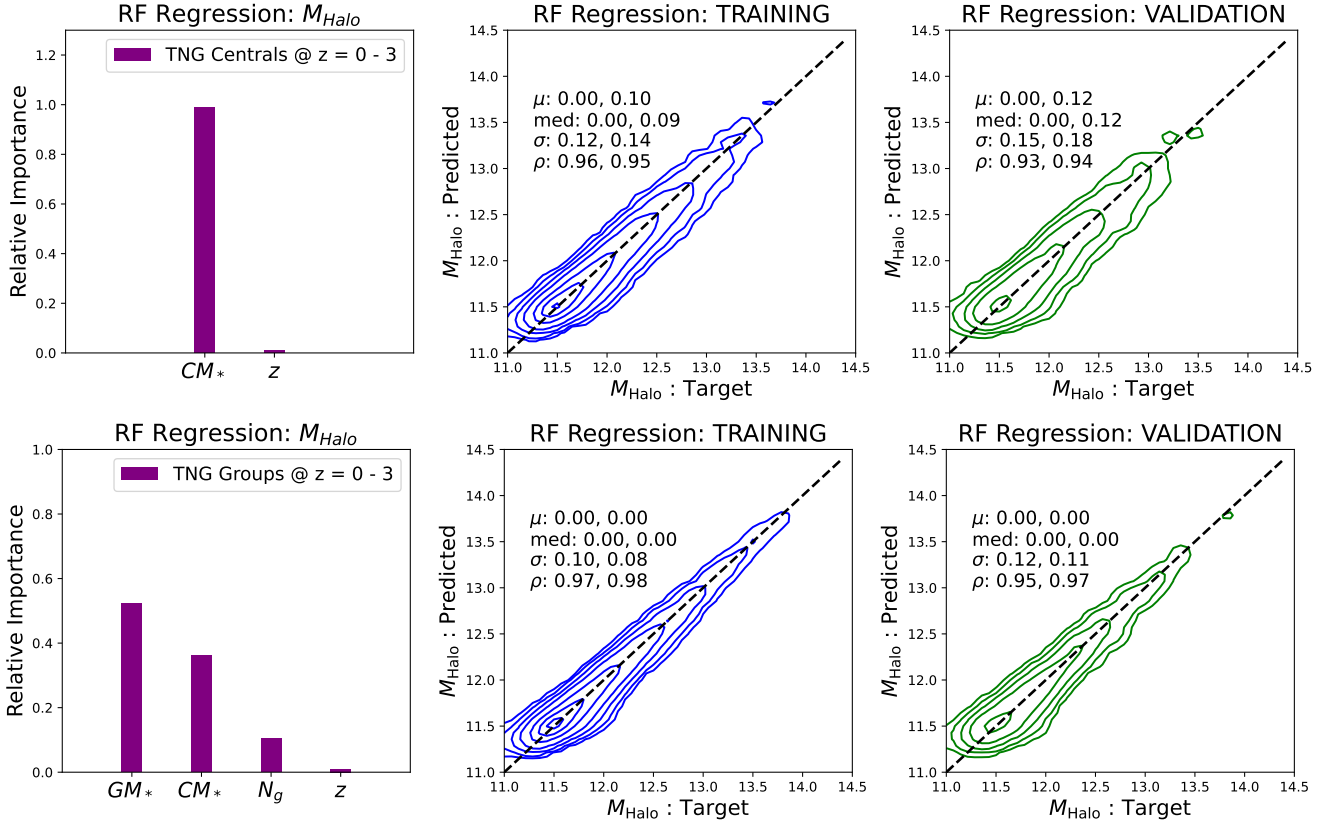
The preliminary group identification allows us to estimate the total group stellar mass and number of galaxies within the group, which are then used in addition to the central’s stellar mass and redshift to refine the prediction for the halo mass in the second stage. In DfL, we iterate this procedure to find the optimal size of the group and halo mass estimate. The full details of this pipeline are discussed in Section 3.5, but first we show examples of the two versions of the random forest regression models which are incorporated into this strategy.

#### 3.4.1. For centrals:

In Fig. 2 (top panels) we show results from an example trained random forest regression model to predict halo masses from central stellar mass and redshift alone in the TNG simulation. In this case, almost all feature importance is given to stellar mass, with a small importance found in redshift (as indicated in the left panel). The weak importance of redshift is interesting, especially given that both the halo and stellar mass functions evolve significantly as a function of cosmic time (see, e.g., Fig. B.2). However, the impact of varying redshift on the  $M_* - M_{\text{Halo}}$  relationship is actually quite weak in both TNG and EAGLE. At a fixed stellar mass, the expected halo mass does not change dramatically with redshift. Conversely, the probability of reaching a given stellar (and halo) mass does change significantly with redshift, but in a manner which keeps the relationship close to constant.

In training, the regressor is capable of learning a bias-free mapping from stellar mass to halo mass, with a statistical uncertainty of  $\sigma = 0.12$  dex (see center panel). We next test the performance of this model on unseen data from the same simulation. Explicitly, we split the input simulation data into 70% training and 30% testing, to check for over-fitting. For the testing data the mapping is still essentially bias free, but the random uncertainty on the halo mass increases to  $\sigma = 0.15$  dex (see left panel). This is still a reasonably accurate recovery of halo mass, given a perfect knowledge of the central’s stellar mass. Indeed, this is a comparable (or slightly improved) performance compared to that achievable in this simulation through abundance matching (see Appendix B, especially Fig. B.2).

However, when one considers the performance of the central-only RF model on multi-galaxy groups (statistics presented to the right of commas on the center and right panels of Fig. 2), it is clear that there is a problem. Multi-galaxy groups are biased in their halo masses to  $b = +0.12$  dex, and have a significantly worse statistical uncertainty of  $\sigma = 0.18$  dex. This strongly



**Fig. 2.** Training and validation of the RF regression models in the TNG simulation at  $z = 0-3$ . The top panels show the results from the central-only RF model and the bottom panels show the results from the full group model. The left panels show the distribution in relative importance for the parameters used in training. This includes central galaxy stellar mass ( $CM_*$ ) and redshift ( $z$ ) for the central-only run, and adds group stellar mass ( $GM_*$ ) and the number of galaxies in the group ( $N_g$ ) in the case of the full group run (see Section 3 for full definitions of these parameters). The center panels show the recovery of halo mass in the training stage and the right panels show the recovery of halo mass in the unseen validation sample (which is used to optimize the performance). The mean ( $\mu$ ) and median (med) offset, standard deviation ( $\sigma$ ), and Spearman rank correlation strength ( $\rho$ ) are displayed for all groups (left) and multi-galaxy groups (right, after comma) for both training and testing. There is a modest improvement in performance in the full group RF model compared to the central-only model for the full data set. Crucially, however, the improvement is much more significant for the multi-galaxy groups. Hence, there is a clear need to incorporate the group information in order to have unbiased and more accurate halo mass estimation for groups and clusters.

motivates us to incorporate group information into DfL, in addition to that of the central galaxy.

We perform an analogous training of central RF models for the EAGLE simulation. The plots from this simulation are not shown here for the sake of brevity, but the performance statistics achieved are presented alongside TNG in Table 1. Additionally, in Table 1 we compare the performance from RF to an ANN approach, which demonstrates that the former is as successful as the latter for solving this problem (with significant additional benefits, as outlined in Section 3.2).

### 3.4.2. For groups:

In Fig. 2 (bottom panels) we show the results from an example random forest regression model trained on groups from the TNG simulation at  $z = 0-3$ . Here we utilize:  $GM_*$ ,  $N_g$ ,  $CM_*$  and  $z$  to train the model. The relative importance of each parameter is shown in the left panel. Here the group total stellar mass clearly outperforms central stellar mass, further justifying this second stage in the halo mass estimation.

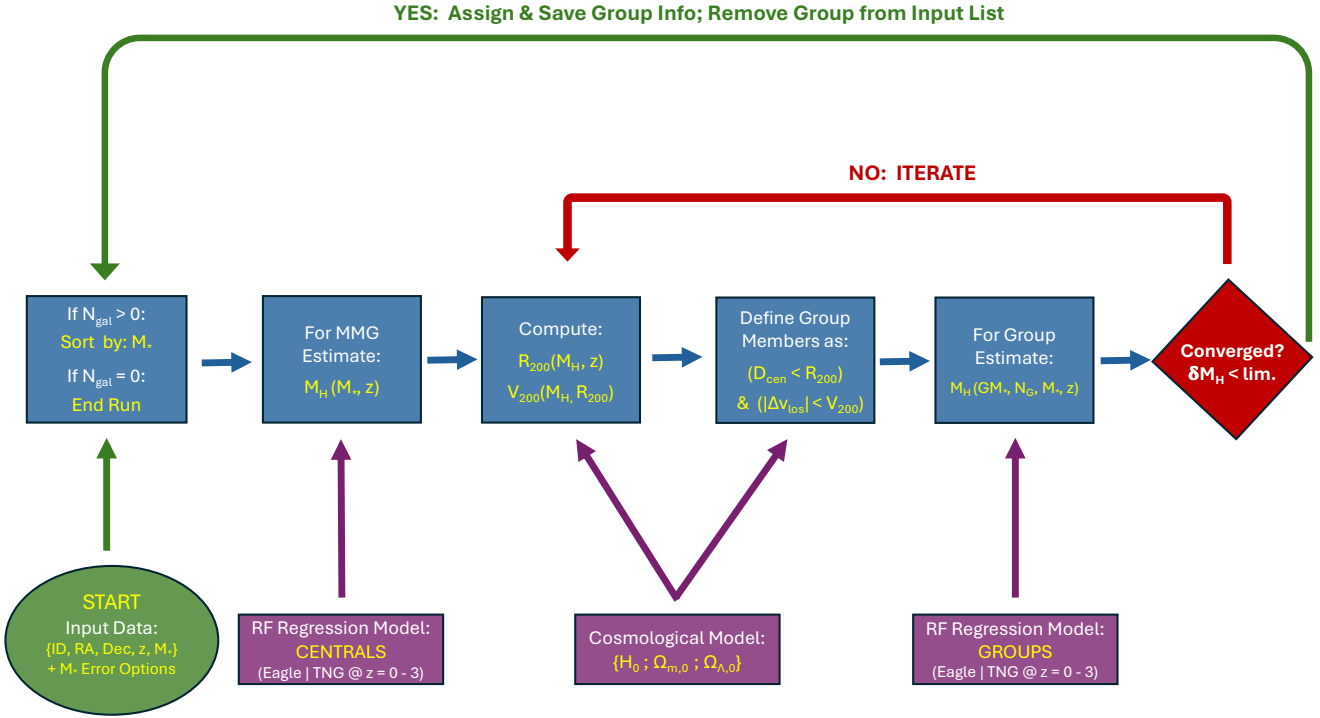
In the center panel we show the performance of halo mass recovery in the training data from TNG. This yields an essentially bias free result with a random uncertainty of  $\sigma = 0.1$  dex. On the right panel, we show the performance in unseen vali-

ation data from TNG. Compared to the results from centrals alone, we note an improved accuracy in halo mass recovery in the unseen data to  $\sigma = 0.12$  dex (TNG). More importantly, we find that the group RF model can reproduce the halo masses of both single and multi-galaxy groups with no bias, and achieves a much improved statistical uncertainty in the multi-galaxy case compared to the central RF model. Taken together, this strongly motivates us to incorporate group finding into the assignment of halo masses, which is the essential idea behind the DfL pipeline.

We perform an analogous training of group RF models for the EAGLE simulation. The plots from this simulation are not shown here for the sake of brevity, but the performance statistics achieved are presented alongside TNG in Table 1. Additionally, in Table 1 we compare the group performance from RF to an ANN approach, which demonstrates that the former is as successful as the latter for solving this problem. See Appendix C for full details on the ANN structure used here as a test.

### 3.5. The halo finder algorithm: DfL

The DfL halo finder pipeline is the main contribution of this work. It provides a fast, accurate, and straightforward method to assign galaxies from a wide-field spectroscopic survey into groups and clusters, while also estimating the dark matter



**Fig. 3.** Flow chart depicting the DfL pipeline, used to assign galaxies to haloes. The pipeline only requires as input data a unique galaxy identifier (ID), coordinates (RA, Dec), spectroscopic redshift ( $z$ ), and a stellar mass estimate ( $M_*$ ), plus optional uncertainties, for each galaxy in the survey. Additionally, one may specify the preferred cosmological parameters and the preferred simulation to use (here either TNG or EAGLE). First, the number of galaxies in the survey remaining to be assigned to a group is checked ( $N_{\text{gal}}$ ), provided this is greater than zero the remaining galaxies are sorted by stellar mass and the most massive galaxy in the survey is taken to be a central. Second, the halo mass of the central is estimated by applying the central galaxy regression model (see Section 3.4.1). Third, the virial radius and virial velocity of the halo are computed using the initial halo mass estimate, redshift, and the specified cosmology. Fourth, all galaxies within one virial radius and one virial velocity of the central galaxy are added to the proto-group. Fifth, the halo mass is re-estimated using the full regression model (see Section 3.4.2), here incorporating the total stellar mass of the group and the number of galaxies in the group with  $M_* > 10^{9.5} M_\odot$ , in addition to the central’s stellar mass and redshift. Finally, the two halo mass estimates are compared. If they are within a variable tolerance limit of each other (default value of  $\text{lim.} = 0.05$  dex), the group is fixed and removed from the input data set. The pipeline then continues from the start for the most massive galaxy remaining in the survey. If the estimates are not converged, the virial radius and velocity are updated with the new halo mass estimate and the group is redefined accordingly. Through iteration the halo may grow or shrink in phase space until convergence is reached (or else an optional limit is exceeded, default value:  $N_{\text{lim.}} = 10$ ). Once converged, the pipeline begins again from the start on the reduced data set until every single galaxy within the survey is classified (as either a central or a satellite) and the halo properties are provided for each group.

halo masses and virial radii, velocities, and temperatures of each group. The pipeline provides a classification of central and satellite galaxies, and hence also allows one to calculate the 2D physical distances of each group member to its central.

The pipeline takes as input only the following five parameters for each galaxy in the survey: (i) unique galaxy ID; (ii) RA; (iii) Dec; (iv) spectroscopic redshift ( $z$ ); and (v) an estimate of stellar mass ( $M_*$ ). Additionally, one may specify the uncertainty on the stellar mass as either a single number (which is interpreted as the  $1\sigma$  width of a Gaussian probability distribution) or else a full probability distribution for each stellar mass value. Furthermore, the desired cosmological parameters (explicitly,  $H_0$  and  $\Omega_{M,0}$ ) may be set by the user. The pipeline assumes a spatially flat  $\Lambda$ CDM cosmology, and hence  $\Omega_{k,0} = 0$  and  $\Omega_{\Lambda,0} = 1 - \Omega_{M,0}$ . The pipeline also makes use of the trained multi-snapshot RF regression models for centrals-only and for groups (outlined in Section 3.4).


We illustrate the logical flow of the DfL pipeline by a flow-chart in Fig. 3. Hereafter, in this section, we present each step in some detail. We recommend readers to follow along using the flow-chart to aid in understanding.

### 3.5.1. Step 1:

First, the number of galaxies remaining to be sorted into groups is checked. Provided this is greater than zero the pipeline begins (or continues, since it is a loop). Thereafter, the galaxies in the input catalog are sorted by stellar mass. The most massive galaxy in the survey is taken to be a central galaxy. Given that central galaxies are defined in this work to be the most massive galaxy in a given group, the most massive galaxy in the survey is highly likely to be a central.

There are only two ways this could be untrue. First, the most massive survey galaxy (the MMG) may be part of a group which extends beyond the survey limit and hosts an even more massive galaxy than the MMG. Second, the uncertainty on the stellar mass estimate could result in the actual stellar mass being lower than some other group member. We quantify both of these possible issues in what follows.

For the use on observational data, we recommend selecting a contiguous area of a wide-field survey with a buffer zone of at least 1 Mpc (physical) around the target region. The halo finder will target only galaxies within the inner region for halo assignment, but have access to the wider survey to search for potential neighbors. Through this method we have found that fewer than



**Fig. 4.** Dependence of virial radius ( $R_{200}$ , left panel), virial velocity ( $V_{200}$ , center panel), and virial temperature ( $T_{200}$ , right panel) on halo mass ( $M_{200}$ ) and cosmological redshift ( $z$ , see legends). Each panel shows results for four orders of magnitude in halo mass and across 13 Gyr of cosmic time (see legends). These panels illustrate how an estimated value of halo mass can be used to infer the physical and phase-space extent of groups and clusters. Additionally, these plots illustrate how various important halo properties may be inferred directly from halo mass estimates and a redshift measurement. The cosmological parameters used in this figure are as follows:  $\Omega_M = 0.3$ ,  $\Omega_\Lambda = 0.7$ ,  $H_0 = 70$  [km/s Mpc $^{-1}$ ]. For the virial temperature, we additionally assume a mean particle mass of  $0.6 m_p$ , which is correct for a primordial plasma. However, in reality this must be updated by constraints on the individual group or cluster metallicity, especially at lower redshifts.

1% of MMGs are erroneously identified as central galaxies due to edge effects. For the impact of stellar mass uncertainty on the overall results we defer to Section 4.5, where this is thoroughly analyzed and discussed.

### 3.5.2. Step 2:

For the most massive galaxy, we apply the central galaxy random forest regression model (see Section 3.4.1). This yields an estimate of halo mass, explicitly  $M_{200}$ , i.e., the total mass contained within  $R_{200}$  (see next step for further details). This part of the pipeline would function just as well with a simple nonlinear fit, or indeed with abundance matching. However, for consistency with what follows (and to provide the benefit of utilizing the individual tree predictions for error analysis) we also use random forest regression here.

### 3.5.3. Step 3:

Using the initial estimate of the dark matter halo mass, we utilize the spectroscopic redshift of the MMG and a chosen cosmology to infer the virial radius and velocity of the dark matter halo (explicitly,  $R_{200}$  and  $V_{200}$ ). This is an important step in the pipeline so we describe it in some detail here.

The dark matter halo mass is taken to be  $M_{200}$  in the simulations used for training the regression models. This is defined as the mass contained within the radius,  $R_{200}$ , at which the average density within this limit is 200 times the critical density of the Universe at this epoch (see, e.g., Mo et al. 2010). Explicitly,

$$M_{200} \equiv 200 \rho_{\text{crit.}}(z) \times \frac{4}{3} \pi R_{200}^3, \quad (7)$$

where from a simple rearrangement of the second Friedmann equation for cosmology, the critical density is given as

$$\rho_{\text{crit.}} = \frac{3H(z)^2}{8\pi G} = \frac{3H_0^2 (\Omega_{M,0} (1+z)^3 + \Omega_{\Lambda,0})}{8\pi G}, \quad (8)$$

where  $H(z)$  is the redshift dependent Hubble parameter. In the second step above we expand  $H(z)$  in terms of the current Hubble parameter value ( $H_0$ ), using the assumption of a spatially flat two component Universe, containing only matter (quantified by the  $z = 0$  contribution to the critical density of matter,  $\Omega_{M,0}$ ) and dark energy (quantified by the  $z = 0$  contribution to the critical density of dark energy,  $\Omega_{\Lambda,0}$ ). This allows the virial radius to be computed as

$$R_{200} = \left( \frac{GM_{200}}{100H(z)^2} \right)^{1/3} = \left( \frac{GM_{200}}{100H_0^2 (\Omega_{M,0} (1+z)^3 + \Omega_{\Lambda,0})} \right)^{1/3}. \quad (9)$$

We then compute the virial velocity as the circular velocity at  $R_{200}$  around the halo. Explicitly,

$$V_{200} = \sqrt{\frac{GM_{200}}{R_{200}}}. \quad (10)$$

In Fig. 4 we illustrate the relationships between halo mass and virial radius, velocity, and temperature for a variety of redshifts. We note that while the virial temperature ( $T_{200}$ ) is not used by our pipeline, it is an output parameter. This is computed from setting the thermal energy equal to the mean kinetic energy per particle in the plasma contained within the virial radius. Explicitly, the virial temperature is given as

$$T_{200} = \frac{\mu m_p V_{200}^2}{2K_B}, \quad (11)$$

where  $\mu$  is the mean atomic mass per particle. This is a free parameter in our pipeline, but we take as a default value  $\mu = 0.6$ , which is accurate for a primordial fully ionized plasma.

### 3.5.4. Step 4:

This step uses the derived halo properties of the previous step to identify group members. In the training of the regression models from simulations, we define the core group as

$$D_{c,3D} < R_{200}, \quad (12)$$

where explicitly

$$D_{c,3D} \equiv \sqrt{(X_i - X_c)^2 + (Y_i - Y_c)^2 + (Z_i - Z_c)^2} \quad (13)$$

and  $(X, Y, Z)$  indicate the coordinates of arbitrary galaxies in the survey (with subscripts,  $i$ ) and the central in question (with subscript,  $c$ ). Both are given in physical units, i.e., in the rest-frame of the simulation snapshot.

In observational, or observational-like data (which are used to test the pipeline), we measure the angular separation between any two galaxies using the haversine formula. Explicitly, this is defined as

$$\Delta\chi = 2 \sin^{-1} \left\{ \sqrt{\sin^2 \left( \frac{\theta_i - \theta_c}{2} \right) + \cos(\theta_i) \cos(\theta_c) \sin^2 \left( \frac{\phi_i - \phi_c}{2} \right)} \right\}, \quad (14)$$

where  $\theta$  indicates declination (Dec) and  $\phi$  indicates right ascension (RA), with subscript “ $i$ ” indicating any arbitrary galaxy in the survey and subscript “ $c$ ” indicating the central galaxy in question. We next convert this to a physical distance at the epoch of the MMG. Explicitly, we compute

$$D_{c,2D} = \Delta\chi D_A(z_c), \quad (15)$$

where  $D_A(z_c)$  is the angular diameter distance evaluated at the redshift of the central galaxy, which is defined as:  $D_A(z) = D_C(z)/(1+z)$ , where  $D_C(z)$  is the comoving radial distance. This is defined explicitly as

$$D_C(z) \equiv \frac{c}{H_0} \int_0^z \frac{dz'}{E(z')}, \quad (16)$$

where

$$E(z) \equiv \sqrt{\Omega_{M,0}(1+z)^3 + \Omega_{\Lambda,0}} \quad (17)$$

for a spatially flat  $\Lambda$ CDM universe.

We next compute the line-of-sight velocity difference between each galaxy and the central as a function of the difference in their observed redshifts. Explicitly, we compute

$$\Delta V_{\text{los}} = \frac{c|\Delta z|}{(1+\bar{z})} \text{ for } \Delta V_{\text{los}} \ll c, \quad (18)$$

where  $\Delta z = (z_i - z_c)$ , and  $\bar{z}$  is taken as the mean redshift of the two galaxies, rather than the (unknown) true cosmological redshift. We note that for physically associated galaxies, the velocity difference is always much lower than the speed of light and hence the above expression is sufficient for our purposes.

Using the above distances and velocity offsets, we define the core group in the pipeline as

$$(D_{c,2D} < R_{200}) \ \& \ (\Delta V_{\text{los}} < V_{200}). \quad (19)$$

The galaxies which meet these criteria are used in the next step to infer  $M_{200}$  on the basis of the total group stellar mass, the total number of galaxies within the group, and the central galaxy stellar mass and redshift (as before).

We also construct an extended group region within both the simulations and observational-like data, which will in principle contain both satellites and non-satellites. This is motivated in part by the literature on the splash-back radius, which demonstrates that satellite galaxies may extend beyond the virial radius out to  $\sim 2R_{200}$  (e.g., Adhikari et al. 2014; More et al. 2015, 2016). Additionally, the impact velocity from infinity at

the virial radius is  $V_{\text{imp}} = \sqrt{2} V_{200}$  (a classic result in Newtonian mechanics), and hence some true satellites will be moving at velocities exceeding virial. Moreover, three body interactions and nonvirial group mergers can both lead to enhanced super-virial velocities. Unfortunately, this region of phase space will also contain chance alignments of true centrals as well. We test the accuracy of central – satellite classification via this method in Section 4.

Explicitly, in simulations, we define the extended satellite region as

$$D_{c,3D} < 2R_{200}. \quad (20)$$

In the rare case of satellites in the simulations residing at  $D_{c,3D} > 2R_{200}$ , we treat these objects as independent centrals and utilize their sub-halo mass as their group halo mass. These cases ( $< 10\%$  of satellites) are likely not truly connected to the group at the snapshot in question, despite being identified as satellites via the SUBFIND algorithm. For example, in Goubert et al. (2024) we find that satellites beyond  $\sim 1.5 R_{200}$  behave identically to isolated centrals in both EAGLE and TNG across a wide variety of parameters, despite being identified as part of a group or cluster.

For observations (and observational-like simulated data), we define the extended satellite region as

$$(D_{c,2D} < 2R_{200}) \ \& \ (\Delta V_{\text{los}} < 2V_{200}). \quad (21)$$

This essentially provides a buffer region around each group. The extended satellite region is not used to infer halo masses directly, but it is used to further isolate each group from truly independent structures. Hence, the extended satellite population is removed with the core group after each completed iteration.

In summary of Step 4, based on the initial halo estimate we now have a proto-group defined, with both core and extended satellite populations, in addition to the central galaxy.

### 3.5.5. Step 5:

Utilizing the proto-group, i.e., the IDs of galaxies within the core group identified in Step 4, we evaluate the total group stellar mass and the total number of galaxies within the group (with  $M_* > 10^{9.5} M_{\odot}$ ). These parameters are then included with the central galaxy’s stellar mass and redshift as input to the group RF regression model (see Section 3.4.2). This yields a new estimate of  $M_{200}$ , now based on the group properties (not the properties of the central alone).

### 3.5.6. Step 6:

The two halo mass estimates (from central-only and the full group) are then compared. If they are sufficiently similar in value ( $|\delta M_{\text{Halo}}| < \zeta$ ) the group is fixed and the pipeline returns to Step 1, with this group removed from the input data. The tolerance limit ( $\zeta$ ) is a user variable parameter in DfL. Through direct experimentation we have found that an optimal value is reached between  $\zeta = 0.01$ – $0.05$  dex (depending on simulation and model pairing). However, no results are strongly dependent on this parameter within this range.

We note that after removing the first group (that of the most massive member of the survey), the pipeline is free to begin again with the next most massive galaxy. Provided this galaxy should not have been included with the first group, the next most massive galaxy is also highly likely to be a central by the same logic as outlined in Step 1. The pipeline will continue iterating until every single galaxy in the survey is assigned to a group, and

a halo mass (plus virial parameters) are estimated. We note that groups may contain only one galaxy, i.e., for isolated centrals. Indeed, this is the most common group type.

It is important to appreciate that it is highly unlikely that the next most massive galaxy ought to have been included in the previous group because it is (by definition) at a distance greater than two virial radii projected, and/ or at a velocity difference of greater than two virial velocities line-of-sight from the first group's central galaxy.

On the other hand, if the two halo mass estimates disagree by more than the tolerance, the pipeline returns to Step 3. That is, a new virial radius and virial velocity is inferred from the updated halo mass estimate. This allows a new definition of the group to be made and hence a further halo mass estimation, based upon the updated group membership information.


Once again, the two most recent halo mass estimates are compared, and if convergence is still not reached the pipeline iterates again back to Step 3. This is allowed to go on for up to  $N_{\text{lim}}$  iterations. This parameter is specified by the user. The default value is a maximum of ten iterations, which we find to be sufficient to reach convergence in  $\sim 99\%$  of cases. This means that the groups may grow or shrink in phase space based on the distribution of galaxies around the MMG. Once convergence is reached, the pipeline removes this group from the survey and returns to Step 1 to continue for the next group in the survey (as outlined above). Once the number of galaxies remaining in the sample reaches zero, the DfL pipeline ceases running and outputs the group catalogs with halo mass estimates.

### 3.5.7. Uncertainty on halo mass predictions

The RF technique is inherently probabilistic and hence offers a natural way to model and propagate uncertainties throughout the DfL pipeline. DfL enables three levels of error propagation from the input stellar masses, which we outline here:

Case 1: No stellar mass uncertainties. DfL does not require uncertainties on the input stellar masses in order to run. However, we strongly recommend to provide uncertainties in order to get the best results out of the pipeline. In the case of no input uncertainties, the uncertainties on halo mass are derived directly and solely from the individual decision trees within the random forest. A total of 250 decision trees are utilized in the determination of each halo mass estimate. Each tree is different due to bootstrapped random sampling of the training data set, and through random partial access to input features (in the group model). The normalized distribution of halo mass predictions is taken to be a probability distribution function (PDF) of the output halo mass. This is saved to a FITS file by DfL along with the point halo mass prediction (the geometric mean of the PDF) and the 1 and 2  $\sigma$  uncertainties, calculated in a nonparametric manner as the 16–84 and 2.5–97.5 percentiles of the data, respectively. This procedure quantifies the uncertainty inherent from the mapping from luminous tracers to dark matter halo mass in the simulation alone.

Case 2: Point stellar mass uncertainties. In the case of a single numerical value for the uncertainty on stellar masses, DfL treats this as the standard deviation of a Gaussian centered on the stellar mass input value. A variable number ( $N_{\text{MC}}$ ) of random draws from this Gaussian is taken during the processing of each galaxy. The default number of random draws is 100. The propagation of uncertainties is then implemented via a Monte Carlo approach. For each of the random draws in stellar mass a halo mass distribution function is created from the individual decision tree outputs (as in Case 1). All of the tree predictions



**Fig. 5.** Example of output halo mass PDFs from DfL from the TNG model applied to the TNG testing data at  $z = 1$ . The shaded histograms show the probability distributions of the halo mass for a randomly selected low-, intermediate-, and high-mass group (colored blue, green, and red, respectively). The true halo mass from the simulation is indicated by a vertical solid line, with the DfL geometric mean output from the pipeline indicated by a vertical dashed line. The PDF can be used to extract the mode, linear mean, or any other indicator for the best mass prediction for  $M_{200}$ , as desired.

from all of the random draws are concatenated, yielding an output distribution in halo mass predictions with  $N_{\text{trees}} \times N_{\text{MC}}$  (i.e., a default of 25 k) individual values. This is then normalized and treated as the PDF output for each halo mass estimate.

Case 3: Full stellar mass PDFs. DfL enables (and we strongly encourage) users to input full PDFs of stellar masses when available (see Appendix A for practical details). In the presence of a full PDF, the pipeline randomly samples the PDF exactly as in Case 2 for the simple Gaussian case, but now with the benefit of a more accurate characterization of the uncertainties on stellar mass. As in Case 2, an estimate of halo mass from each of the 250 trees in each RF model run is created for each of the  $N_{\text{MC}}$  random draws of the stellar mass PDF. These are then concatenated together and normalized to yield the final output PDF for each halo mass estimate, in addition to its default point value and uncertainties (exactly as in Case 1).

In Fig. 5 we show example output PDFs for the halo mass predictions, operating with a bespoke (randomly generated) input PDF for the stellar mass inputs. The dashed vertical lines indicate the DfL halo mass prediction and the solid vertical lines indicate the simulation truth. In this figure three random PDFs are drawn at low mass (blue), intermediate mass (green), and high mass (red) from TNG at  $z = 1$ .

### 3.5.8. Note on the structure of the pipeline

The pipeline is essentially powered by two nested while loops programmed in PYTHON-3. The first loop operates while there are unclassified galaxies into groups within the survey (illustrated by the green arrow in Fig. 3). Within this, the second loop operates while there is a lack of convergence in the halo mass estimators for a single proto-group (illustrated by the red arrow in Fig. 3). In conjunction, the pipeline takes a relatively simple set of observational input data and assigns a physically motivated group structure to the entire galaxy survey. In the remainder of

this paper we carefully test the performance of this method for estimating the dark sector in observational-like simulated data. In Appendix A we provide a user guide to DfL, enabling the straightforward application to wide-field galaxy surveys.

### 3.6. Observational-like simulated data: The testing sample

In order to test the DfL pipeline, we construct an observational-like data set from each simulation utilizing an unseen testing sample. The testing sample is defined as all data for which:  $X > 50$  cMpc. This represents approximately half of the TNG and EAGLE simulations. Crucially, this subsample is not used in either training or validating the RF models for DfL. These data are completely unseen by the machine learning prescriptions prior to testing. Moreover, we also transform the testing sample into an observational-like sample by restricting the dimensionality of the data from 6D phase space to 2D +  $z$  observer space (discussed in detail below). Hence, the purpose of the testing sample is: (i) to have an independent data set unseen by the machine learning models in training and validation, with which to test the performance of DfL; and (ii) to have an observational-like data set with which to test the impact of observational limitations on the recovery of halo masses with DfL.

For the testing sample, we randomly project each simulation snapshot into a 2D plane (hereafter labeled as  $X&Y$ ). The  $X$  and  $Y$  coordinates are then converted to sky coordinates by applying the angular diameter distance and placing onto the sky at an appropriate location. Explicitly, we position the edge of each simulation snapshot 2D plane ( $X = Y = 0$ ) at  $RA = 0$  and  $Dec = 0$ .

We then make the simplifying assumption that the sky is flat on the small angular scales subtended by the relatively small simulation boxes (when placed near the celestial equator). We have tested that at the scale of the most massive clusters this simplification incurs  $<1\%$  distance errors at the lowest redshift positioning (taken as  $z = 0.1$ )<sup>8</sup>. At higher redshifts the difference becomes completely negligible.

Explicitly, we compute the following mappings:

$$\theta_i = \theta_0 + \frac{Y_i}{D_A(z_i)}, \quad (22)$$

$$\phi_i = \phi_0 + \frac{X_i}{D_A(z_i) \cos(\theta_i)}. \quad (23)$$

Here  $\phi$  represents RA and  $\theta$  represents Dec. In practice we set the zero points to zero, and the cosine function evaluates to very close to unity. As mentioned above, this assumes orthogonality between sky coordinates on the scales of relevance in the simulations. This is a very reasonable assumption in practice here due to the relatively small comoving box size of the simulations, though (of course) this is not true in general<sup>9</sup>.

One advantage of the simplifying approach outlined above is that the  $Z$  coordinate may now be taken as perpendicular to

<sup>8</sup> Obviously, one cannot literally place a galaxy at  $z = 0$ , since this would mean it resides on the Milky Way. To combat this technicality we position the  $z = 0$  snapshot data at  $z = 0.1$ . This is only used to allow the conversion from simulation to observer-like distances and velocities, and is not used to modify the cosmological parameters. Hence, more concretely, we test the  $z = 0$  epoch at a  $z = 0.1$  distance. No changes are needed for the other redshift snapshots used in this work.

<sup>9</sup> More quantitatively, we note that the Haversine formula (Eq. (14)) reduces to:  $\Delta\chi \approx \sqrt{(\theta_2 - \theta_1)^2 + (\phi_2 - \phi_1)^2 \cos^2(\langle\theta_{12}\rangle)}$ , using the small angle formula. At the celestial equator the cosine term equals unity and orthogonality is achieved. Using Eqs. (22) & (23), this yields:  $D_{c,2D} = \Delta\chi D_A(z) = \sqrt{X^2 + Y^2}$ , as desired.

the ‘sky plane’. We use this to construct a pseudo cosmological recession velocity (and hence redshift), which is then added to the real peculiar line-of-sight velocity in the simulation to construct an observational-like redshift. Explicitly, we construct the observer-like cosmological redshift as

$$z_{\text{cos}} = z(D_c(z_{\text{snapshot}}) + Z_i(1 + z_{\text{snapshot}})), \quad (24)$$

where we evaluate redshift as a function of comoving radial distance by numerically solving the defining integral equation (see back to Eq. (16)). The true comoving distance is taken as the actual comoving distance to the snapshot redshift plus the comoving radial coordinate of the galaxy. We note that  $Z_i$  is given in physical coordinates, which is why it is multiplied by the  $(1+z)$  factor. We take this redshift to be equal to the snapshot redshift because: (i) this is approximately true; and (ii) the cosmological redshift is not defined prior to this computation. The above expression thus accounts for the positioning of galaxies within the simulation box in the assignment of a cosmological redshift to each galaxy.

Next, we add in the contribution from the line-of-sight peculiar velocity of the galaxy, to better approximate a true observer-like scenario. Explicitly, we compute the observation-like redshift as

$$z_{\text{obs}} = z_{\text{cos}} + z_{\text{pec}}, \quad (25)$$

where

$$z_{\text{pec}} = \frac{V_Z(1 + z_{\text{snapshot}})}{c}, \quad (26)$$

and  $V_Z$  is the  $Z$ -component physical velocity of each galaxy in the rest-frame of the simulation snapshot’s redshift. We note that when placed into Eq. (18) (used in the pipeline, see above), for any two galaxies with the same cosmological redshift, Eq. (26) yields the rest-frame velocity difference, as desired. This is accurate provided  $z_{\text{obs}} \approx z_{\text{cos}} \approx z_{\text{snapshot}}$ , which is invariably the case in these data.

Finally, the actual stellar mass of each galaxy is given to the pipeline as an input. Initially, this is used to test the recovery of halo masses under a perfect knowledge of the stellar component. However, later in the next section we test the impact of varying uncertainty in stellar mass estimates, in order to assess halo recovery under realistic galaxy survey conditions. Additionally, we also test the impact of incompleteness in galaxy detection on the recovery of halo masses (see Section 4.5).

Crucially, no information on the halo masses of galaxies, or membership of galaxies to groups, is provided to the pipeline. Nonetheless, this information is used post analysis to check the performance of the pipeline in estimating these properties. This is why we test the pipeline on observational-like simulated data (where the true halo properties are known) instead of real observational data (where the true halo properties are unknown). That said, we do additionally test the recovery of the observed  $M_* - M_{\text{Halo}}$  relationship in DfL from gravitational lensing data later (see Section 4.4).

Finally, for testing, we remove edge-effects from the observation-like simulated data by selecting galaxies which reside at least 1 cMpc from the nearest edge of the survey (which represents a distance approximately equal to the largest cluster virial radius at all epochs). We recommend that observational applications mimic this approach. This essentially prevents edge effects from impacting the results of the halo finder pipeline. We note that galaxies beyond this limit may be used to infer group properties, but these are only trusted if the central lies within the inner contiguous region (and hence all group members may in principle be identified).

## 4. DfL performance results

In this section we test the performance of the DfL halo finder pipeline on observational-like simulated data from TNG and EAGLE at example redshifts of  $z = 0, 1, 2$ . First, we consider the performance of the pipeline when using the appropriate model and data pairing (Section 4.1). In addition to quantifying the precision of halo mass estimates under different circumstances, we also quantify the accuracy of central – satellite classification and determine the fidelity of the overall halo mass function as recovered through the pipeline. Second, we consider the impact of deliberately applying a different model to the data to test the impact of model choice on the recovery of halo masses (Section 4.2). Third, we compare the performance results from DfL to abundance matching applied to the groups found from a friends-of-friends algorithm (Section 4.3). Fourth, we compare the output  $M_* - M_{200}$  relation from DfL at  $z = 0$  to observational constraints from strong gravitational lensing (Section 4.4). Finally, we test the impact on the recovery of halo masses with DfL of systematically varying stellar mass uncertainty and survey incompleteness (Section 4.5).

### 4.1. Single model performance

In Fig. 6, we present the recovery of halo masses from the DfL halo finder pipeline method. Density contours of true  $M_{200}$  values, from the simulations, are plotted against the estimated  $M_{200}$  values output from the pipeline. Due to the steepness of the halo mass function, the density contours are dominated by lower-mass haloes. To combat this issue, we additionally plot as star-shaped points (with error bars) the location of high-mass groups and clusters (explicitly those with  $N_g > 5$ ).

The top row of Fig. 6 shows results for TNG (using the appropriate TNG models) and the second row shows the results for EAGLE (using the appropriate EAGLE models). We note, however, that different data are used in the training and validation of the RF models and in the testing of DfL, shown here (as explained in Section 3.4). The lower two rows show cross-validation between models, which is discussed later in Section 4.2. In that section, completely different simulations (in addition to samples) are used in training the RF models and in testing the performance of DfL. Results are shown separately for  $z = 0, 1, 2$  (ordered by column). Dashed red lines indicate the one-to-one relationship, with dotted red lines indicating  $\pm 1$  dex from the one-to-one relation, to help indicate scale.

Visually, the recovery of halo masses is very good in the top two rows. The contours are well centered on the one-to-one relationship and clearly reasonably tight. We quantify the strength of the relationship with the Spearman rank correlation ( $\rho$ ), displayed in each panel. The average correlation strength is very high, at  $\langle \rho \rangle = 0.91$ , indicating an excellent recovery of halo masses from the pipeline on average. Additionally, we show the fraction of catastrophic outliers ( $f_{\text{out}}$ ), cases where the discrepancy is an order of magnitude or more. These are very rare occurrences ( $\langle f_{\text{out}} \rangle < 0.01$ ), but are important to keep track of in the overall performance. We also show the correlation strengths in parenthesis for the  $\sim 99\%$  of the data which is not an outlier. Here the correlation strength rises to  $\langle \rho \rangle = 0.94$ .

In Fig. 7, we show the distributions in the offsets in halo masses, defined explicitly as

$$\delta M_{\text{Halo}} \equiv M_{200}(\text{True}) - M_{200}(\text{Predicted}). \quad (27)$$

In each panel we present the bias ( $b$ , computed as the median value of  $\delta M_{\text{Halo}}$ ), the  $1\sigma$  uncertainty ( $\sigma_1$ , spanning the 16th–

84th percentile of the PDF), the  $2\sigma$  uncertainty ( $\sigma_2$ , spanning the 2.5th–97.5th percentile of the PDF), and the catastrophic outlier fraction ( $f_{\text{out}}$ , defined above).

Considering just the top two rows, where the same model is used for each data set, we see that the halo finder recovers halo masses with essentially no bias ( $|b| \leq 0.02$  dex) and an average uncertainty across redshifts of  $\langle \sigma_1 \rangle (\langle \sigma_2 \rangle) = 0.12 (0.28)$  dex. This is an extremely good performance, which substantially improves on the current status quo in halo mass estimation for wide-field galaxy surveys (e.g.,  $\sigma_1 \sim 0.2\text{--}0.5$  dex in Yang et al. 2007, 2009; Woo et al. 2013).

In Fig. 8, we show distributions of the offsets in estimated halo masses ( $\delta M_{200}$ ) for the full group model from DfL (shaded histograms) and the central-only model (open black histograms). The top panel shows results for single galaxy ‘groups’. The bottom panel shows results for multi-galaxy groups (also including clusters). In both panels the TNG data are shown as the top row and the EAGLE data are shown as the bottom row.

For single galaxy groups, the group and the central models yield qualitatively very similar results, as expected. However, for the multi-parameter groups, the central galaxy model is significantly biased to lower halo mass predictions (higher values of  $\delta M_{200}$ ) and has increased uncertainty, relative to the group model. This highlights the clear value in incorporating group information into the estimation of the halo masses within the DfL pipeline. Typically, the halo masses are biased by  $\sim 0.1$  dex and there is an increase in uncertainty of 20–35%.

In Fig. 9, we present the halo mass functions for the actual simulations (shown as shaded regions) compared to those output from the DfL halo finder pipeline (shown as red stars with error bars). This figure is structured as in Figs 6 & 7, whereby different redshifts are shown along each column and different data – model pairings are shown along each row. Here we discuss the first two rows, which show correctly paired models and data.

Visually, there is excellent agreement between the actual and predicted halo mass functions for each snapshot and in both simulations. We quantify the similarity between the halo mass functions with the mean offset in number densities and the standard deviation in number density offsets (both presented in each panel).

Finally, in Fig. 10 we present the confusion matrices for central – satellite classification from the halo finder pipeline. The confusion matrix plots true centrals and satellites against predicted centrals and satellites, resulting in four quadrants (from top left to bottom right): (i) true predicted satellites; (ii) false predicted centrals; (iii) false predicted satellites; and (iv) true predicted centrals. Hence, a perfect classification would have all data on the diagonal. The fraction of data in each quadrant is indicated by the darkness of each rectangular region, as indicated by the color bar. Additionally, we present the accuracy, precision, and recall of central – (core) satellite classification in each panel<sup>10</sup>.

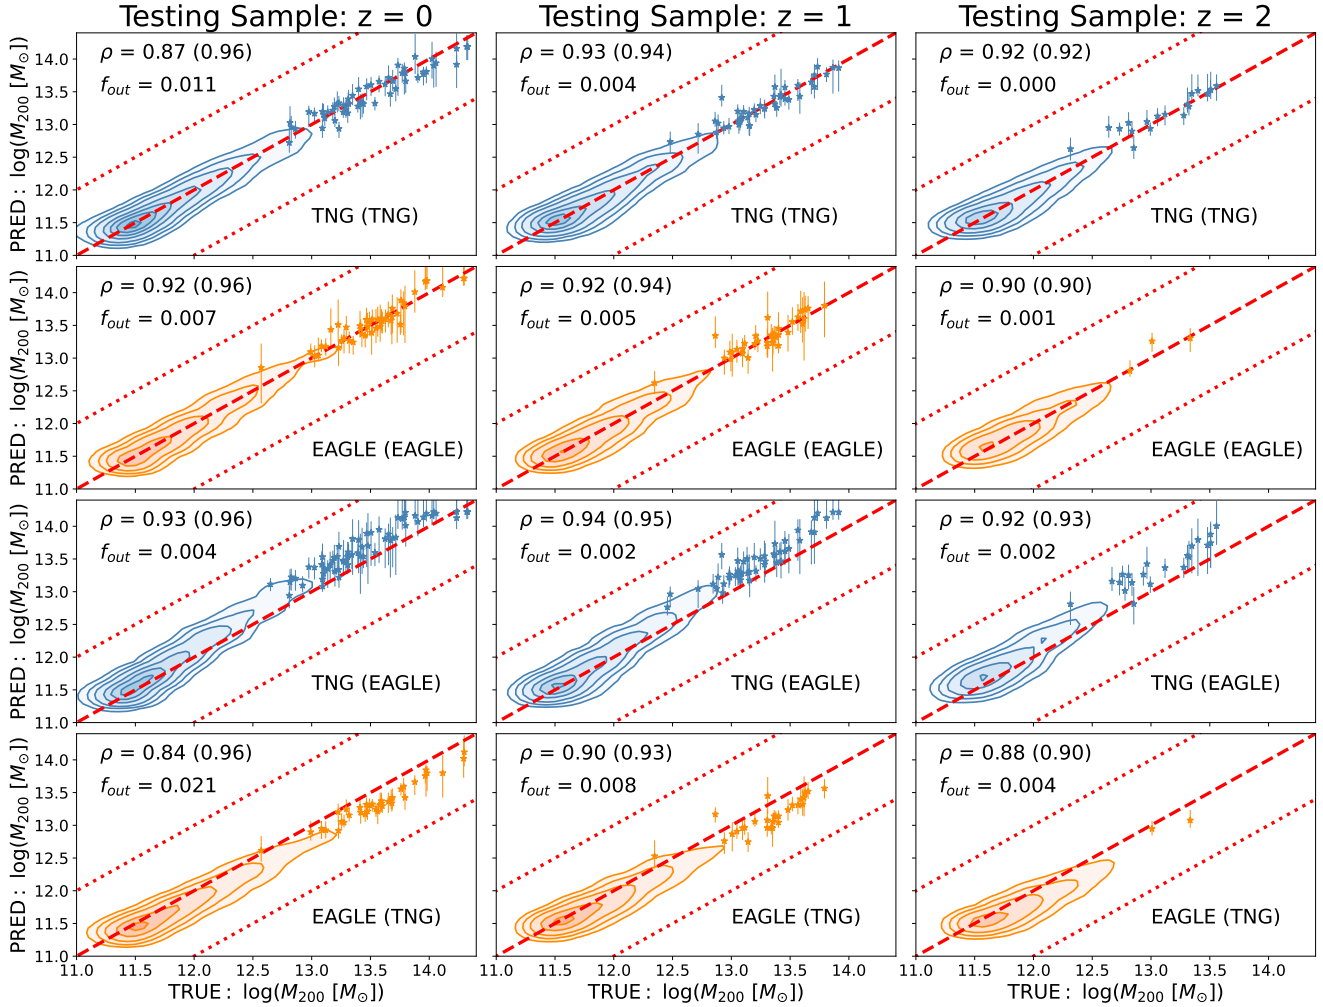
These statistics are defined as

$$\text{Accuracy (Acc.)} = \frac{N_{\text{True Centrals}} + N_{\text{True Satellites}}}{N_{\text{Total Classified}}}, \quad (28)$$

$$\text{Precision (Prec.)} = \frac{N_{\text{True Centrals}}}{N_{\text{True Centrals}} + N_{\text{False Centrals}}}, \quad (29)$$

$$\text{Recall (Rec.)} = \frac{N_{\text{True Centrals}}}{N_{\text{True Centrals}} + N_{\text{False Satellites}}}. \quad (30)$$

<sup>10</sup> For results incorporating the buffer zone of extended satellites, see Fig. B.6, where they are discussed in comparison to the FOF-AM technique.



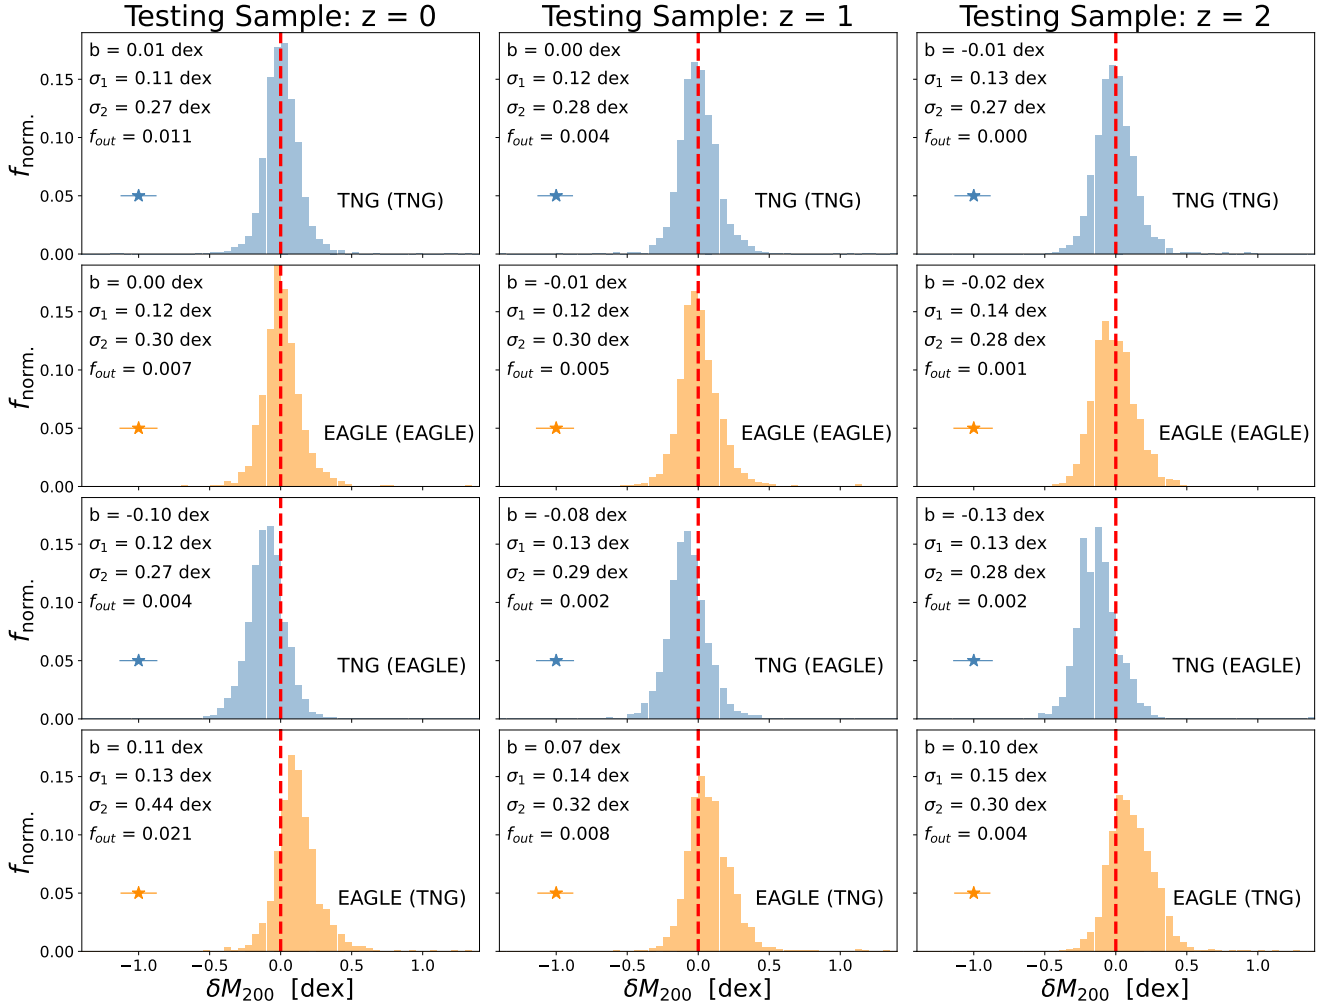
**Fig. 6.** Halo mass recovery performance with DfL I: DfL Predicted vs. simulation truth  $M_{200}$  density contours for all output galaxy groups (including isolated galaxies). In addition to contours, we also show as star-shaped points (with error bars) the location of the largest groups and clusters at each epoch (those with  $N_g > 5$ ). This allows one to assess the recovery of this small but important population, which are not well represented by the density contours, due to these being dominated by the most numerous lower-mass haloes. The figure is structured as follows: Columns indicate redshift snapshots (as labeled at the top), and rows indicate data – model pairings (as indicated in the panels). The top row shows results for TNG data using the TNG regression model (trained and validated on different subsets of the data). The second row shows the equivalent results for EAGLE data using an EAGLE trained model. The third and fourth rows show the results for TNG data using an EAGLE model and vice versa, respectively. In each panel the one-to-one relation is indicated by a dashed red line, and  $\pm 1$  dex is indicated by dotted red lines to help indicate scale. Also in each panel the Spearman correlation strength ( $\rho$ ) and the fraction of catastrophic outliers ( $f_{out}$ , i.e., those with  $|\delta M_{Halo}| > 1$  dex) are shown. In parentheses, we additionally show the correlation strength excluding the very rare case of catastrophic outliers. We note that visually the recovery of halo masses from luminous tracers is excellent when the appropriate model and data pairing is used, but that a modest bias is induced when the model and the data are mismatched.

The accuracy is intuitive, but the precision and recall can be a little confusing. Essentially, the recall measures the completeness of the output sample and the precision measures its purity. Optimizing for precision often worsens the recall, whereas optimizing for recall often worsens precision. Therefore, our methodology is to optimize the overall accuracy, which is the best compromise between recall and precision.

Excluding the extended satellite region, and focusing on the matched models and data, the group finder pipeline achieves an excellent accuracy of  $\langle \text{Acc} \rangle = 96\%$ . This indicates that just 4% of galaxies are misclassified as centrals or satellites in the pipeline within one virial radius of each identified central galaxy. Including the extended satellite region (i.e., galaxies within  $1 < D_c/R_{200} < 2$  and  $1 < \Delta V_{los}/V_{200} < 2$  of their central) the accuracy is significantly lowered to  $\langle \text{Acc} \rangle = 88\%$  (see

Fig. B.6). But we emphasize that this region is flagged and hence may be excluded from any given analysis.

The extended satellite region is a highly complex part of the parameter space, where true satellites and true centrals mingle. This region contains splash-back satellites, which have reached beyond the virial radius having previously been inside it (see, e.g., Adhikari et al. 2014; More et al. 2015), satellites falling into the group/ cluster for the first time (which can have  $|\Delta V_{los}/V_{200}| \leq \sqrt{2}$ ), and genuine central galaxies which are independent from the group, though close to it. This is further exacerbated by this extended phase space cut being more prone to chance alignments of galaxies which are at large radial distances from the central, but at similar line-of-sight radial velocities. Hence, this region of parameter space is inevitably going to lead to more uncertainty in central – satellite classification.



**Fig. 7.** Halo mass recovery performance with DfL Iia: Distributions in the offset from true-to-predicted  $M_{200}$  values ( $\delta M_{200} \equiv M_{200}(\text{true}) - M_{200}(\text{predicted})$ ). The structure of this figure is identical to Fig. 6. In each panel a dashed red line shows the perfect regression value of  $\delta M_{200} = 0$ , and the mean uncertainty on halo masses output from DfL is indicated by a horizontal error bar. Additionally, the bias (median offset), full  $1\sigma$  dispersion ( $\sigma_1$ , containing 68% of the data set),  $2\sigma$  dispersion ( $\sigma_2$ , containing 95% of the data), and the fraction of catastrophic outliers (those with  $|\delta M_{200}| > 1$  dex) are indicated in each panel. For the cases of data and models being matched (top two rows), the pipeline recovers halo masses with essentially no bias ( $\langle b \rangle \leq 0.02$  dex) and an average uncertainty of  $\langle \sigma_1 \rangle = 0.12$  dex. To test the impact of model uncertainty, in the lower two rows we show the  $\delta M_{200}$  distributions for TNG data with an EAGLE model, and vice versa. Here an average bias of  $\langle b \rangle = -0.10$  dex is found by applying the EAGLE model to TNG data, and an average bias of  $\langle b \rangle = +0.10$  dex is found by applying the TNG model to EAGLE data. The overall random uncertainty remains essentially unchanged.

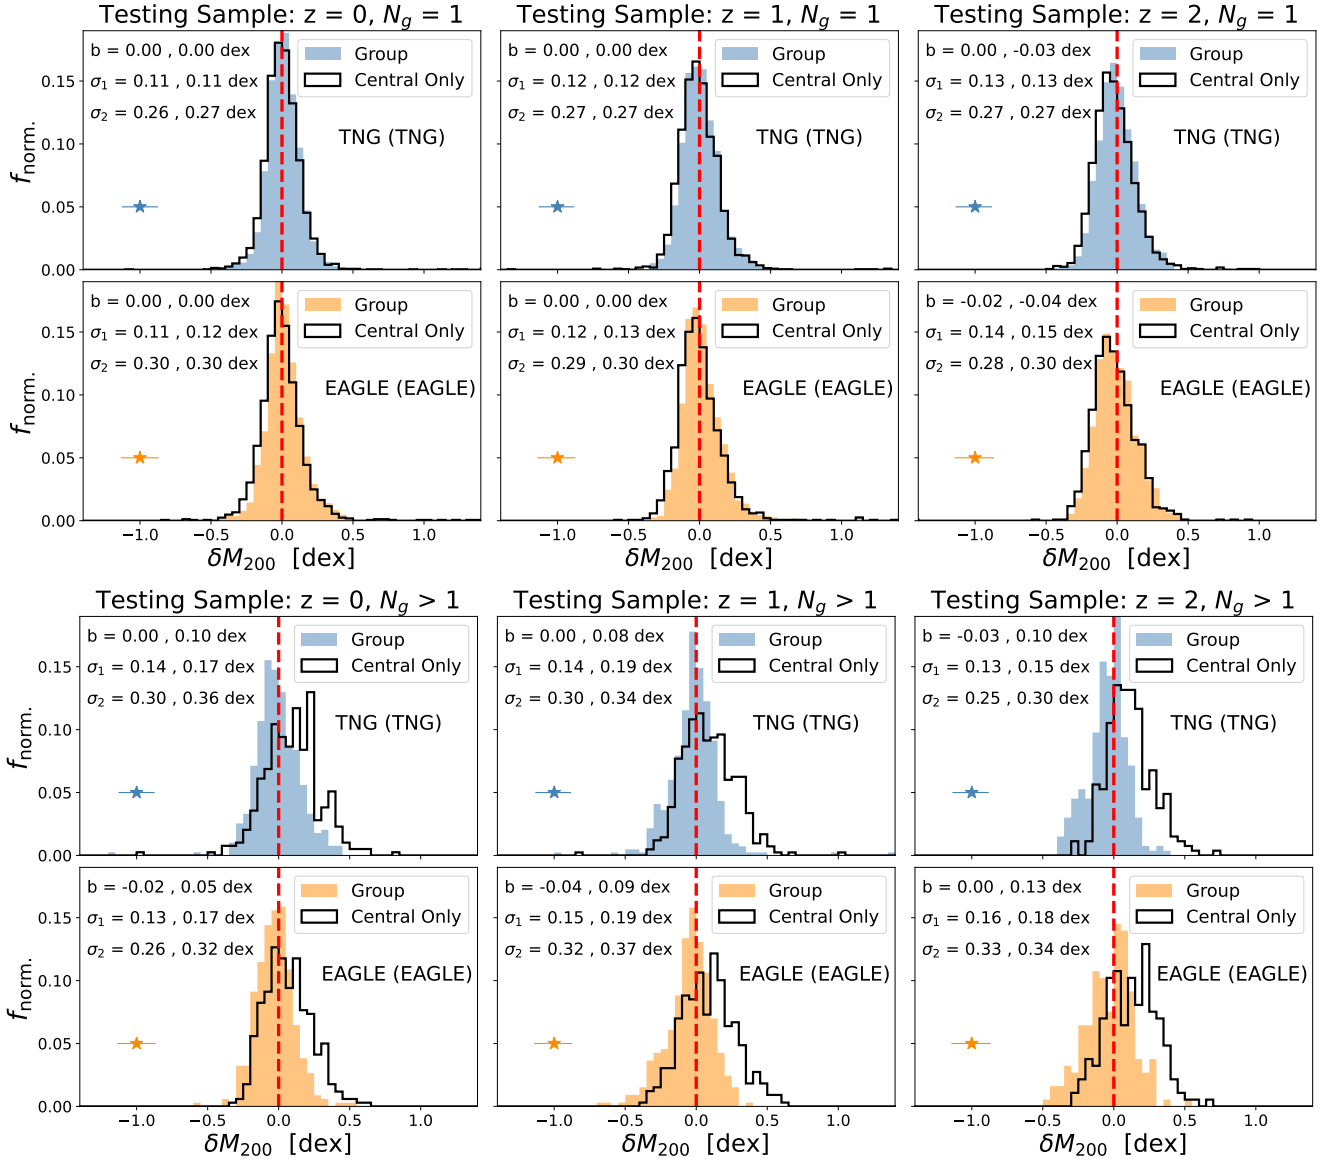
We have tested treating this region as all centrals or as all satellites, and the overall accuracy is essentially identical in both cases. We have also tried reducing, and increasing, the buffer zone region and here again the overall accuracy of central – satellite classification is not significantly impacted. Hence, we conclude that this is an unavoidable uncertainty resulting from restricting the true 6D phase space information (in simulations) to observational 2D +  $V_{\text{los}}$  space. In Section 4.3 we compare to another approach (abundance matching applied to FoF groups) and find similar issues.

As a result of the above tests, we recommend users of the halo finder pipeline to treat the extended satellite region with caution. For applications where purity in central – satellite classification is more important than completeness, we recommend to exclude the extended satellite region from the analysis. This will yield a very high central – satellite classification accuracy. Conversely, for applications where completeness is more important than purity, we recommend to include the extended satellite population, but to appreciate that this results in a  $\sim 8\%$

reduction in overall accuracy. In total  $\sim 20\%$  of the full sample (and  $\sim 50\%$  of all satellites) are in the extended satellite region.

#### 4.2. Cross validation between models

In the previous subsection we considered the performance of the DfL halo finder pipeline in the case where the random forest regression models are trained on the same simulation as used for the test data (although, crucially, not the same data set or data format; see Section 3.4). This would be adequate if a given model was known to be correct in all aspects. However, in reality the true underlying galaxy formation model of the Universe is unknown. Therefore, it is essential to attempt to quantify the uncertainty on halo mass recovery induced by the choice of model. To achieve this, we apply the EAGLE model to TNG data, and vice versa. This enables us to determine a preliminary estimate of the impact of the model choice on the final halo mass estimates, and group membership.




**Fig. 8.** Halo mass recovery performance with DfL IIb. Distributions in the offsets from true-to-predicted  $M_{200}$  values ( $\delta M_{200}$ ), shown here for single galaxy groups (top panels) and multi-galaxy groups (bottom panels). In each panel the distributions in halo mass offsets are shown for the full group RF model (shaded histograms) and compared to the central-only RF model (open black histograms), both of which are outputs from DfL. Statistics are presented in each panel for the group model, followed by the central only model. In the case of single galaxy groups, the two distributions are very similar, as expected. However, for multi-galaxy groups, the two halo mass offset distributions are significantly different. The central-only RF models yield halo masses which are biased to higher values of  $\delta M_{200}$  (lower values of  $M_{200}$  predicted) and have a higher dispersion relative to the group results. Moreover, the full group RF models have significantly better agreement with the true halo masses from each simulation. This demonstrates the critical importance of incorporating the group information into the DfL pipeline and hence justifies the need for the two stage approach, discussed in Section 3.

We present the cross-validation of mismatched model and data pairings as the bottom two rows in Figs. 6, 7, 9 & 10. The third row indicates the EAGLE model applied to TNG observational-like data, and the bottom row indicates the TNG model applied to EAGLE observational-like data.

Unlike in the case of matched models and data, in Fig. 6 we see that the contours (and data points) are slightly offset from the one-to-one relationship. The EAGLE model applied to TNG data systematically over-estimates the true values of halo masses, whereas the TNG model applied to EAGLE data systematically under-estimates the true values of halo masses. Nonetheless, the correlation strengths remain extremely high and the fraction of

catastrophic outliers remains very low (see values in lower panels in Fig. 6).

More quantitatively, in the lower half of Fig. 7 we display the bias and error from applying a mismatched model to the simulated observational-like data. From applying the EAGLE model to TNG data, we find an average bias of  $\langle b \rangle = -0.10$  dex. Conversely, for the TNG model applied to EAGLE data we find an average bias of  $\langle b \rangle = +0.09$  dex. This implies that the choice of model results in a systematic uncertainty of  $\sim \pm 0.1$  dex on the halo mass estimates. The random error (computed as the  $\sigma_1$  statistic) is essentially unchanged by applying mismatched models ( $\langle \sigma \rangle = 0.13$  dex). This indicates that the model uncertainty



**Fig. 9.** Halo mass recovery performance with DfL III: Halo mass functions. This figure is structured as in Figs. 6 and 7. In each panel the actual halo mass function for the simulation at each epoch (shaded regions) and the estimated halo mass function from the DfL pipeline (data points with errors) are compared. In the case of applying the same model and data pairing (albeit trained and tested on different subsamples), there is excellent agreement between the two distributions. In the case of applying different model and data pairings, there is more noticeable disagreement, originating from the differences in the true halo mass functions between simulations. Therefore, one can recover the halo mass function as it would be in either simulation from observational-like data to high fidelity, but the output halo mass function will be biased by the model choice. Additionally, in each panel we show the mean offset and standard deviation in offsets between the mass functions in order to quantify the statistical difference, and to facilitate comparison to the FOF-AM approach (see Section 4.4).

results in a systematic bias, but very little increase in random uncertainty after this is taken into account.

Of course, exploration of further galaxy formation models may result in different biases in principle. Indeed, it is our intent to add to the usefulness of DfL as a tool by incorporating more simulations into later versions, enabling a greater sampling of the contemporary model parameter space. However, for now, we note that EAGLE and TNG are very different models in terms of the baryonic physics applied, especially in terms of feedback (see Section 2). Hence, EAGLE and TNG provide a good baseline comparison for exploring how significant variation in the baryonic subgrid physics of feedback is on the estimation of halo masses from luminous tracers.

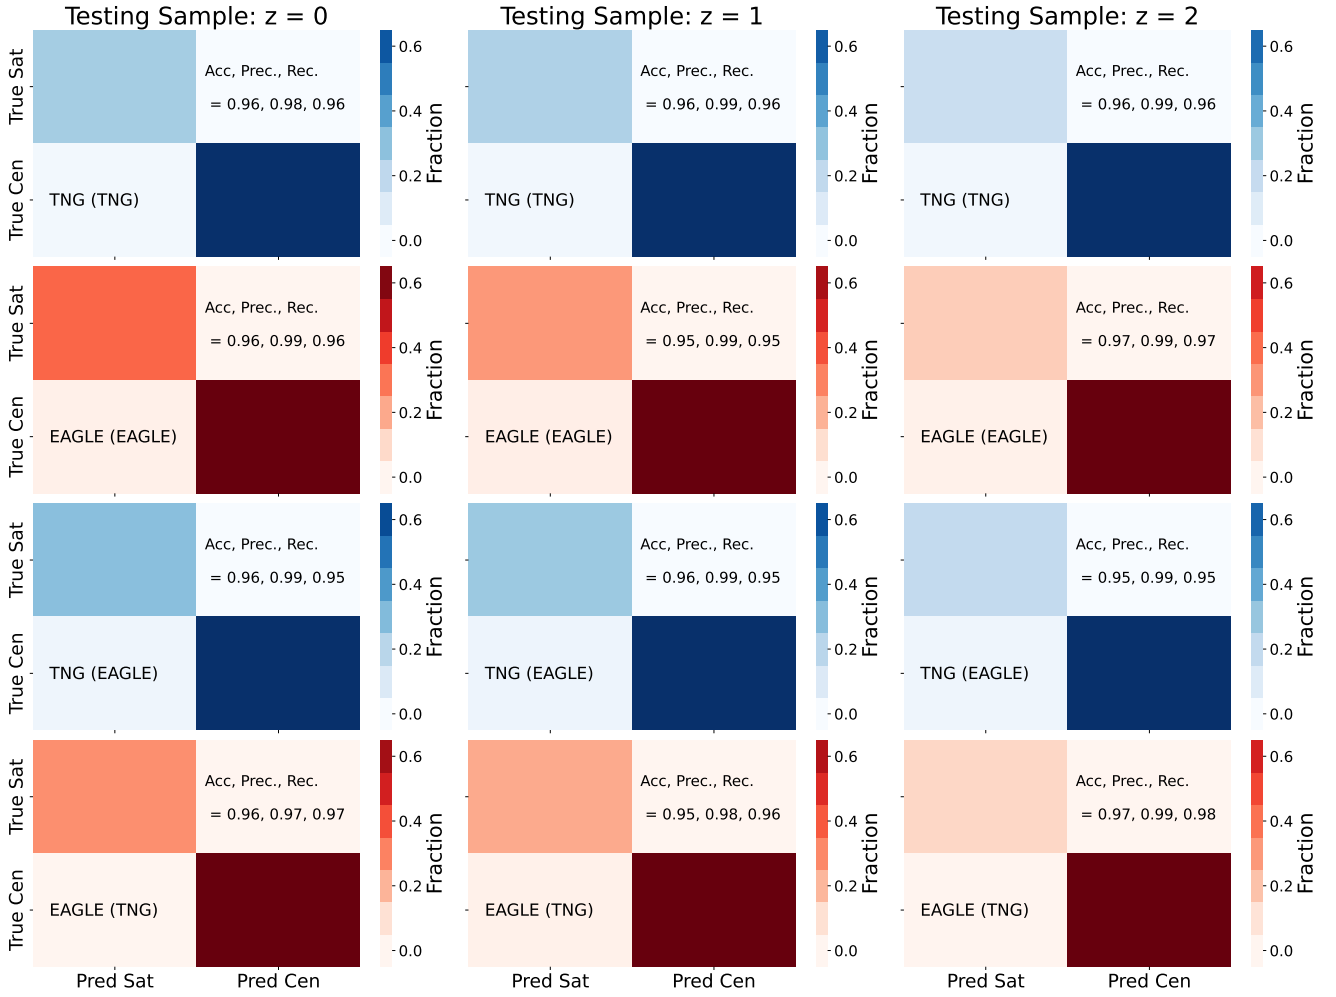
In the lower half of Fig. 9, we compare the estimated halo mass functions from the DfL halo finder pipeline (trained with mismatched models) to the true halo mass functions from the simulations. Here the results are quite instructive. While the halo mass function recovery is certainly not unreasonable, it is noticeably worse than in the case of matched model and data pairings. The EAGLE model tends to over-estimate the number densities of high-mass haloes relative to TNG, whereas the TNG

model tends to under-estimate the number densities of high-mass haloes.

Finally, in the lower half of Fig. 10, we show the performance of the halo finder on central – (core) satellite classification for the case of mismatched models and data. The accuracy in central – satellite classification, both including and excluding the extended satellite population, is extremely similar in the mismatched case to the matched case. This implies that the identification of galaxy classes is not strongly impacted by the galaxy formation model choice. Ultimately, this is logical because the impact of the model will only (slightly) bias the halo mass estimates, and not impact directly the relative location of centrals and satellites, or their stellar masses. It is the latter which is most critical for determining group membership in most cases.

#### 4.3. Benchmarking: performance comparison with FOF-AM

In this subsection we compare the performance of DfL to a more established set of techniques. More explicitly, we combine friends-of-friends group finding (FOF; see, e.g., Davis et al. 1985; Springel et al. 2001, 2005; More et al. 2011;



**Fig. 10.** Central – (core) satellite classification results from the DfL pipeline. The structure of this figure is identical to Figs. 6, 7, and 9. Each panel displays the confusion matrix, which plots true centrals and satellites against predicted centrals and satellites, with the color of each rectangular region indicating the fraction of data in each quadrant (as labeled by the color bars). The sum of all quadrants equals unity. A perfect classification would have all data on the diagonal in each panel, or equivalently dark colored rectangles on the diagonal and white space in the off-diagonal regions. Three statistics are indicated in each panel (see Eqs. (28)–(30) for definitions): the accuracy in separation between central and satellite galaxies (Acc.), the precision in identifying centrals (Prec.), and the recall (Rec.), which quantifies the completeness of the recovered central population. All results are shown here for centrals compared to the core satellite population. The average accuracy across all epochs is 96%, and this is largely unaffected by model choice.

Duarte & Mamon 2014; Tempel et al. 2016) with abundance matching (AM; see, e.g., Vale & Ostriker 2004; Moster et al. 2010; Behroozi et al. 2010; Hearin et al. 2013) in order to make halo mass predictions on the basis of total group stellar mass. This approach has much precedence in observational applications (see especially Yang et al. 2007, 2009).

In order for a fair comparison we must conduct abundance matching with total group stellar mass, not central stellar mass. This is because we have already established that for multi-galaxy haloes information on the group stellar mass yields superior results to the central galaxy alone (see back to Section 4.1). To achieve this one must already have groups defined. Indeed, even to reliably run on central galaxies one must already have a reasonable estimate of the group structure of the survey. Therefore, for this comparison, we first run FOF on the observational-like testing data. Using the outputs of the FOF group-finding algorithm, we next use AM from the total stellar mass of each group compared to the halo mass, as constrained in TNG and EAGLE (as with DfL). We refer to this combined approach as: FOF-AM.

Full details on our application of FOF-AM in TNG and EAGLE are provided in Appendix B, including information on optimizing the linking-lengths used on sky and in redshift space (see Fig. B.1), and a demonstration of the abundance matching procedure with the simulations (see Fig. B.2). We also show in Appendix B reproductions of the main performance plots for FOF-AM (see Figs. B.3–B.6). For the sake of brevity, in this subsection we concentrate on the bottom line – a direct comparison of performance statistics.

In Table 2 we show a side-by-side comparison between DfL and FOF-AM for the various training samples and model pairings at  $z = 0, 1, 2$ . Within each cell of Table 2 we show four performance statistics:

- i  $b$ : the median value of  $\delta M_{200}$ ,
  - ii  $\sigma_1$ : the range in  $\delta M_{200}$  containing 68% of the full sample,
  - iii  $\sigma_2$ : the range in  $\delta M_{200}$  containing 95% of the full sample,
  - iv  $f_{\text{out}}$ : the fraction of data with  $|\delta M_{200}| > 1$  dex.
- In both DfL and FOF-AM, halo masses are recovered without any significant bias when comparing runs with the same testing data and model pairings. In the case of mismatched model and

**Table 2.** Performance comparison between DfL and FOF-AM.

| Method/Sim.   | $z = 0$                   | $z = 1$                     | $z = 2$                   |
|---------------|---------------------------|-----------------------------|---------------------------|
| DfL:          |                           |                             |                           |
| TNG (TNG)     | (0.01, 0.11, 0.27, 0.01)  | (0.00, 0.12, 0.28, 0.00)    | (−0.01, 0.13, 0.27, 0.00) |
| EAGLE (EAGLE) | (0.00, 0.12, 0.30, 0.01)  | (−0.01, 0.12, 0.30, 0.01)   | (−0.02, 0.14, 0.28, 0.00) |
| TNG (EAGLE)   | (−0.10, 0.12, 0.27, 0.00) | (−0.08, 0.13, 0.29, 0.00)   | (−0.13, 0.13, 0.28, 0.00) |
| EAGLE (TNG)   | (0.11, 0.13, 0.44, 0.02)  | (0.07, 0.14, 0.32, 0.01)    | (0.10, 0.15, 0.30, 0.00)  |
| FOF-AM:       |                           |                             |                           |
| TNG (TNG)     | (0.01, 0.13, 1.13, 0.06)  | (−0.01, 0.15, 0.83, 0.04)   | (−0.01, 0.15, 0.55, 0.02) |
| EAGLE (EAGLE) | (0.00, 0.15, 1.06, 0.06)  | (−0.02, 0.15, 0.71, 0.03)   | (−0.01, 0.16, 0.39, 0.01) |
| TNG (EAGLE)   | (−0.10, 0.15, 1.15, 0.06) | (−0.08, 0.17, 0.86, 0.04, ) | (−0.13, 0.16, 0.58, 0.01) |
| EAGLE (TNG)   | (0.12, 0.15, 1.14, 0.07)  | (0.05, 0.15, 0.70, 0.03)    | (0.10, 0.15, 0.39, 0.01)  |

**Notes.** For each cell the performance metrics are given as  $(b, \sigma_1, \sigma_2, f_{\text{out}})$ . The biases and dispersions are all given in units of [dex], while the catastrophic outlier fractions are unitless.

data pairings, both DfL and FOF-AM experience very similar (almost identical) biases. This indicates that the biases learned from model choices will equally impact these different methodologies to infer halo masses. This is not surprising given that both approaches require a training set for the halo masses, and hence will learn from that a structure which will not in general be identical to other simulations and models.

In terms of  $\sigma_1$ , the performance of DfL and FOF-AM is quite similar, but DfL is marginally superior across all of the different testing modes and redshift snapshots. However, when viewing  $\sigma_2$  and  $f_{\text{out}}$ , we see that DfL performs with much greater accuracy than the more conventional FOF-AM approach. This indicates that DfL is much less likely to erroneously misplace galaxies into groups than the more standard FOF-AM approach, and hence is less likely to lead to catastrophic outliers and long tails in the  $\delta M_{200}$  distributions.

By comparing Fig. 7 to Fig. B.4, one can immediately see that both FOF-AM and DfL recover narrow Gaussian-like distributions centered on  $\delta M_{200}$ , but that there is more of a power-law tail in the former than the latter. This is further evidenced by the contour relationships (compare Fig. 6 to Fig. B.3), which show significantly reduced correlation strengths in FOF-AM compared to DfL, except when excluding catastrophic outliers. Additionally, the recovery of the halo mass functions is visibly less accurate in FOF-AM compared to DfL (compare Fig. 9 to Fig. B.5).

Taken in concert, the results outlined in Table 2 indicate that DfL performs at least as well as FOF-AM in all metrics, and clearly outperforms FOF-AM in the wings of the  $\delta M_{200}$  distributions. Ultimately, this is achieved by DfL accounting for the virial structure of the haloes by incorporating physically motivated group sizes (in both physical and velocity space) into the group finding algorithm (see Section 3). Additionally, the buffer zone incorporated into the DfL group structure helps to insulate groups from the uncertain region on their edges, preventing the over- or under- assignment of haloes to galaxies. This is a key aspect of why there are significantly fewer catastrophic outliers in DfL compared to FOF-AM.

In principle one could incorporate a similar outlier region in FOF-AM. However, it is important to appreciate that this is not normally done and, moreover, that this is not a natural way to implement a linking-length algorithm. In more detail, FOF operates by assigning group membership based on the distance to any group member in comparison to a universal threshold, whereas DfL operates by assigning group membership based on the dis-

tance to the central galaxy compared to its estimated virial properties. In order to define the outlier region one needs to establish the distance from the central to each galaxy in addition to the virial parameters, and hence one ends up with an approach which is essentially a hybrid of DfL and FOF-AM. This is not suitable for a clear comparison of methodologies.

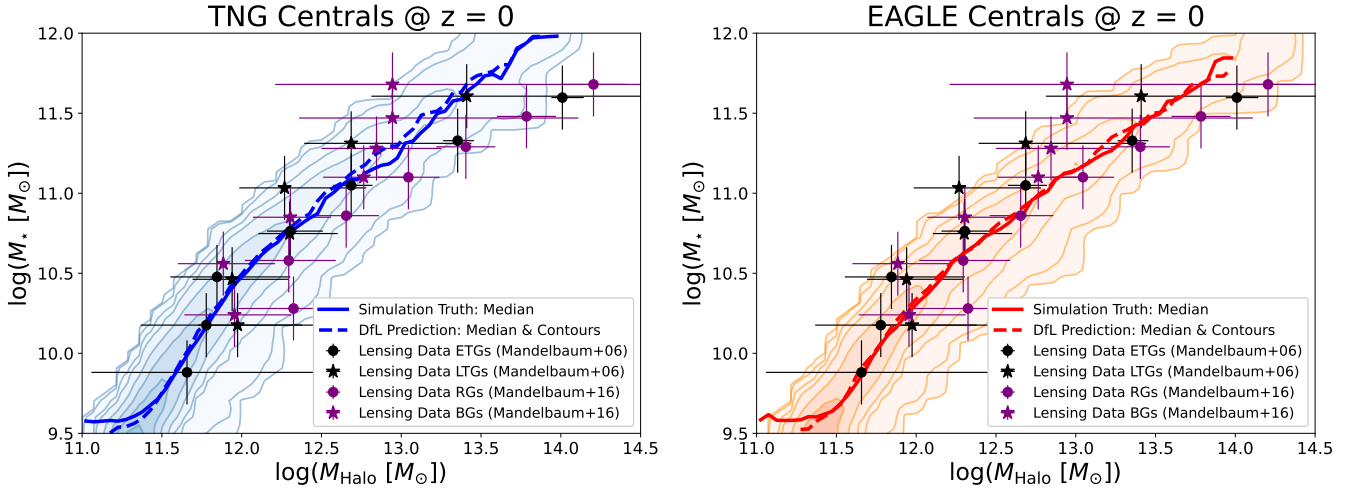
One final thing to note is that DfL provides a single use algorithm to infer the group structure, central – satellite classification, and estimation of halo masses (plus virial parameters) in observational data. In comparison, FOF-AM requires a three stage approach – (i) optimizing the group finder to the observational data set, (ii) running the group finder, and (iii) running abundance matching on the group outputs in comparison to a given simulation or model. Hence, in addition to outperforming FOF-AM in a number of performance metrics, DfL is also more efficient and faster to run on observational data.

#### 4.4. Comparison to observations

So far in the Results section we have assessed the performance of DfL on various simulated data sets, and in comparison to a more established technique (FOF-AM). In this subsection we compare the outputs of DfL applied to TNG and EAGLE to observational constraints at  $z \sim 0$  from gravitational lensing.

In Fig. 11 we present the stellar mass – halo mass ( $M_* - M_{\text{Halo}}$ ) relationship for TNG (left panel) and EAGLE (right panel), both at  $z = 0$ . In both panels the estimates derived from DfL are shown as contours, with the median relationship shown as a dashed solid line. We compare this with the simulation truth (shown as a solid line of the same color). Clearly, DfL does an excellent job of recovering the true  $M_* - M_{\text{Halo}}$  relationship from unseen, observational-like input data in both TNG and EAGLE. This is very reassuring as it demonstrates that one can reconstruct this key scaling relationship in observer ( $2D + z$ ) space to high fidelity using the DfL methodology.

Furthermore, on both panels in Fig. 11 we show observational constraints on the  $M_* - M_{\text{Halo}}$  relationship from strong gravitational lensing. Data points are taken from Mandelbaum et al. (2006, 2016). Observational measurements are shown separately for early-type galaxies (ETGs), late-type galaxies (LTGs), red galaxies (RGs), and blue galaxies (BGs). This means that these data span a high diversity of galaxy types and a large range in both stellar and halo masses. Uncertainties on the halo masses are taken from Mandelbaum et al. (2006, 2016) directly. However, no uncertainties on stellar masses are



**Fig. 11.** Test of DfL with observational data at  $z \sim 0$ . The stellar mass – halo mass relation for central galaxies in TNG (left panel) and EAGLE (right panel), compared to observations from strong gravitational lensing. In both panels the solid colored line indicates the true stellar mass – halo mass relation from each simulation; the dashed colored line indicates this relationship as recovered from observational-like data with the DfL pipeline. Clearly, there is excellent agreement, demonstrating that the DfL pipeline can reproduce the true stellar mass – halo mass relationship from each simulation to high fidelity. The contours in each panel show the full range in the stellar mass – halo mass relationship as ascertained through the DfL pipeline. Additionally, observational constraints from gravitational lensing are shown in each panel, taken from Mandelbaum et al. (2006, 2016). In both cases, DfL does a good job at reproducing the observations (both the general trend and the observed scatter). However, the TNG simulation appears visually to be slightly more accurate compared to EAGLE at  $z = 0$ .

provided in those publications. So, we add a blanket stellar mass uncertainty of 0.2 dex to all data points as the  $Y$ -axis error bars, which is well established as a reasonable baseline (see, e.g., Brinchmann et al. 2004; Peng et al. 2010; Mendel et al. 2014).

The simulated  $M_* - M_{\text{Halo}}$  relationship (and most importantly, the recovered DfL version) are in very good accord with the observational constraints. In the case of TNG, the DfL output contours are almost perfectly spanned by the observational constraints. In the case of EAGLE, there is a slight tendency for the observational constraints to be higher in stellar mass at a fixed halo mass (i.e., more data points lying above the DfL median relationship than below). Nonetheless, the observational data are still reasonably well supported by the DfL recovered simulated extent of this relationship.

Taken in concert, Fig. 11 demonstrates the potential of DfL to recover the true underlying  $M_* - M_{\text{Halo}}$  relationship from simulations under observational limitations. Moreover, Fig. 11 further establishes that this approach leads to a very reasonable estimate of the  $M_* - M_{\text{Halo}}$  relationship as constrained by strong gravitational lensing observations at  $z \sim 0$ . This is very encouraging for the application of DfL at higher redshifts, where the observational constraints from gravitational lensing are much more sparse, and biased to the highest masses.

#### 4.5. The impact of stellar mass uncertainty and incompleteness on halo mass recovery

Up to this point, we have analyzed the performance of DfL with perfect input parameters, albeit limited to observational-like formats. In this subsection we relax that assumption and systematically assess the impact of observational limitations on the performance of DfL under various conditions.

In observational data, modern astrometry combined with precision constraints on cosmological parameters is sufficient to yield physical distance accuracy to at least  $\sim$ kpc scales at the highest redshifts probed in this work (Groggin et al. 2011; Koekemoer et al. 2011; Cirasuolo et al.

2014; Planck Collaboration VI 2020). This is at least two orders of magnitude smaller than the typical virial radii of groups (0.1–1 Mpc). Additionally, through application of spectroscopic redshifts, line-of-sight velocities are obtainable to  $\sim$ 25–75 km/s accuracy at cosmic noon (e.g., Steidel et al. 2014; Cirasuolo et al. 2014; Kriek et al. 2015; Maiolino et al. 2020). This is significantly smaller than typical group virial velocities at these redshifts ( $\sim$ 100–1000 km/s). Hence, we do not consider uncertainties in either coordinates or redshifts here as significant limitations to DfL performance in observational data.

On the other hand, stellar mass estimates in observational data are notoriously imprecise, even when using spectroscopic redshifts or highly sampled multi-band photometry (see, e.g., Mendel et al. 2014; Sánchez et al. 2016; Dimauro et al. 2018; Mucesh et al. 2021). This is in large part due to uncertainty in stellar population synthesis models arising from both stellar population library uncertainty and incompleteness, and from inherent degeneracies between stellar population age and metallicity (among several other parameters). Accounting for all of this uncertainty, most modern stellar mass estimates applied with known (spectroscopic) redshifts are  $\sigma(M_*) \sim 0.2$ – $0.3$  dex (e.g., Mendel et al. 2014; Dimauro et al. 2018; Sánchez et al. 2016; Sánchez 2020). Hence, the impact of stellar mass uncertainty in halo mass recovery with DfL is an essential issue to investigate.

Additionally, incompleteness in observational data due to both spectroscopic fiber collisions and photometric magnitude limits may bias the results obtained through the halo finder pipeline (e.g., Cirasuolo et al. 2020). Here we consider the case where the survey is stellar mass complete at  $M_* > 10^{9.5} M_\odot$  (the approximate mass completeness limit of VLT-MOONRISE, see Maiolino et al. 2020), but incurs a variable (mass-independent) incompleteness, due to fiber collisions and survey planning limitations.

To simulate stellar mass uncertainty we construct a new stellar masses from the original data set by taking a random draw from a Gaussian distribution centered on the true stellar mass, with a standard deviation set equal to  $\sigma(M_*) = 0$ – $0.5$  dex (in

steps of 0.1 dex). That is, before sending to the pipeline in each run, we set

$$M_* \rightarrow \mathcal{R}(\mathcal{N}[\mu = M_*; \sigma = \sigma(M_*)]) \quad (31)$$

for each individual galaxy in the observational-like simulated galaxy survey. This is a rather simplistic modeling of stellar mass uncertainty. However, this approach is appropriate for modeling stellar mass uncertainties in spectroscopic data where the full PDFs are typically uni-modal, due to the accurate constraints on redshifts. Moreover, the use of more complex PDF shapes is not strongly motivated here. That said, DfL enables users to input any known PDF format for stellar masses into the pipeline. This is then fully propagated through the pipeline to yield the impact on the final halo mass PDF (see Section 3). Here our intent is simply to assess the impact of the width of the stellar mass PDF on the resultant errors incurred on halo masses through the DfL pipeline. This is most straightforwardly achieved with a normal distribution.

To simulate incompleteness we randomly remove 0–50% of the data from the input sample (in steps of 10%) prior to running the halo finder pipeline. Of course, in any real galaxy survey the loss in galaxies will not be strictly random. However, until a specific observing strategy is available this is the most logical manner in which to assess the impact of incompleteness. Obviously, this results in cases where there is no halo estimate at all. However, our interest lies in the case where there is an estimate (i.e., a predicted central galaxy exists in the survey data). Hence, we aim to quantify how the incompleteness of other galaxies impacts the halo mass estimate of surviving galaxies.

In Fig. 12 we systematically explore the impact of stellar mass uncertainty and survey incompleteness on the recovery of halo mass and central – satellite classification with DfL. The left column shows results for the TNG simulation, with the right column showing results for the EAGLE simulation. The top row shows the impact on the bias ( $b$ ), the second row shows the impact on the dispersion ( $\sigma_1$ ), and the bottom row shows the impact on central – satellite classification accuracy (Acc). In each panel, the data are used from all redshift ranges from  $z = 0-2$ , and the statistics are measured as offsets relative to the performance with  $\sigma(M_*) = 0$  dex and incompleteness (INC) = 0%.

In terms of bias (top row of Fig. 12), we note a weak systematic bias to overestimate halo masses with increasing stellar mass uncertainty, and a weak systematic bias to underestimate halo masses with increasing incompleteness. Both of these trends are expected. In the case of missing galaxies, in general the stellar mass of groups will be lower leading to underestimates in halo masses. Furthermore, due to the steepness of the stellar mass function, increasing stellar mass uncertainty leads to a greater number of low-mass galaxies moving higher in mass estimates (beyond the mass limit) than the other way around. This in turn leads to a small systematic bias to overestimate the halo masses of groups. That said, the biases incurred by increasing stellar mass uncertainty and survey incompleteness are generally quite small. For example, in VLT-MOONRISE, with an estimated  $\sigma(M_*) = 0.3$  dex and INC = 30% (see Maiolino et al. 2020), we would expect to engender no significant additional bias in the recovery of halo masses. This region is highlighted with a solid box in Fig. 12 on all panels for reference. We note that the bias incurred due to model choice will still be an important factor (see Section 4.2).

In terms of dispersion (center row of Fig. 12), we note a systematic increase in the uncertainty on halo mass recovery with

both increasing stellar mass uncertainty and survey incompleteness. Again, this is as expected. Increasing the uncertainty on the stellar mass inputs will lead to more uncertainty on the halo mass outputs. Additionally, increasing incompleteness will add more noise into the group membership, leading to more uncertainty in halo mass recovery. However, the dependence is not symmetric. The impact of stellar mass uncertainty is greater than the impact of survey incompleteness. For the VLT-MOONRISE specifications (mentioned above), we anticipate increasing the uncertainty on the halo mass estimates over their optimal values (see Fig. 7) by  $\sim 0.15$  dex. This leads to a total uncertainty of  $\sim 0.25-0.3$  dex expected in this upcoming survey. We note that this is a significant improvement on halo mass recovery with stellar mass uncertainty at  $z \sim 0$  with abundance matching techniques ( $\sim 0.3-0.5$  dex, see, e.g., Yang et al. 2007; Woo et al. 2013).

Finally, in the bottom panels we explore the impact of stellar mass uncertainty and incompleteness on the classification of central and satellite galaxies in DfL. We find that the classification accuracy systematically decreases with both increasing stellar mass uncertainty and increasing incompleteness in the survey. Here, however, the impact of incompleteness is similarly important to that of stellar mass uncertainty (unlike in the case of dispersion, above). Again this is as expected. Higher uncertainty on stellar masses will lead to a greater chance that a satellite overtakes the true central in measured stellar mass, leading to a misclassification. Additionally, absent data may on occasion remove a central from a group, which will lead to a satellite taking its place (as the most massive galaxy), again leading to misclassification. For VLT-MOONRISE specifications, we anticipate a reduction in central – satellite classification accuracy of 5–8%. This would lead to a final central – (core) satellite classification accuracy of 88–91%.


Taken in concert, Fig. 12 can be used to infer how survey limitations (explicitly due to stellar mass uncertainty and incompleteness) will impact the performance of DfL in recovering halo masses and in central – satellite classification. We note that these effects are not strongly redshift dependent, and hence we present in Fig. 12 the results for the entire redshift range for the sake of brevity. However, redshift specific performance matrices are available from the corresponding author upon request.

## 5. Discussion: applications of DfL in wide-field spectroscopic galaxy surveys

One of the immediate intended applications of the DfL halo finder pipeline presented in this work is to the VLT-MOONRISE wide-field spectroscopic galaxy survey targeting cosmic noon, explicitly  $z = 0.8-2.0$  (see Cirasuolo et al. 2020; Maiolino et al. 2020). MOONRISE will obtain rest-frame optical spectra for  $\sim 500$  k galaxies across three wide-field regions within (i) COSMOS; (ii) EGS; and (iii) VIDEO, each with  $\geq 1$  deg<sup>2</sup> area on sky. This means that MOONRISE will provide near optimal input data for DfL, accessing a unique cosmological parameter space hitherto unexplored with dense spectroscopic sampling across a large volume. Initial data from this unprecedented galaxy survey will become available to the community by the end of 2026, with the final survey data publicly available by the end of 2031<sup>11</sup>.

It is anticipated that MOONRISE will achieve approximate stellar mass completeness above  $M_* > 10^{9.5} M_\odot$ , and spectroscopic sampling completeness of  $\sim 70\%$  for a large fraction of

<sup>11</sup> Dates may be subject to change due to technical and survey planning issues.



**Fig. 12.** Halo recovery performance from DfL as a function of spectroscopic incompleteness and stellar mass uncertainty. The left column shows results for the TNG simulation and the right column shows results for the EAGLE simulation. The top row shows the systematic shift in bias of halo mass recovery ( $\Delta \text{Bias}$ ), the middle row shows the systematic shift in statistical uncertainty of halo mass recovery ( $\Delta \sigma_1(M_{200})$ ), and the bottom row shows the systematic shift in central – satellite classification accuracy ( $\Delta \text{Acc}(C/S)$ ). These statistics are all calculated as a function of increasing stellar mass uncertainty and survey incompleteness, relative to the case of a complete sample with zero stellar mass uncertainty. Hence, this figure enables one to ascertain how these two critical observational limitations will impact the output halo properties from DfL under realistic observational conditions. The expected survey incompleteness and stellar mass uncertainty of MOONRISE is indicated by a black box on each grid. We expect DfL run on MOONRISE not to incur increased bias, to have increased statistical uncertainty of  $\sim 0.15$  dex, and to have decreased accuracy in central satellite classification of 5–8% (dependent on simulation). These values are measured relative to the case of running on a complete observational spectroscopic sample, with perfect knowledge of stellar masses.

the survey (Maiolino et al. 2020). Additionally, we estimate that stellar masses will be recovered using extant multi-band photometry plus new spectroscopic redshifts with an accuracy of  $\sigma(M_*) \leq 0.3$  dex. This will likely be improved upon by the subset of data for which full continuum fitting is feasible (anticipated to be  $\sim 50\%$  of galaxies), and for galaxies with exten-

sive JWST multi-band follow-up. The latter will improve stellar mass estimates by constraining the redder part of the spectrum, enabling more stringent constraints on dust content.

The performance expectations for DfL under the planned limitations of MOONRISE are indicated by bold rectangular regions in Fig. 12 (as discussed in the previous subsection).

We anticipate to be able to achieve halo masses from  $M_{\text{Halo}} = 10^{11} - 10^{14} M_{\odot}$  with uncertainties of  $\sigma(M_{\text{halo}}) \leq 0.3$  dex and a model-dependent bias of  $\pm 0.1$  dex. Moreover, we will provide full PDFs for all halo mass estimates, enabling rigorous accounting for uncertainty in various scientific applications. Furthermore, we also anticipate to achieve a central – (core) satellite classification accuracy of  $\sim 90\%$ . In this subsection we list several potential scientific applications of the MOONRISE halo mass and group catalog from DfL. Many, if not all, of these applications will also be feasible with other wide-field galaxy surveys moving forwards.

(i) Halo mass functions & the growth of large scale structure: Perhaps the most direct application of the halo finder pipeline to MOONRISE will be the formation of halo mass functions from  $z = 0.8 - 2.0$ . This will be the first time that accurate halo mass functions can be empirically determined at these relatively early cosmic times across a large dynamic range in halo mass (up to three orders of magnitude). This will directly probe the growth of large scale structure as a function of cosmic time. Moreover, these halo mass functions can be directly compared to measured halo mass functions at lower redshifts, enabling strong empirical constraints across  $\sim 10$  Gyr of cosmic history.

Furthermore, one can also use these constraints on the growth of large scale structure to test and constrain cosmological models. In the first instance, one can compare to dark matter only simulations to gauge the accuracy of the underlying cosmological model and hierarchical assembly physics from gravitation alone. Interestingly, one can even use each hydrodynamical simulation in question to predict the best light-to-dark mapping. Differences in the halo mass functions which persist will point to statistically meaningful differences between the galaxy formation model and the real Universe, which can be used to constrain the feedback mechanisms used in future generations of cosmological simulations. In the most extreme case, a fundamental mismatch between predicted and inferred halo mass distributions may point towards new fundamental physics (see also point (v)).

(ii) The role of cosmic environment in galaxy formation: In the hierarchical theory of structure formation, galaxies form within dark matter haloes and grow as a result of both accretion of dark matter and baryons and through halo (and galaxy) mergers (e.g., Cole et al. 2000; Bower et al. 2006; Vogelsberger et al. 2014a; Schaye et al. 2015; Davé et al. 2019). Using the DfL halo finder presented in this work applied to MOONRISE, in conjunction with computational morphology on extant imaging and close pair statistics in MOONRISE fields, one may investigate how the merger rate of galaxies varies as a function of halo mass and galaxy class (central vs. satellite) across cosmic time. This will reveal where and when in the history of the Universe galaxy formation occurs and is most efficient. Furthermore, through combining this data with star formation rate tracers (e.g., extinction corrected  $H\alpha$  luminosity, see Kennicutt 1998), one can also quantify whether in situ star formation or mergers dominate the growth of stellar mass in galaxies as a function of location within the cosmic web.

Ultimately, probing the relationship between halo mass and stellar mass as a function of cosmic time will provide precision new constraints on the formation of galaxies within their dark matter haloes. Comparison with simulations will be particularly constraining of feedback mechanisms, both stellar and AGN. We note that even though stellar mass is used as an input to the halo finder, the number of galaxies within a group and the total group stellar mass significantly ameliorate the issue of comparing the same parameter to itself. For instance, using extant halo catalogs

based on stellar mass inputs at  $z \sim 0$ , important differences in the dependencies of galaxy properties on stellar and halo mass have been found (e.g., Yang et al. 2009; Woo et al. 2013; Bluck et al. 2014, 2016).

(iii) The role of cosmic environment in galaxy quenching: One of the two primary routes to quench a galaxy is through environment, the other being via AGN feedback (see, e.g., Peng et al. 2010, 2012; Bluck et al. 2020a,b, 2022; Piotrowska et al. 2022; Goubert et al. 2024). Using the group catalogs output by the halo finder, in addition to the halo mass estimates, one can probe how galactic star formation and its cessation through quenching is connected with cosmic environment as a function of cosmic time.

We will be able to reveal where and when in the history of the Universe dense environments are conducive to star formation and when they become hostile to future star formation, resulting in galaxy quenching. Moreover, we will also be able to investigate the relationship of galactic star formation and quenching to both the mass of groups and clusters and to the location of satellites within their groups. This is unprecedented science at cosmic noon, due largely to the prior lack of wide-field spectroscopic galaxy surveys targeting this epoch. This is essential for accurate group identification on large scales and for the estimation of the dark sector through DfL.

(iv) The role of cosmic environment in triggering active galactic nuclei (AGN): Using the DfL halo mass catalogs for MOONRISE in addition to optical AGN identification through emission lines (e.g., Baldwin et al. 1981; Kewley et al. 2001, 2006; Kauffmann et al. 2003), one will be able to search for possible links between AGN and cosmic environment. Questions we can answer with DfL applied to MOONRISE include: are the gas rich cores of high-mass groups and proto-clusters the origin of the most powerful quasars in the Universe?; how and when do these environments become stabilized through feedback processes?; is cosmic environment the fundamental driver of supermassive black hole growth?; and, furthermore, which grows first, the dark matter halo, the galaxy, or the supermassive black hole?

Supermassive black hole masses may be inferred in MOONRISE for broad-line AGN utilizing emission line calibrations (e.g., Maiolino et al. 2024a,b). Additionally, black hole masses may be statistically constrained on a population basis through calibrations with photometric (e.g., the stellar potential,  $\phi_* \sim M_*/R_e$ ) and kinematic (e.g., the central velocity dispersion,  $\sigma_*$ ) measurements (see Bluck et al. 2022, 2023, 2024). Hence, another powerful science application will be to explore how the halo mass – black hole mass relation varies with cosmic time, and location of galaxies within the cosmic web. This will reveal key new insights on whether supermassive black holes grow early or late in the history of the Universe, ultimately answering whether they are the end product of galaxy formation, or, perhaps, a trigger for it.

(v) Direct tests of  $\Lambda$ CDM & alternative cosmologies: Perhaps the most exciting potential application of the halo finder pipeline applied to VLT-MOONRISE data is the potential to constrain the underlying background cosmology. The evolution in the halo mass function at high-mass scales is largely unaffected by baryonic physics (e.g., Vogelsberger et al. 2014a,b; Schaye et al. 2015; Henriques et al. 2015, 2019). Hence, variations in the observed halo mass functions at large scales, when compared to  $\Lambda$ CDM models, may indicate the need for new physics, if significant tensions emerge. Given the precise knowledge we have gained in this work on both the systematic and random uncertainties of the recovered halo masses, we can rigorously quantify any high-mass deviations from the expected halo

mass functions, potentially revealing the need for modifications to  $\Lambda$ CDM and/or new physics.

Yet another powerful cosmological constraint can be obtained through the halo merger rate, which is strongly dependent upon cosmological parameters (e.g., [Conselice et al. 2014](#)). Through kinematics of central – central close pairs, using the classifications from the halo finder, we will be able to make the first direct measurements of halo assembly at cosmic noon. The dependence of halo merger rates is especially sensitive to the dark energy equation of state. Therefore, by comparing the cosmological dark matter halo merger rate at high redshifts to the local Universe, we can obtain independent constraints on the transition from matter to dark energy domination in the Universe, and potentially find evidence for, or against, quintessence (e.g., [Steinhardt et al. 1999](#); [Copeland et al. 2006](#)).

(vi) Environmental physics beyond the virial radius: During the course of this paper the extended satellite population has been more of a challenge than an advantage. However, identifying the region on the edge of galaxy groups and clusters has many potential science benefits. For instance, looking for correlations between star formation enhancements and/or suppression as a function of physical distance from the group center will yield new constraints on the splashback radius (e.g., [Adhikari et al. 2014](#); [Baxter et al. 2017](#)), and the fraction of non-virialized high-mass objects at cosmic noon (e.g., [Chiang et al. 2013](#); [Overzier 2016](#)). Ultimately, studies of the extended satellite population (despite its relative impurity) will yield new insights on how environment shapes galaxy evolution beyond the virial radius.

(vii) Cross validation of halo properties with alternative methods: In addition to the many scientific applications outlined above there are also important technical applications. Researchers may also infer halo masses from clustering statistics (e.g., [Berlind & Weinberg 2002](#); [Zheng et al. 2007](#)), sub-halo and/or conditional abundance matching (e.g., [Hearin et al. 2013, 2014](#)), in addition to the machine learning approach presented in this work (in comparison to the more standard FOF-AM technique, see [Vale & Ostriker 2004](#); [Conroy et al. 2006](#)). Cross validation through comparison of different approaches will be highly beneficial in exposing advantages and disadvantages between methods and establishing whether or not key insights are method dependent.

There are, of course, many more scientific applications which will become tractable once a halo catalog with well-defined errors is available at high redshifts through DfL. Nonetheless, we hope the above provides a strong motivation for the approach outlined in this paper, and provides some useful ideas for applications within the extragalactic community in the coming years.

In future work we plan to add more simulations to the DfL library, especially SIMBA ([Davé et al. 2019](#)), Horizon-AGN ([Dubois et al. 2014](#)), MassiveBlack-II ([Khandai et al. 2015](#)), and Romulus ([Tremmel et al. 2017](#)), among several others. Moreover, we also plan to add greater flexibility to DfL by adding the option to include more galaxy parameters in halo estimation, and by enabling an ANN (in addition to RF) module for mapping between light and dark sectors. Furthermore, we also plan to make further benchmarking comparisons to a host of other methodologies from CAM to HOD, in addition to FOF-AM (considered in [Appendix B](#)). Finally, we anticipate adding a magnitude limit option to the DfL halo finder, enabling straightforward applications to galaxy surveys beyond MOONRISE. Hence, DfL is intended as living code repository. What is presented in this work is essentially version-1 (see [Appendix A](#) for

access), with proof of concept, initial testing, and a comprehensive discussion of the methodology. Much more is planned for the coming years.

## 6. Summary

In this paper we presented DfL, a novel method for inferring the dark sector from luminous tracers at all epochs from the present to cosmic noon. DfL leverages a machine learning-based approach utilizing random forest regression, trained on cosmological simulations (specifically EAGLE and IllustrisTNG). The essential idea behind DfL is to learn an effective mapping from observational-like input data to dark matter halo properties, including masses, group memberships, and other derived properties (e.g., virial radii, velocities, and temperatures).

In [Section 3](#) we presented a detailed overview of the DfL halo finder pipeline. Briefly, two regression models are prepared for each simulation, spanning  $z = 0-3$  in snapshots. The first traces a direct mapping from stellar to halo mass for central galaxies (as in abundance matching), but with redshift information added. The second incorporates the total stellar mass of the group and the number of galaxies in the group, in addition to the stellar mass and redshift of the central, in order to yield a more accurate halo mass estimate (see [Fig. 2](#)).

The first regression model is used to estimate the physical and phase-space size of the group from just the central galaxy stellar mass and redshift. Central galaxies are identified initially as the most massive galaxy in the survey and then sequentially as the most massive galaxy remaining in the survey after each group is removed in turn. Satellites are identified as galaxies within one virial radius and one virial velocity of the central galaxy, with an extended satellite population counted out to two virial radii and two virial velocities. This information is then used to infer the inputs for the second stage from the group membership.

This approach results in two estimates of halo mass: from the central and from the group as a whole. If they are similar enough ( $\delta M_{\text{Halo}} < \text{lim.}$ ), the halo mass and group membership are set, and the group is removed from the pipeline. If they are not sufficiently similar the pipeline iterates, allowing the group to grow or shrink in phase-space based on the new group membership. This process continues until convergence is reached, or else a hard cutoff is exceeded.

Crucially, DfL requires only coordinates, spectroscopic redshifts, and a stellar mass estimate for every galaxy in the survey in order to run. This makes it extremely convenient for use on wide-field spectroscopic galaxy surveys. Additionally, stellar mass uncertainties may be specified as either the width of a Gaussian PDF or else as a full nonparametric PDF. DfL propagates the uncertainty on stellar masses through the pipeline, outputting full PDFs for the halo mass estimates, which also take account of the distribution of individual tree predictions for each group from the RF regression analysis.

In [Section 4](#) we describe how we tested the performance of the halo finder on observational-like inputs drawn from the two simulations, which are unseen by the RF models in either training or validation. This allowed us to compare the halo mass estimates to the actual halo masses in each simulation, testing the performance of DfL under observational conditions. When applying the same model and data pairing, we recovered a bias-free estimate of halo mass with a random uncertainty of  $\langle \sigma \rangle = 0.12$  dex. Furthermore, we also obtained a 96% accuracy in central – (core) satellite classification, which is lowered to 88% including the extended satellite population.

We performed cross-validation tests of the pipeline, where we mismatched the model and data pairing to test the impact of the galaxy formation model choice on the recovery of halo properties (see Section 4.2). We find that the random error is largely unaffected by this, but that the bias rises to  $\langle b \rangle = \pm 0.1$  dex. This is taken as a preliminary estimate of the systematic uncertainty induced by model choice, resulting from the fundamental uncertainty of the true underlying galaxy formation model of the Universe. We plan to investigate this further in future work by comparing the training and testing of DfL on other cosmological simulations.

We compared the performance of DfL to the more standard FOF-AM technique (abundance matching applied to the total group stellar mass, inferred through a friends-of-friends group finder; see Appendix B). DfL and FOF-AM perform similarly in terms of bias and the central dispersion of the offset in halo masses ( $\delta M_{200}$ ). However, DfL yields significantly fewer catastrophic outliers, and hence leads to a much improved recovery of halo masses in the outskirts of the distribution (see Section 4.3). Additionally, we compared the output  $M_* - M_{\text{Halo}}$  relationships from DfL to observational constraints from strong gravitational lensing in Section 4.4. The recovered relationship is in very good agreement with these observational constraints at low redshifts.

We systematically explored how stellar mass uncertainty and spectroscopic incompleteness impact the recovery of halo masses in DfL (see Section 4.5). For VLT-MOONRISE, we anticipate a reduction in accuracy of  $\sim 0.15$  dex due to the combination of these effects, but little impact on the overall bias in halo mass recovery. Additionally, we anticipate a reduction in central – satellite accuracy of  $\sim 5\text{--}8\%$ .

Finally, in Section 5 we presented some potential scientific applications of the DfL halo finder applied to VLT-MOONRISE, ranging from cosmology to galaxy formation and evolution. In the Appendix we provide information on how to access and use the halo finder pipeline (Appendix A), details on the FOF-AM method and results that are used to compare to DfL (Appendix B), and a list of hyper-parameters used in the RF and ANN approaches for reproducibility (Appendix C).

In conclusion, DfL leverages machine learning trained on cosmological hydrodynamical simulations to provide a fast, straightforward, and accurate route to infer the dark sector from luminous tracers at all epochs up to cosmic noon. This software is publicly available for use in research (see below for access).

## Data availability

All of the data used in this paper is publicly available, or else currently in preparation for public release. For the DfL Halo Finder<sup>12</sup> see Appendix A for full details. For simulations, see: Eagle data access<sup>13</sup>, IllustrisTNG data access<sup>14</sup>, Docker<sup>15</sup>. All analyses in this work were performed using PYTHON-3<sup>16</sup>, including NUMPY, ASTROPY, SCIPY, PANDAS, SEABORN, MATPLOTLIB. All machine learning analyses were performed using SCIKIT-LEARN<sup>17</sup>.

*Acknowledgements.* AFLB gratefully acknowledges support from an NSF research grant: NSF-AST 2408009, in addition to an ORAU Junior Faculty

Enhancement Award in Physical Sciences, and research start-up funds from FIU. PG also acknowledges support from NSF-AST 2408009. RM acknowledges support from the Science and Technology Facilities Council (STFC), ERC Advanced Grant 695671 “QUENCH”, and by the UKRI Frontier Research grant RISEandFALL. RM also acknowledges funding from a research professorship from the Royal Society.

## References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Phys. Rev. Lett.*, **116**, 061102
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *Phys. Rev. Lett.*, **119**, 161101
- Adhikari, S., Dalal, N., & Chamberlain, R. T. 2014, *JCAP*, **2014**, 019
- Ansoldi, S., Antonelli, L. A., Arcaro, C., et al. 2018, *ApJ*, **863**, L10
- Artale, M. C., Zehavi, I., Contreras, S., & Norberg, P. 2018, *MNRAS*, **480**, 3978
- Atek, H., Richard, J., Kneib, J.-P., & Schaerer, D. 2018, *MNRAS*, **479**, 5184
- Auger, M. W., Treu, T., Gavazzi, R., et al. 2010, *ApJ*, **721**, L163
- Baldwin, J. A., Phillips, M. M., & Terlevich, R. 1981, *PASP*, **93**, 5
- Behroozi, M., Lahav, O., Lintott, C. J., et al. 2010, *MNRAS*, **406**, 342
- Baron, D. 2019, ArXiv e-prints [arXiv:1904.07248]
- Baxter, E., Chang, C., Jain, B., et al. 2017, *ApJ*, **841**, 18
- Begeman, K. G. 1989, *A&A*, **223**, 47
- Behroozi, P. S., Conroy, C., & Wechsler, R. H. 2010, *ApJ*, **717**, 379
- Benson, A. J., Cole, S., Frenk, C. S., Baugh, C. M., & Lacey, C. G. 2000, *MNRAS*, **311**, 793
- Berlind, A. A., & Weinberg, D. H. 2002, *ApJ*, **575**, 587
- Bertotti, B., Iess, L., & Tortora, P. 2003, *Nature*, **425**, 374
- Bluck, A. F. L., Mendel, J. T., Ellison, S. L., et al. 2014, *MNRAS*, **441**, 599
- Bluck, A. F. L., Mendel, J. T., Ellison, S. L., et al. 2016, *MNRAS*, **462**, 2559
- Bluck, A. F. L., Maiolino, R., Piotrowska, J. M., et al. 2020a, *MNRAS*, **499**, 230
- Bluck, A. F. L., Maiolino, R., Sánchez, S. F., et al. 2020b, *MNRAS*, **492**, 96
- Bluck, A. F. L., Maiolino, R., Brownson, S., et al. 2022, *A&A*, **659**, A160
- Bluck, A. F. L., Piotrowska, J. M., & Maiolino, R. 2023, *ApJ*, **944**, 108
- Bluck, A. F. L., Conselice, C. J., Ormerod, K., et al. 2024, *ApJ*, **961**, 163
- Bocquet, S., Dietrich, J. P., Schrabback, T., et al. 2019, *ApJ*, **878**, 55
- Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, *ApJ*, **638**, 703
- Bondi, H., & Hoyle, F. 1944, *MNRAS*, **104**, 273
- Bower, R. G., Benson, A. J., Malbon, R., et al. 2006, *MNRAS*, **370**, 645
- Bower, R. G., McCarthy, I. G., & Benson, A. J. 2008, *MNRAS*, **390**, 1399
- Breiman, L. 2001, *Mach. Learn.*, **45**, 5
- Brinchmann, J., Charlot, S., White, S. D. M., et al. 2004, *MNRAS*, **351**, 1151
- Brink, H., Richards, J. W., Poznanski, D., et al. 2013, *MNRAS*, **435**, 1047
- Campbell, D., van den Bosch, F. C., Padmanabhan, N., et al. 2018, *MNRAS*, **477**, 359
- Cappellari, M., Bacon, R., Bureau, M., et al. 2006, *MNRAS*, **366**, 1126
- Cappellari, M., McDermid, R. M., Alatalo, K., et al. 2013, *MNRAS*, **432**, 1862
- Carlberg, R. G., Yee, H. K. C., Ellingson, E., et al. 1996, *ApJ*, **462**, 32
- Carlstrom, J. E., Holder, G. P., & Reese, E. D. 2002, *ARA&A*, **40**, 643
- Chabrier, G. 2003, *PASP*, **115**, 763
- Chaves-Montero, J., Angulo, R. E., Schaye, J., et al. 2016, *MNRAS*, **460**, 3100
- Chiang, Y.-K., Overzier, R., & Gebhardt, K. 2013, *ApJ*, **779**, 127
- Chua, A. J. K., & Vallisneri, M. 2020, *Phys. Rev. Lett.*, **124**, 041102
- Cirasuolo, M., Afonso, J., Carollo, M., et al. 2014, in Ground-based and Airborne Instrumentation for Astronomy V, eds. S. K. Ramsay, I. S. McLean, & H. Takami, *SPIE Conf. Ser.*, **9147**, 91470N
- Cirasuolo, M., Fairley, A., Rees, P., et al. 2020, *Messenger*, **180**, 10
- Cochrane, R. K., Best, P. N., Sobral, D., et al. 2018, *MNRAS*, **475**, 3730
- Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000, *MNRAS*, **319**, 168
- Conroy, C., Wechsler, R. H., & Kravtsov, A. V. 2006, *ApJ*, **647**, 201
- Conselice, C. J., Bluck, A. F. L., Mortlock, A., Palamara, D., & Benson, A. J. 2014, *MNRAS*, **444**, 1125
- Copeland, E. J., Sami, M., & Tsujikawa, S. 2006, *Int. J. Mod. Phys. D*, **15**, 1753
- Crain, R. A., Schaye, J., Bower, R. G., et al. 2015, *MNRAS*, **450**, 1937
- Croton, D. J., Gao, L., & White, S. D. M. 2007, *MNRAS*, **374**, 1303
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, **486**, 2827
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
- de Blok, W. J. G., Walter, F., Brinks, E., et al. 2008, *AJ*, **136**, 2648
- Dieleman, S., Willett, K. W., & Dambre, J. 2015, *MNRAS*, **450**, 1441
- Dimaro, P., Huertas-Company, M., Daddi, E., et al. 2018, *MNRAS*, **478**, 5410
- Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, *MNRAS*, **399**, 497
- Duarte, M., & Mamon, G. A. 2014, *MNRAS*, **440**, 1763
- Dubois, Y., Pichon, C., Welker, C., et al. 2014, *MNRAS*, **444**, 1453
- Efstathiou, G., Sutherland, W. J., & Maddox, S. J. 1990, *Nature*, **348**, 705
- Erb, D. K., Steidel, C. C., Shapley, A. E., et al. 2006, *ApJ*, **646**, 107
- Everitt, C. W. F., Debra, D. B., Parkinson, B. W., et al. 2011, *Phys. Rev. Lett.*, **106**, 221101
- Evrard, A. E., Metzler, C. A., & Navarro, J. F. 1996, *ApJ*, **469**, 494

- Frenk, C. S., White, S. D. M., Davis, M., & Efstathiou, G. 1988, *ApJ*, **327**, 507
- Gerke, B. F., Newman, J. A., Davis, M., et al. 2005, *ApJ*, **625**, 6
- Giodini, S., Pierini, D., Finoguenov, A., et al. 2009, *ApJ*, **703**, 982
- Goubert, P. H., Bluck, A. F. L., Piotrowska, J. M., & Maiolino, R. 2024, *MNRAS*, **528**, 4891
- Grogin, N. A., Kocevski, D. D., Faber, S. M., et al. 2011, *ApJS*, **197**, 35
- Guo, Q., White, S., Li, C., & Boylan-Kolchin, M. 2010, *MNRAS*, **404**, 1111
- Hadzhiyska, B., Bose, S., Eisenstein, D., Hernquist, L., & Spergel, D. N. 2020, *MNRAS*, **493**, 5506
- Hadzhiyska, B., Bose, S., Eisenstein, D., & Hernquist, L. 2021, *MNRAS*, **501**, 1603
- Hadzhiyska, B., Eisenstein, D., Bose, S., Garrison, L. H., & Maksimova, N. 2022, *MNRAS*, **509**, 501
- Hadzhiyska, B., Eisenstein, D., Hernquist, L., et al. 2023a, *MNRAS*, **524**, 2507
- Hadzhiyska, B., Hernquist, L., Eisenstein, D., et al. 2023b, *MNRAS*, **524**, 2524
- Hahn, C., Bottrell, C., & Lee, K.-G. 2024, *ApJ*, **968**, 90
- Hearin, A. P., Zentner, A. R., Berlind, A. A., & Newman, J. A. 2013, *MNRAS*, **433**, 659
- Hearin, A. P., Watson, D. F., Becker, M. R., et al. 2014, *MNRAS*, **444**, 729
- Henriques, B. M. B., White, S. D. M., Thomas, P. A., et al. 2015, *MNRAS*, **451**, 2663
- Henriques, B. M. B., White, S. D. M., Lilly, S. J., et al. 2019, *MNRAS*, **485**, 3446
- Hickson, P., Mendes de Oliveira, C., Huchra, J. P., & Palumbo, G. G. 1992, *ApJ*, **399**, 353
- Ho, M., Bartlett, D. J., Chartier, N., et al. 2024, *Open J. Astrophys.*, **7**, 54
- Hoekstra, H., Bartelmann, M., Dahle, H., et al. 2013, *Space Sci. Rev.*, **177**, 75
- Hortúa, H. J., García, L. Á., & Castañeda, L. 2023, *Front. Astron. Space Sci.*, **10**, 1139120
- Hoyle, F., & Lyttleton, R. A. 1939, *Proc. Camb. Philos. Soc.*, **35**, 405
- Hu, W., & Sugiyama, N. 1995, *ApJ*, **444**, 489
- Huertas-Company, M., Iyer, K. G., Angeloudi, E., et al. 2024, *A&A*, **685**, A48
- Ilbert, O., Arnouts, S., McCracken, H. J., et al. 2006, *A&A*, **457**, 841
- Jones, E., Do, T., Boscoe, B., et al. 2024, *ApJ*, **964**, 130
- Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, *MNRAS*, **346**, 1055
- Kennicutt, R. C., Jr. 1998, *ARA&A*, **36**, 189
- Kewley, L. J., Dopita, M. A., Sutherland, R. S., Heisler, C. A., & Trevena, J. 2001, *ApJ*, **556**, 121
- Kewley, L. J., Groves, B., Kauffmann, G., & Heckman, T. 2006, *MNRAS*, **372**, 961
- Khandai, N., Di Matteo, T., Croft, R., et al. 2015, *MNRAS*, **450**, 1349
- Klypin, A. A., Trujillo-Gomez, S., & Primack, J. 2011, *ApJ*, **740**, 102
- Klypin, A., Prada, F., Yepes, G., Heß, S., & Gottlöber, S. 2015, *MNRAS*, **447**, 3693
- Knebe, A., Knollmann, S. R., Muldrew, S. I., et al. 2011, *MNRAS*, **415**, 2293
- Knobel, C., Lilly, S. J., Iovino, A., et al. 2009, *ApJ*, **697**, 1842
- Koekemoer, A. M., Faber, S. M., Ferguson, H. C., et al. 2011, *ApJS*, **197**, 36
- Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, *ApJS*, **192**, 18
- Kraivtsov, A. V., Berlind, A. A., Wechsler, R. H., et al. 2004, *ApJ*, **609**, 35
- Kriek, M., Shapley, A. E., Reddy, N. A., et al. 2015, *ApJS*, **218**, 15
- Lehmann, B. V., Mao, Y.-Y., Becker, M. R., Skillman, S. W., & Wechsler, R. H. 2017, *ApJ*, **834**, 37
- Looser, T. J., Lilly, S. J., Sin, L. P. T., et al. 2021, *MNRAS*, **504**, 3029
- Maiolino, R., Cirasuolo, M., Afonso, J., et al. 2020, *Messenger*, **180**, 24
- Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2024a, *A&A*, **691**, A145
- Maiolino, R., Scholtz, J., Witsok, J., et al. 2024b, *Nature*, **627**, 59
- Mandelbaum, R. 2018, *ARA&A*, **56**, 393
- Mandelbaum, R., Seljak, U., Kauffmann, G., Hirata, C. M., & Brinkmann, J. 2006, *MNRAS*, **368**, 715
- Mandelbaum, R., Wang, W., Zu, Y., et al. 2016, *MNRAS*, **457**, 3200
- Mannheim, K., Protheroe, R. J., & Rachen, J. P. 2000, *Phys. Rev. D*, **63**, 023003
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, **480**, 5113
- McAlpine, S., Helly, J. C., Schaller, M., et al. 2016, *Astron. Comput.*, **15**, 72
- Mendel, J. T., Simard, L., Palmer, M., Ellison, S. L., & Patton, D. R. 2014, *ApJS*, **210**, 3
- Mo, H., van den Bosch, F. C., & White, S. 2010, *Galaxy Formation and Evolution* (Cambridge)
- More, S., Kravtsov, A. V., Dalal, N., & Gottlöber, S. 2011, *ApJS*, **195**, 4
- More, S., Diemer, B., & Kravtsov, A. V. 2015, *ApJ*, **810**, 36
- More, S., Miyatake, H., Takada, M., et al. 2016, *ApJ*, **825**, 39
- Moster, B. P., Somerville, R. S., Maulbetsch, C., et al. 2010, *ApJ*, **710**, 903
- Mucesh, S., Hartley, W. G., Palmese, A., et al. 2021, *MNRAS*, **502**, 2770
- Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **477**, 1206
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, *ApJ*, **462**, 563
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **475**, 624
- Overzier, R. A. 2016, *A&ARv*, **24**, 14
- Peacock, J. A., & Smith, R. E. 2000, *MNRAS*, **318**, 1144
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Peng, Y.-J., Lilly, S. J., Kovač, K., et al. 2010, *ApJ*, **721**, 193
- Peng, Y.-J., Lilly, S. J., Renzini, A., & Carollo, M. 2012, *ApJ*, **757**, 4
- Persic, M., Salucci, P., & Stel, F. 1996, *MNRAS*, **281**, 27
- Pillepich, A., Springel, V., Nelson, D., et al. 2018, *MNRAS*, **473**, 4077
- Piotrowska, J. M., Bluck, A. F. L., Maiolino, R., & Peng, Y. 2022, *MNRAS*, **512**, 1052
- Planck Collaboration VI. 2020, *A&A*, **641**, A6
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Planck Collaboration XVI. 2014, *A&A*, **571**, A16
- Pratt, G. W., Arnaud, M., Piffaretti, R., et al. 2010, *A&A*, **511**, A85
- Press, W. H., & Davis, M. 1982, *ApJ*, **259**, 449
- Refsdal, S. 1964, *MNRAS*, **128**, 307
- Remy, B., Lanusse, F., Jeffrey, N., et al. 2023, *A&A*, **672**, A51
- Rubin, V. C., Ford, W. K., Jr., & Thonnard, N. 1978, *ApJ*, **225**, L107
- Sánchez, S. F. 2020, *ARA&A*, **58**, 99
- Sánchez, S. F., Pérez, E., Sánchez-Blázquez, P., et al. 2016, *Rev. Mex. Astron. Astrofis.*, **52**, 171
- Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521
- Schneider, P., Ehlers, J., & Falco, E. E. 1992, *Gravitational Lenses* (Springer)
- Shapiro, I. I. 1964, *Phys. Rev. Lett.*, **13**, 789
- Somerville, R. S., & Davé, R. 2015, *ARA&A*, **53**, 51
- Steinhardt, D. N., Verde, L., Peiris, H. V., et al. 2003, *ApJS*, **148**, 175
- Spergel, D. N., Bean, R., Doré, O., et al. 2007, *ApJS*, **170**, 377
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, **328**, 726
- Springel, V., White, S. D. M., Jenkins, A., et al. 2005, *Nature*, **435**, 629
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, **475**, 676
- Steidel, C. C., Rudie, G. C., Strom, A. L., et al. 2014, *ApJ*, **795**, 165
- Steinhardt, P. J., Wang, L., & Zlatev, I. 1999, *Phys. Rev. D*, **59**, 123504
- Sunyaev, R. A., & Zeldovich, Y. B. 1972, *Comm. Astrophys. Space Phys.*, **4**, 173
- Teimoorinia, H., Bluck, A. F. L., & Ellison, S. L. 2016, *MNRAS*, **457**, 2086
- Tempel, E., Kipper, R., Tamm, A., et al. 2016, *A&A*, **588**, A14
- Tremmel, M., Karcher, M., Governato, F., et al. 2017, *MNRAS*, **470**, 1121
- Trujillo-Gomez, S., Klypin, A., Primack, J., & Romanowsky, A. J. 2011, *ApJ*, **742**, 16
- Turner, E. L., Ostriker, J. P., & Gott, J. R., III. 1984, *ApJ*, **284**, 1
- Vale, A., & Ostriker, J. P. 2004, *MNRAS*, **353**, 189
- Vegetti, S., Koopmans, L. V. E., Auger, M. W., Treu, T., & Bolton, A. S. 2014, *MNRAS*, **442**, 2017
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014a, *MNRAS*, **444**, 1518
- Vogelsberger, M., Genel, S., Springel, V., et al. 2014b, *Nature*, **509**, 177
- Walsh, M. J. 1979, *AIAA J.*, **17**, 783
- Warren, M. S., Abazajian, K., Holz, D. E., & Teodoro, L. 2006, *ApJ*, **646**, 881
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291
- Weinberger, R., Springel, V., Pakmor, R., et al. 2018, *MNRAS*, **479**, 4056
- Wolf, J., Martinez, G. D., Bullock, J. S., et al. 2010, *MNRAS*, **406**, 1220
- Woo, J., Dekel, A., Faber, S. M., et al. 2013, *MNRAS*, **428**, 3306
- Xue, X. X., Rix, H. W., Zhao, G., et al. 2008, *ApJ*, **684**, 1143
- Yang, X., Mo, H. J., van den Bosch, F. C., et al. 2007, *ApJ*, **671**, 153
- Yang, X., Mo, H. J., & van den Bosch, F. C. 2009, *ApJ*, **695**, 900
- Zheng, Z., Coil, A. L., & Zehavi, I. 2007, *ApJ*, **667**, 760
- Zinger, E., Pillepich, A., Nelson, D., et al. 2020, *MNRAS*, **499**, 768
- Zwicky, F. 1933, *Helv. Phys. Acta*, **6**, 110

## Appendix A: Guide to DfL

In this appendix we provide a brief guide to the PYTHON package, Dark from Light (DfL), which is publicly available for use.<sup>18</sup>

This software provides a fast, efficient, and straightforward method to go from standard spectroscopic survey data on galaxies to a physically motivated group catalog, with halo mass estimates and virial properties of the groups. See Section 3 for a full discussion of the pipeline mechanics. Here we focus on the practicalities for using and running DfL.

Required function arguments:

`input_data_path` : str

A string specifying a path to the input data file, either in ASCII or FITS format. A template for input column names, formats, and units is presented in Table A.1.

`output_data_path` : str

A string specifying a path to the output file in FITS format. A summary of output names, formats, and units of columns appended to the input file is listed in Table A.2.

`model_type`: str

A string specifying regression model to use in Random Forest halo mass regression. Choice among "HYBRID", "EAGLE", and "TNG", where "HYBRID" yields an average prediction between "EAGLE" and "TNG". This enables one to construct model-dependent halo mass estimates which are optimized to either EAGLE or IllustrisTNG, or else opt simply for the average of the two.

`seed`: int

An integer specifying the seed for random number generation. Required to ensure reproducibility.

`include_error`: bool

A boolean switch indicating a choice between Case 1 and Cases 2 & 3 in Section 3.

Optional arguments:

In addition to the five necessary function arguments, a user can specify the following optional features:

`hubble_param0`: float [default = 70.0]

Hubble parameter value at redshift  $z = 0$  in units of km/s/Mpc for a spatially flat  $\Lambda$ CDM cosmology

`omega_matter0`: float [default = 0.3]

Mean redshift  $z = 0$  fractional energy density of all forms of matter for a spatially flat  $\Lambda$ CDM cosmology

`verbose`: bool [default is False]

A boolean verbosity switch.

`convergence_tol`: float [default = 0.05]

Required tolerance for halo mass convergence in the iterative part of the group finding algorithm in units of dex.

`iteration_limit`: int [default = 10]

Maximum number of iterations allowed per halo finding loop.

`num_montecarlo` = int [default = 100]

Number of random draws for Monte Carlo stellar mass error propagation (Cases 2 and 3).

`zrange`: tuple[float, float] [default is None]

Redshift range to impose on the input data.

## Appendix B: Halo mass estimation with FOF-AM

In this subsection we present the FOF-AM method, which stands for abundance matching applied to total group stellar masses, extracted from friends-of-friends group finding. A comparison of performance results in TNG and EAGLE from FOF-AM to DfL is provided in Table 2. Here we discuss the details of this methodology and present additional training and testing plots. We first discuss the optimization of the FOF algorithm (Appendix B.1.1), next we discuss the AM method (Appendix B.1.2), and then we discuss combining the two techniques (Appendix B.1.3). Thereafter we present performance tests of FOF-AM in halo mass estimation, and compare to DfL (Appendix B.2).

### B.1. FOF-AM method

The FOF-AM method is essentially a three stage process. First we optimize the linking lengths in FOF for application to observational-like simulated data from TNG and EAGLE. Second, we learn a one-to-one mapping between total group stellar masses and halo masses in the simulations, via matching on cumulative comoving number densities. Finally, we apply the FOF algorithm followed by the abundance matching procedure to the observational-like testing data, in order to extract both group information and halo mass estimates.

#### B.1.1. FOF group finder

We follow an approach to FOF similar to Gerke et al. (2005), Knobel et al. (2009); and Looser et al. (2021). Galaxies are grouped together if they satisfy a linking-length criterion:

$$l_{\text{sky}} < l_{\text{sky lim.}} \quad \& \quad l_z < l_z \text{ lim.} \quad (\text{B.1})$$

Different linking lengths are used on sky and in redshift space to account for the famous ‘fingers of god’ effect (e.g., Knobel et al. 2009). Moreover, we use physical distances for sky and velocities for redshift linking lengths. These limits are usually determined as some fraction of the average separation between galaxies in a survey, which allows it to be applied to other surveys with different depths. However, since we are running this on simulated data (which is stellar mass complete) we opt to optimize directly on physical distances and velocities. In essence, we test the potential of FOF to constrain groups in observer space (2D+z), compared to the full 6D phase space.

In order to optimize the linking lengths we must choose a parameter to be optimized. Here we select the overall accuracy of central - satellite classification. We then systematically vary the on-sky physical and redshift linking lengths, running the FOF group finder for each configuration. We select the peak of the optimization curves as the optimal values to use in FOF-AM for comparison to DfL.

<sup>18</sup> <https://github.com/abluck/Dark-from-Light>

**Table A.1.** INPUT DATA FORMAT

| Column Name                                    | Format       | Units                   | Comments:                             |
|--|--------------|-------------------------|---------------------------------------|
| GAL_ID   | int          | [Unitless]              | Unique Galaxy Identifier              |
| Z_SPEC   | float        | [Unitless]              | Spectroscopic Redshift                |
| RA   | float        | [DEG]                   | Right Ascension of Galaxy             |
| DEC  | float        | [DEG]                   | Declination of Galaxy                 |
| STELLAR_MASS                                   | float        | $[\log_{10} M_{\odot}]$ | Stellar Mass of Galaxy                |
| — stellar mass uncertainty choices, optional — |              |                         |                                       |
| STELLAR_MASS_ERROR                             | float        | $[\log_{10} M_{\odot}]$ | $1\sigma$ Point Error on Stellar Mass |
| STELLAR_MASS_PDF_BINS                          | array[float] | $[\log_{10} M_{\odot}]$ | PDF Bin Centers                       |
| STELLAR_MASS_PDF_VALS                          | array[float] | [Unitless]              | PDF Probabilities                     |

**Table A.2.** OUTPUT DATA FORMAT

| Column Name        | Format       | Units                   | Comments:                                 |
|--------------------|--------------|-------------------------|---|
| groupID            | int          | [Unitless]              | Unique Group Identifier                   |
| group_total_Mstar  | float        | $[\log_{10} M_{\odot}]$ | Total Group Stellar Mass                  |
| group_count        | float        | [Unitless]              | Number of Galaxies in Group               |
| Mhalo              | float        | $[\log_{10} M_{\odot}]$ | Predicted Group Halo Mass ( $M_{200}$ )   |
| Mhalo_2sigma_neg   | float        | $[\log_{10} M_{\odot}]$ | $2\sigma$ Negative Uncertainty            |
| Mhalo_1sigma_neg   | float        | $[\log_{10} M_{\odot}]$ | $1\sigma$ Negative Uncertainty            |
| Mhalo_1sigma_pos   | float        | $[\log_{10} M_{\odot}]$ | $1\sigma$ Positive Uncertainty            |
| Mhalo_2sigma_pos   | float        | $[\log_{10} M_{\odot}]$ | $2\sigma$ Positive Uncertainty            |
| Mhalo_pdf          | array[float] | [Unitless]              | Output Probabilities for Halo Mass PDF    |
| Mhalo_pdf_bins     | array[float] | $[\log_{10} M_{\odot}]$ | Output Bin Centers for Halo Mass PDF      |
| Rvir               | float        | [kpc]                   | Virial Radius of Group ( $R_{200}$ )      |
| Vvir               | float        | [km/s]                  | Virial Velocity of Group ( $V_{200}$ )    |
| Tvir               | float        | $[\log_{10} K]$         | Virial Temperature of Group ( $T_{200}$ ) |
| galaxy_class       | int          | [Unitless]              | 1 = Cen.; 0 = Core Sat.; -1 = Ext. Sat.   |
| num_iteration      | int          | [Unitless]              | No. of DfL Iterations                     |
| Mhalo_central_only | float        | $[\log_{10} M_{\odot}]$ | Halo Mass Prediction from Central Only    |

In Fig. B.1 we plot the FOF output central - satellite classification accuracy (as defined in eq. 28) as a function of  $l_{\text{sky}}$ . We additionally split each curve into different values of  $l_z$  (as labeled by the legends). We run this optimization procedure separately at different redshift snapshots (columns in Fig. B1) for TNG (top row) and EAGLE (bottom row). The peak of the optimization curve is then used to determine the optimal linking-lengths for each bespoke run of FOF-AM. The final sky and redshift limits used in our analysis are presented in each panel.

### B.1.2. AM technique

For the abundance matching technique, we follow an approach similar to Behroozi et al. (2010) and Moster et al. (2010). In Fig. B.2 we show the cumulative comoving number densities for TNG galaxies as a function of central stellar mass (left panels), group total stellar mass (center panels), and halo mass (explicitly  $M_{200}$ , right panels). Results are shown for  $z = 0$  (top row),  $z = 1$  (middle row), and  $z = 2$  (bottom row). We also perform an analogous procedure in EAGLE (not shown here for the sake of brevity).

Leveraging the basic assumption of abundance matching, i.e., that the most massive galaxies (or groups) reside in the most

massive haloes, one can link stellar and halo properties as follows. For central galaxies we define the AM halo mass as


$$M_{\text{Halo-AM},i} = M_{\text{Halo List}} [n(M_{\text{Halo List}}) = n(CM_{*,i})], \quad (\text{B.2})$$

and for groups we define the AM halo mass as

$$M_{\text{Halo-AM},i} = M_{\text{Halo List}} [n(M_{\text{Halo List}}) = n(GM_{*,i})], \quad (\text{B.3})$$

where  $M_{\text{Halo-AM},i}$  indicates the AM estimate of halo mass for galaxy  $i$ ,  $M_{\text{Halo List}}$  indicates the list of halo masses from the cumulative distribution function (CDF) plot, and  $n(M_{\text{Halo List}})$  indicates the CDF number densities for each halo mass in the list. In eq. B2,  $n(CM_{*,i})$  indicates the number density of the central galaxy stellar mass in question. In eq. B3,  $n(GM_{*,i})$  indicates the number density of the group total stellar mass in question.

The AM procedure is much easier to understand visually than in abstract mathematics. For central galaxy abundance matching, to go from a stellar to a halo mass, simply read off the number density of the stellar mass from the CDF. Then read off the halo mass which has the same number density. Similarly, for total group stellar mass abundance matching, simply read off the number density of the total group stellar mass from the CDF. Then read off the halo mass which has the same number density. This procedure is illustrated with dashed and dotted lines



**Fig. B.1.** Optimization of linking lengths for the friends-of-friends (FOF) group finding algorithm. The top row shows results for TNG; the bottom row shows results for EAGLE. Each column displays a redshift snapshot (as labeled for each panel). Separate sky ( $l_{\text{sky}}$ ) and redshift ( $l_z$ ) lengths are used as the maximum threshold for linking galaxies together in the FOF procedure. The X-axis displays varying  $l_{\text{sky}}$  values, which are measured in physical distances, with varying  $l_z$  values shown by different colored lines (see legends), which are measured in velocity space. Optimization is performed with respect to the total accuracy of central - satellite classification (shown on the Y-axis). The optimal values are presented in each panel.

in Fig. B.2. We note that for low-mass galaxies (and groups) the central and group AM procedures typically lead to similar (or identical) halo mass estimates. However, for high-mass galaxies (and groups), the central and group AM halo mass predictions typically differ significantly.

On the panels of Fig. B.2 we display the bias and standard deviation of halo mass recovery from AM for all groups (first number) and for multi-galaxy groups (second number, after comma). We show these results for central galaxy AM on the left panels, and for group AM on the center panels. It is clear that group abundance matching leads to superior results than central galaxy abundance matching, especially for multi-galaxy groups (see also, e.g., Yang et al. 2007, 2009). These results can be compared to those obtained in RF and ANN in Table 1 in the main body of this paper.

### B.1.3. Combining FOF & AM

Having optimized the linking lengths for FOF (see Appendix B.1.1) and having learned a mapping between central and group stellar mass to halo mass in simulations via AM (see Appendix B.1.2), we are now in position to infer halo masses via FOF-AM. We first run FOF on the observational-like testing sample from each simulation at  $z = 0, 1, 2$ . We then extract the group total stellar mass from the FOF group catalogs. Thereafter, we use AM to convert the group total stellar masses into estimated halo masses. We also keep track of the central - satellite classification through FOF for further analyses. In the remainder of this appendix we explore the performance in halo mass estimation of FOF-AM, and compare to DfL.

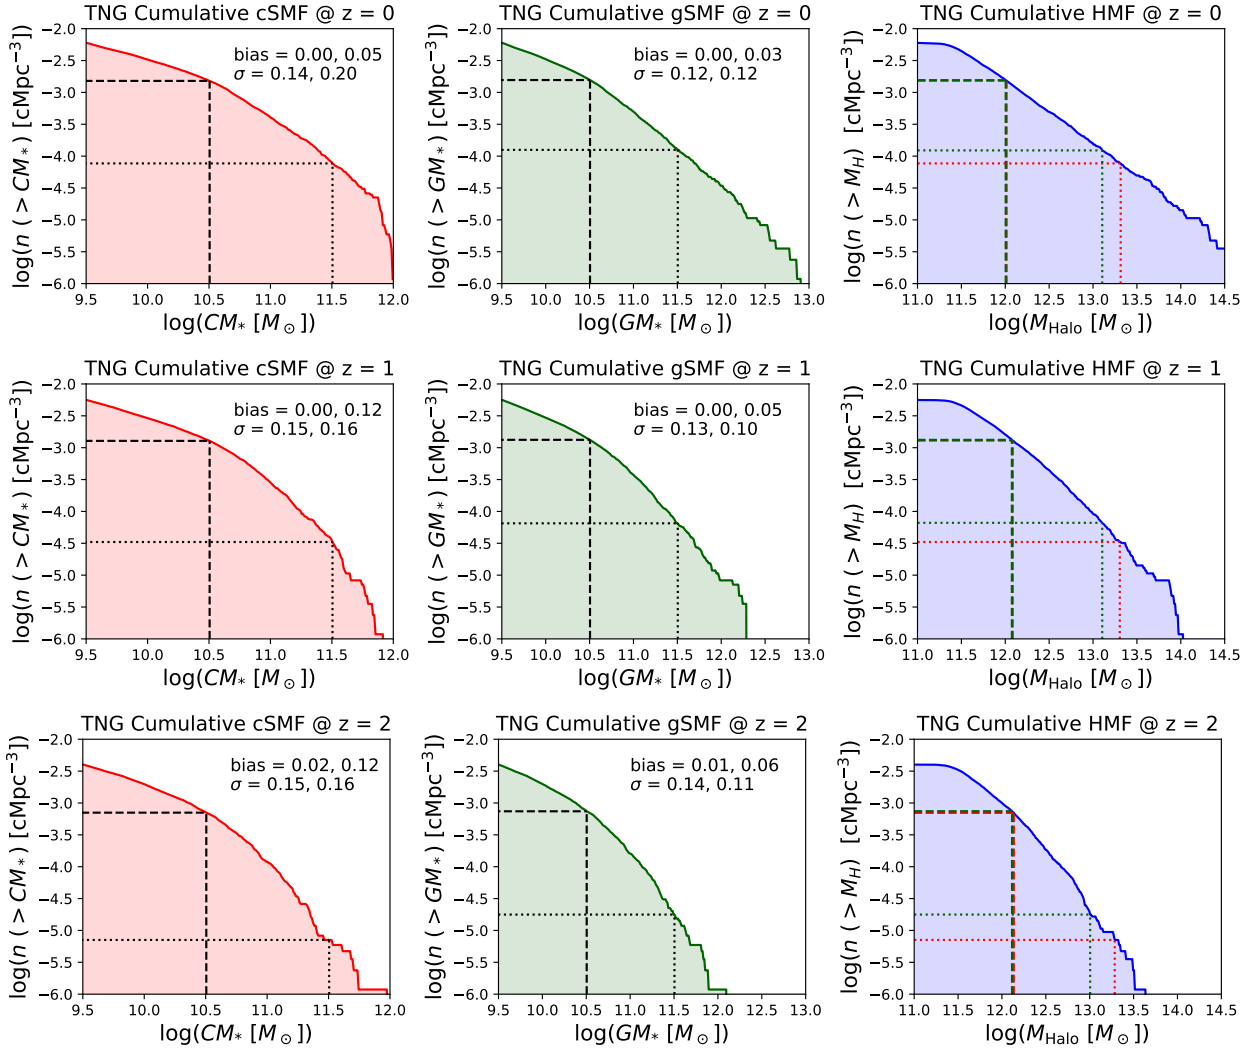
### B.2. FOF-AM performance results

In Figs. B.3, B.4 & B.5 we reproduce Figs. 6, 7 & 9 for the FOF-AM method. In Fig. B.3 we plot predicted vs. true halo masses derived from the FOF-AM methodology. As in DfL (see Fig. 6), we find relatively narrow contour distributions centered around the one-to-one relation. However, we note that the correlation strengths are significantly lower, and the catastrophic outlier fractions are significantly higher, than in DfL (see statistics presented on panels). However, the correlation strengths rise to similar values to that found in DfL in the case of excluding catastrophic outlier (statistics shown in parenthesis). This indicates that FOF-AM is similarly accurate to DfL for predicting halo masses for the majority of the sample, but is significantly worse for a relatively small subsample ( $\sim 5\%$  of the total).

In Fig. B.4 we plot the distributions of  $\delta M_{200}$ , as defined in eq. 27 (but here for the FOF-AM method). The structure of Fig. B.4 is identical to Fig. 7. We find that the FOF-AM method yields essentially bias-free results for the same data and model pairings, as in DfL. Additionally, we find similar (although slightly broader)  $\sigma_1$  values in FOF-AM, compared to DfL. Most interestingly, we find that the  $\sigma_2$  values are much higher from FOF-AM than in DfL. This is a direct consequence of FOF-AM having higher catastrophic outlier fractions than DfL (see also Fig. B.3). This is discussed in Section 4.

In Fig. B.5 we show the FOF-AM estimated halo mass functions compared to the simulation truth. As in DfL, visually there is a reasonable correspondence between the two distributions. However, looking in more detail one can immediately see that DfL achieves a higher fidelity (as confirmed by the offset statistics displayed in each panel in comparison to Fig. 9).

Finally, in Fig. B.6 we compare the confusion matrices for FOF-AM against DfL directly, for the case of matched model



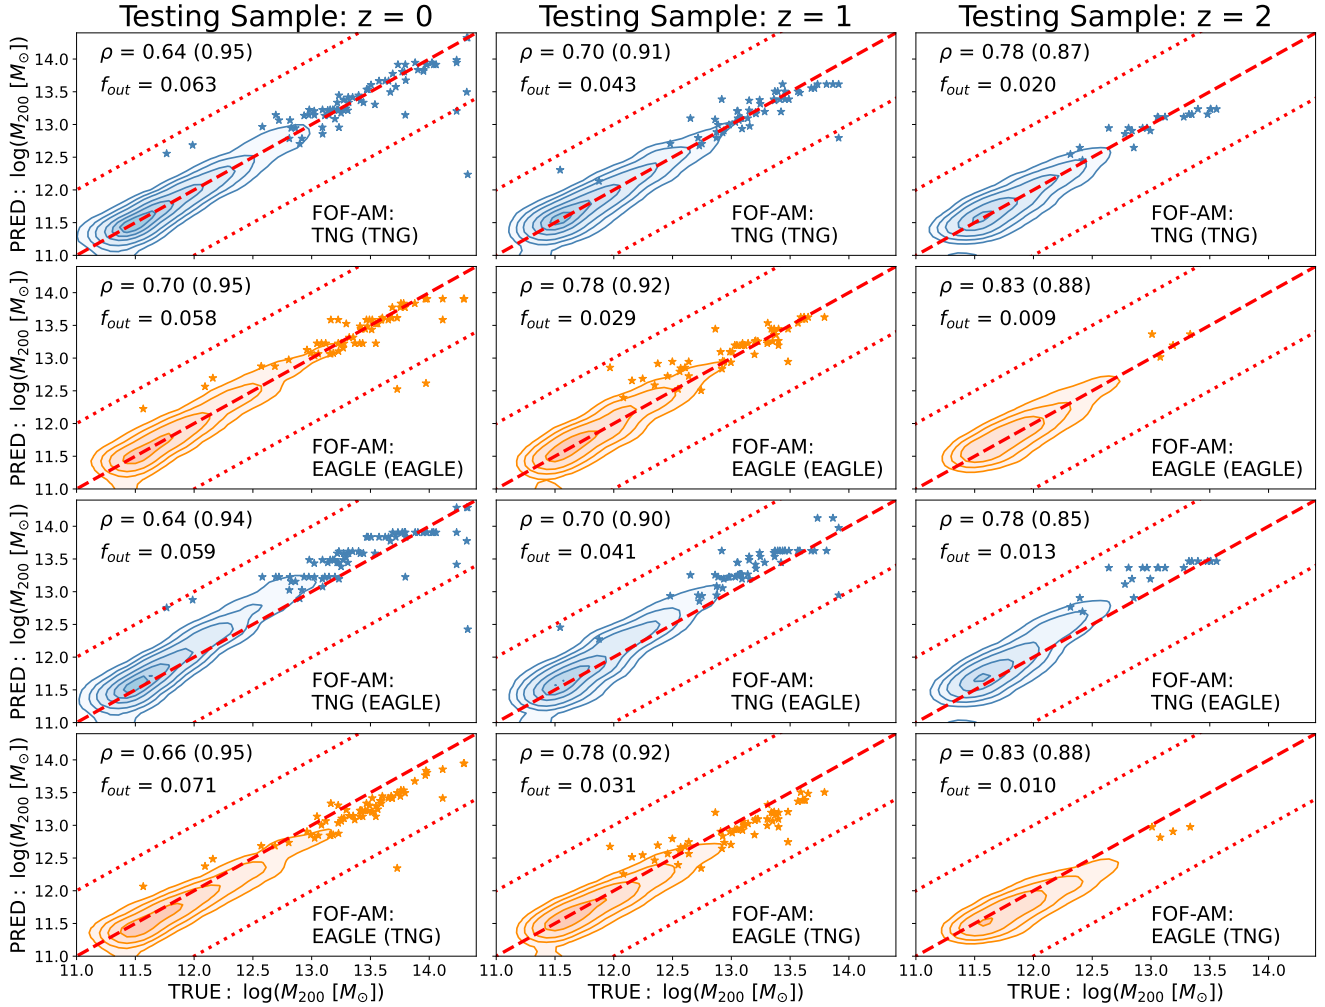
**Fig. B.2.** Illustration of the abundance matching approach applied to centrals and groups for the TNG simulation. The rows show results from different redshift snapshots, and the columns indicate (from left to right) the cumulative central stellar mass function, the cumulative group total stellar mass function, and the cumulative halo mass function (explicitly computed for  $M_{200}$ ). In abundance matching the cumulative number densities of either the central stellar mass or the group total stellar mass are matched to the cumulative number density of the halo mass function. Dashed and dotted lines indicate this matching procedure, and illustrate how the matching between central and group stellar mass is not in general identical (especially at high masses). In this method we train and test on different subsamples from within each simulation. Indicated on the stellar mass function panels are the bias and standard deviation in halo mass recovery from abundance matching within each approach. Additionally, the performance is given for all groups (first value) and for multi-galaxy groups (second value, after comma). Performing abundance matching on the total group stellar mass is superior to with the central stellar mass, and this is particularly evident for multi-galaxy groups (as one might expect).

and data pairings. In the case of central - satellite classification in DfL we found no significant differences with the mismatched model - data pairings, so we exclude that here for the sake of brevity. Here we present the full central - (all) satellite confusion matrices for DfL because this is a fairer comparison to FOF-AM, in which selecting subregions within proto-groups is not well motivated (see Section 4 for further discussion on this).

In Fig. B.6 we find that the overall accuracy in central - satellite classification is similar between DfL and FOF-AM. However, the precision in central galaxy identification is notably higher in DfL compared to FOF-AM. Ultimately, this helps to explain the enhanced performance of DfL in the wings of the  $\delta M_{200}$  distributions (commented on above). In the case of DfL, fewer groups are erroneously constructed than in FOF-AM. The price for this, however, is that DfL performs a little worse than FOF-AM in terms of recall (i.e., completeness in the central galaxy population). Nonetheless, we emphasize that this limi-

tation has little impact on the recovery of halo mass functions, which are in excellent accord with the simulation truth in DfL (see back to Fig. 9).

In summary, DfL and FOF-AM achieve similar halo mass recovery in terms of bias and  $\sigma_1$ . However, DfL significantly outperforms FOF-AM in terms of  $\sigma_2$  and the  $f_{\text{out}}$ . Hence, in terms of halo mass estimation, DfL is overall superior to FOF-AM as a methodology. Considering central - satellite classification, DfL and FOF-AM perform similarly well in overall accuracy. DfL is superior to FOF-AM in precision, yet FOF-AM is superior to DfL in recall. However, given the ability to remove the extended satellite population in DfL (which is not a natural feature of FOF-AM), one can recover a much higher performance in accuracy, precision, and recall in DfL than in FOF-AM (see Fig. 10). This does come at the price of reducing the completeness of the satellite population, yet for many analyses this is not especially problematic.



**Fig. B.3.** Identical in structure to Fig. 6, but here showing results for halo mass estimation from the FOF-AM approach. Generally, a good recovery of halo masses is obtained by this approach, but the performance is significantly poorer than with DfL. We also note that similar biases are seen when changing the model - data pairing in FOF-AM to DfL. No uncertainties are shown in this figure since the abundance matching technique does not provide an obvious framework for error propagation when applied to groups (unlike in the DfL method).

A side-by-side comparison of performance in halo mass estimation between FOF-AM and DfL is presented in Table 2, in the main body of the paper.

### Appendix C: Machine learning reproducibility

To facilitate reproduction of our machine learning results by others, we provide in this appendix details on the hyper-parameters used in the RF (see Table C.1) and ANN (see Table C.2) analyses. All regression models are constructed using SCIKIT-LEARN<sup>19</sup> (Pedregosa et al. 2011).

In the group and all parameter runs we set: `max_features='sqrt'`. This ensures that trees within the forest are different to one another by both random sampling of the available features and through bootstrapped random sampling of the training data (with return). In the case of the central galaxy runs (which have only two parameters: stellar mass and redshift), we set `max_features='None'`. This ensures that both parameters are always available to decision trees in training. Nonetheless, the trees remain distinct due to the bootstrapped random sampling.

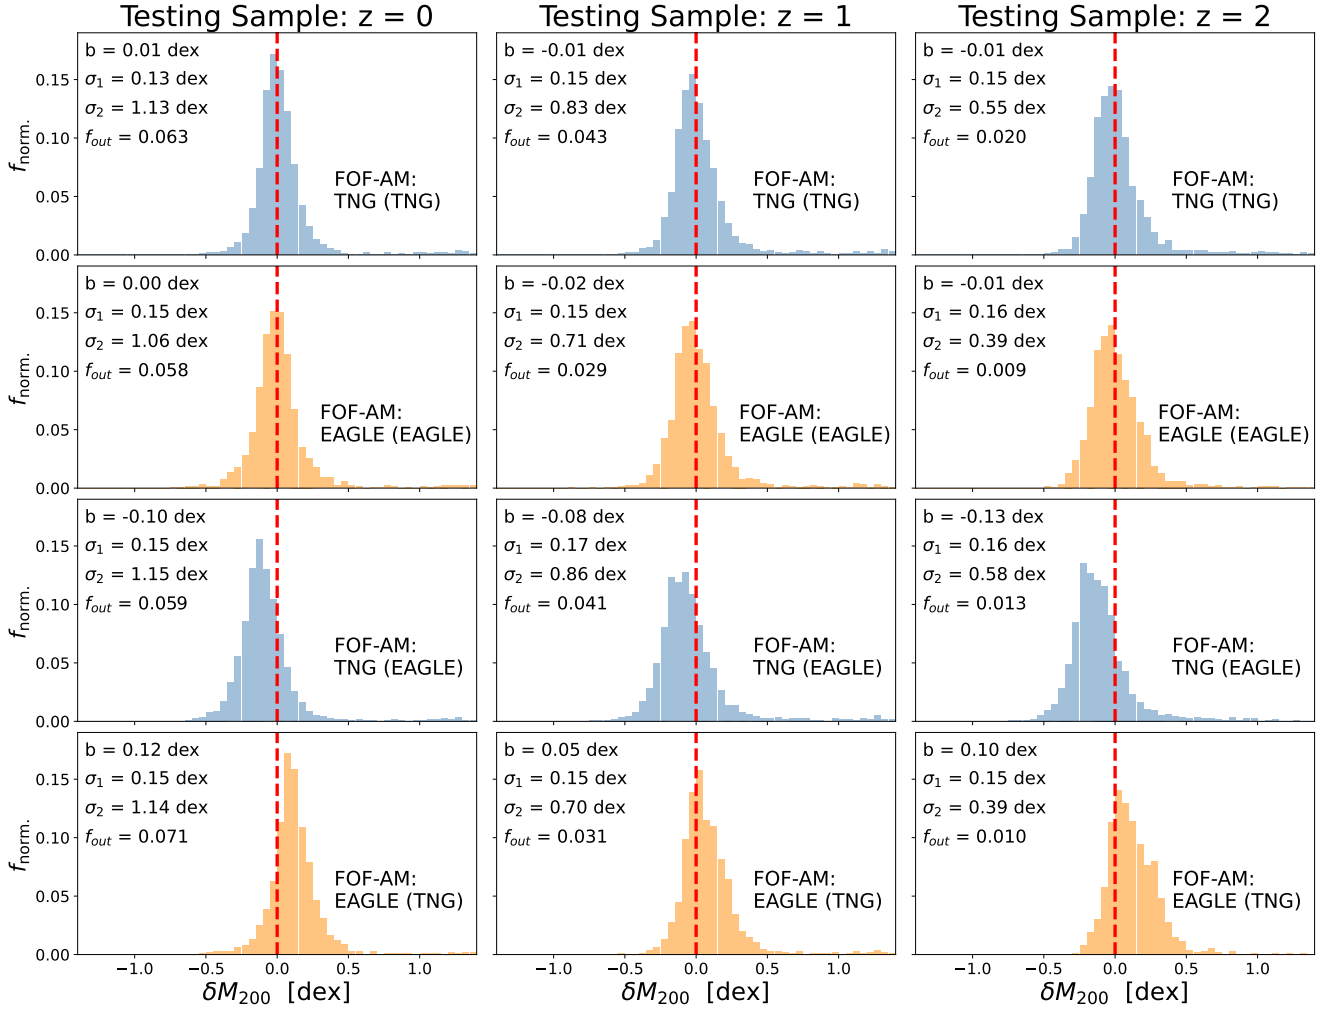
<sup>19</sup> <https://scikit-learn.org>

In Bluck et al. (2022), we have demonstrated that the all parameter ('None') mode represents the most effective RF architecture for disentangling inter-correlated 'nuisance' parameters from the underlying causal relations in simple models, simulations, and observations. However, we note that this is not the default method of RF classification in SCIKIT-LEARN. In this work, we opt to use the more conventional ('sqrt') setting for two reasons. First, we seek the most stable predictions, and this mode is less sensitive to individual data. Second, this mode yields a more thorough exploration of the predicted parameter space, which is advantageous for estimation and propagation of uncertainties.

The full list of hyper-parameters used in RF analyses are listed in Table C.1, with the full list of hyper-parameters used in ANN analyses listed in Table C.2. In the remainder of this appendix we explain what the table headings mean. The table columns are as follows:

Method: Type of machine learning used. Here either RF Regression (Table C.1) or ANN Regression (Table C.2).

$N_{\text{trees}}$ : The number of decision trees used in the RF model.



**Fig. B.4.** Identical in structure to Fig. 7, but here showing results for halo mass estimation from the FOF-AM approach. While many halo masses are accurately estimated via this technique, a significantly higher fraction of galaxies are found to be catastrophic outliers. This in turn leads to less accurate results across redshift ranges than seen in DfL (compare to Fig. 7, especially in terms of  $\sigma_2$ ).


**Table C.1.** RF hyper-parameters for DfL

| Simulation      | Method        | $N_{\text{trees}}$ | Max Depth | Max Features | T : V   | MSL | $\sigma$ | $\Delta\sigma$ |
|-----------------|---------------|--------------------|-----------|--------------|---------|-----|----------|----------------|
| TNG (Central)   | RF Regression | 250                | 1000      | None         | 70 : 30 | 8   | 0.15     | 0.03           |
| TNG (Group)     | RF Regression | 250                | 1000      | 'sqrt'       | 70 : 30 | 5   | 0.12     | 0.02           |
| TNG (All)       | RF Regression | 250                | 1000      | 'sqrt'       | 70 : 30 | 5   | 0.11     | 0.03           |
| EAGLE (Central) | RF Regression | 250                | 1000      | None         | 70 : 30 | 10  | 0.15     | 0.02           |
| EAGLE (Group)   | RF Regression | 250                | 1000      | 'sqrt'       | 70 : 30 | 5   | 0.13     | 0.03           |
| EAGLE (All)     | RF Regression | 250                | 1000      | 'sqrt'       | 70 : 30 | 5   | 0.12     | 0.03           |

**Max Depth:** The number of decision forks permitted within each decision tree in the RF models.

**Max Features:** The maximum number of features sent to each individual tree in the models. Here we use 'None' for centrals and 'sqrt' for all other data sets. The latter indicates that the integer closest to the square root of the total features is used.

**Min-Samples-Leaf (MSL):** The MSL hyper-parameter sets the minimum number of data in a given node of a given decision tree needed to attempt further splits in the data. We optimize MSL to maximize accuracy in the validation sample, while attempting to avoid significant over-fitting through comparison in performance to the validation sample. In all models and data sets we require that the difference in performance between training validation is less than 0.03 dex, which we achieve via systematically varying MSL. This parameter is used only in the RF models.



**Fig. B.5.** Identical in structure to Fig. 9, but here showing results for halo mass estimation from the FOF-AM approach. Clearly, the FOF-AM approach recovers reasonable estimates of the halo mass functions across epochs. However, DfL achieves a significantly tighter reproduction (as quantified by the mean offsets and standard deviations, displayed in each panel, compared to Fig. 9).

**Table C.2.** ANN hyper-parameters for comparison to DfL method

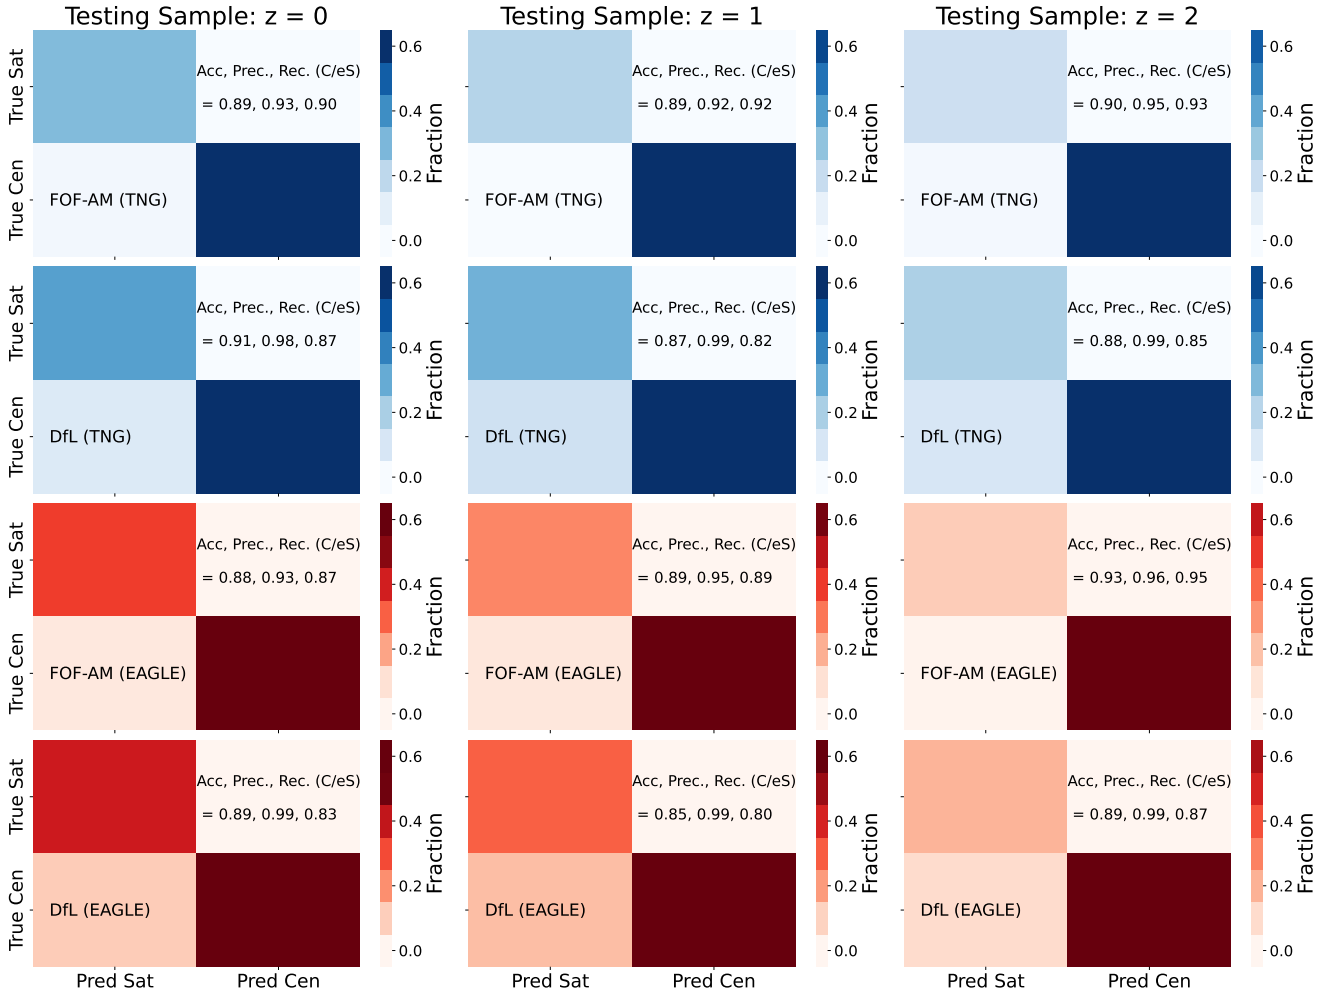
| Simulation      | Method   | Solver  | Activation  | Hidden Layers  | Learning Rate | Early Stop | T : V   | $\sigma$ | $\Delta\sigma$ |
|-----------------|----------|---------|-------------|----------------|---------------|------------|---------|----------|----------------|
| TNG (Central)   | ANN Reg  | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.14     | 0.00           |
| TNG (Group)     | ANN Reg  | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.12     | 0.00           |
| TNG (All)       | ANN Reg  | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.12     | 0.01           |
| EAGLE (Central) | ANN Reg. | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.14     | 0.01           |
| EAGLE (Group)   | ANN Reg. | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.13     | 0.01           |
| EAGLE (All)     | ANN Reg. | 'lbfgs' | <i>tanh</i> | (10 : 50 : 10) | 0.001         | True       | 70 : 30 | 0.12     | 0.00           |

**Solver:** The chosen method to update weights in the ANN back-propagation process. Here we use 'lbfgs' for all ANN runs, which is an optimizer in the quasi-Newtonian method family. We found the best ANN performance in preliminary testing with this mode. However, we note that other methods perform similarly (e.g., ADAM and 'sgd').

**Activation:** The activation function for neuron firing in the ANN. Here we use the *tanh* function throughout the ANN

models. This function is often faster to train with than other popular choices.

**Hidden Layers:** The hidden layers of the fully connected ANN. In preliminary testing we explored a wide variety in complexity of the ANN structure. We found that our chosen structure (10 : 50 : 10) performed well on all data sets and led to no significant over-fitting. Increasing either the depth or complexity of the network does not lead to higher performance



**Fig. B.6.** Confusion matrices from FOF-AM compared to DfL for the same model and data pairings. Here the performance statistics are shown for the full sample of all galaxies, with galaxies in the extended satellite region treated as normal satellites. The overall accuracy of both methods are similar, but DfL achieves notably higher precision in the central galaxy sample, which is important for accurately identifying haloes. On the other hand, the recall in central galaxies is lower in DfL compared to FOF-AM, indicating a trade off between purity and completeness.

on the validation samples. Equally, decreasing the depth or complexity does not significantly lower the difference in performance in training vs. validation. Hence, we are confident that this hidden layer structure is appropriate for the task at hand.

**Learning Rate:** The initial step size in gradient descent. We note that we utilize adaptive learning in all cases, so this is not fixed throughout the process of minimization.

**Early Stop:** In all ANN runs we enable early stopping, which permits the ANN solver to consider whether there is any improvement in a randomly generated 10% of the training sample (not used for training in that step), before deciding whether to continue in gradient descent further. This method is proven to reduce the potential for over-fitting, and clearly works well in these data (see  $\Delta\sigma$  values in table C.2).

**T : V (Train : Validate):** The ratio of data sizes used in training and validating. We note that this is completely separate to the 50% of the data which is converted into an observational-like form, and used to test DfL in Section 4. We emphasize again that the testing data are completely unseen in either training or validation.

$\sigma$ : The standard deviation of offsets in true-to-predicted halo masses in the validation sample (given in [dex]).

$\Delta\sigma$ : The difference in standard deviation between training and validation samples (given in [dex]).