



Full length article

The EAGLE simulations of galaxy formation: Public release of halo and galaxy catalogues[☆]

S. McAlpine^{a,*}, J.C. Helly^a, M. Schaller^a, J.W. Trayford^a, Y. Qu^a, M. Furlong^a, R.G. Bower^a, R.A. Crain^b, J. Schaye^c, T. Theuns^a, C. Dalla Vecchia^{d,e}, C.S. Frenk^a, I.G. McCarthy^b, A. Jenkins^a, Y. Rosas-Guevara^f, S.D.M. White^g, M. Baes^h, P. Camps^h, G. Lemsonⁱ

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

^b Astrophysics Research Institute, Liverpool John Moores University, 146 Brownlow Hill, Liverpool L3 5RF, UK

^c Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

^d Instituto de Astrofísica de Canarias, C/ Vía Láctea s/n, 38205 La Laguna, Tenerife, Spain

^e Departamento de Astrofísica, Universidad de La Laguna, Av. del Astrofísico Francisco Sánchez s/n, 38206 La Laguna, Tenerife, Spain

^f Departamento de Astronomía, Universidad de Chile, Casilla 36-D, Las Condes, Santiago, Chile

^g Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany

^h Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281, B-9000 Gent, Belgium

ⁱ Department of Physics & Astronomy, Johns Hopkins University, Baltimore, MD, 21218, USA

ARTICLE INFO

Article history:

Received 5 October 2015

Accepted 25 February 2016

Available online 17 March 2016

Keywords:

Cosmology: theory

Galaxies: formation

Galaxies: evolution

Method: numerical

ABSTRACT

We present the public data release of halo and galaxy catalogues extracted from the EAGLE suite of cosmological hydrodynamical simulations of galaxy formation. These simulations were performed with an enhanced version of the GADGET code that includes a modified hydrodynamics solver, time-step limiter and subgrid treatments of baryonic physics, such as stellar mass loss, element-by-element radiative cooling, star formation and feedback from star formation and black hole accretion. The simulation suite includes runs performed in volumes ranging from 25 to 100 comoving megaparsecs per side, with numerical resolution chosen to marginally resolve the Jeans mass of the gas at the star formation threshold. The free parameters of the subgrid models for feedback are calibrated to the redshift $z = 0$ galaxy stellar mass function, galaxy sizes and black hole mass–stellar mass relation. The simulations have been shown to match a wide range of observations for present-day and higher-redshift galaxies. The raw particle data have been used to link galaxies across redshifts by creating merger trees. The indexing of the tree produces a simple way to connect a galaxy at one redshift to its progenitors at higher redshift and to identify its descendants at lower redshift. In this paper we present a relational database which we are making available for general use. A large number of properties of haloes and galaxies and their merger trees are stored in the database, including stellar masses, star formation rates, metallicities, photometric measurements and mock *gri* images. Complex queries can be created to explore the evolution of more than 10^5 galaxies, examples of which are provided in the Appendix. The relatively good and broad agreement of the simulations with a wide range of observational datasets makes the database an ideal resource for the analysis of model galaxies through time, and for connecting and interpreting observational datasets.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Galaxy formation is a complex, non-linear process that involves a wide range of physical and astrophysical phenomena, from the

evolution of dark matter clustering to intricate feedback effects coupling gas cooling and outflows to star and black hole formation. Theoretical studies of galaxy formation thus require rigorous detailed modelling to link together these phenomena over a very wide range of scales. Two techniques have been developed for this purpose: semianalytic modelling (White and Frenk, 1991) and hydrodynamical simulations (Carlberg et al., 1990; Katz et al., 1992). Both techniques have been extensively developed over the past 25 years (e.g. Porter et al., 2014; Henriques et al., 2015; Lacey

[☆] <http://www.eaglesim.org/database.php>.

* Corresponding author.

E-mail addresses: s.r.mcalpine@durham.ac.uk (S. McAlpine), j.c.helly@durham.ac.uk (J.C. Helly), matthieu.schaller@durham.ac.uk (M. Schaller).

et al., 2015, for semi-analytic models) and (e.g. Oppenheimer et al., 2010; Puchwein and Springel, 2013; Dubois et al., 2014; Okamoto et al., 2014; Vogelsberger et al., 2014; Khandai et al., 2015, for hydrodynamical simulations).

Recently, the Virgo¹ Consortium’s “Evolution and Assembly of Galaxies and their Environments” simulation suite (EAGLE, Schaye et al., 2015; Crain et al., 2015) has been able to reproduce key observational datasets, such as the present-day stellar mass function of galaxies, the correlation of black hole mass and stellar mass and the dependence of galaxy sizes on stellar mass, with unprecedented fidelity. As well as reproducing these observations, which were used during the calibration of the simulation parameters, the simulation outputs match many other properties of the observed galaxy population and the intergalactic medium both at the present day and at earlier epochs, as we briefly discuss below. These simulations therefore provide a powerful resource for understanding the formation of galaxies and for linking and interpreting observational datasets.

The aim of this paper is to introduce and make available a relational database that can be queried using the Structured Query Language (SQL) to explore and exploit the halo and galaxy catalogues of the main EAGLE simulations. Columns containing integrated quantities describing the galaxies, such as stellar mass, star formation rates, metallicities and luminosities, are provided for more than 10^5 simulated galaxies and these can be individually followed through their evolution across cosmic time. This database is available at the address <http://www.eaglesim.org/database.php>.

The simulations follow the gravitational hydrodynamical equations, tracking the evolution of baryons and dark matter. The initial conditions reflect the small density fluctuations observed in the cosmic microwave background (CMB). By tracking the movement of baryon and dark matter particles, the simulations calculate how these fluctuations are amplified by gravity, and how pressure and radiative cooling of baryons separate these two matter components of the universe. The simulations include subgrid formulations to account for processes that cannot be directly resolved in the calculation and that describe how stars and black holes form and impact the matter distribution around them. EAGLE improves on previous hydrodynamical simulations of representative volumes, through the use of physically motivated subgrid source and sink terms as well as through the adoption of a clear strategy for the calibration of uncertain subgrid parameters (Crain et al., 2015) and by producing a galaxy population that reproduces many of the characteristics of the observed population over a wide range of redshifts.

The usability of the simulation data products is greatly enhanced when presented in a relational database, making it simple and quick to select galaxy samples based on multiple galaxy properties, to connect them to their halos and to follow their evolution over cosmic time (Lemson and Springel, 2006). Such databases were originally designed to host results from large surveys (e.g. the SDSS SkyServer Szalay et al., 2000) and later the halo catalogues from dark matter simulations and galaxy catalogues from semi-analytic models (applied to the *Millennium Simulation*, see Lemson and Virgo Consortium, 2006). They have since been expanded to include the wider range of data available from hydrodynamical simulations (e.g. Dolag et al., 2009; Khandai et al., 2015; Nelson et al., 2015). The database allows multiple indexing of the data that significantly enhances access speed and allows the selection of smaller data subsets that can be quickly analysed using simple scripting languages. This approach avoids

the need for the user to copy the raw simulation data or even just the full galaxy catalogues, reducing data transfer volumes to a manageable level. The galaxy properties stored in the database can be compared to observations or to other models, whilst the physics of galaxy formation can be explored by tracking an individual galaxy’s behaviour and environment through cosmic time.

This paper is intended as a reference guide for accessing the publicly available EAGLE database, and is laid out as follows. Section 2 presents a brief overview of the EAGLE simulation suite, including the list of simulations available in the database and the values of the subgrid parameters that vary, as well as an overview of the construction of the merger trees and database tables. A short tutorial describing how to access the data is presented in Section 3. We give some words of caution and some remarks on the simulations in Section 4 and conclude in Section 5. Some additional examples combining the PYTHON and SQL languages to access the data are given in Appendix A whilst the full list of galaxy and halo properties available in this data release is given in Appendix B together with a list of output redshifts in Appendix C and detailed equations given in Appendix D. Throughout this paper we quote magnitudes in the AB system and use ‘h-free’ units unless stated otherwise.

2. The EAGLE simulation suite

The EAGLE simulation suite is a set of cosmological hydrodynamical simulations in cubic, periodic volumes ranging from 25 to 100 comoving megaparsecs (cMpc) per side that track the evolution of both baryonic (gas, stars and massive black holes) and non-baryonic (dark matter) elements from a starting redshift of $z = 127$ to the present day. All simulations adopt a flat Λ CDM cosmology with parameters taken from the *Planck* mission (Planck Collaboration et al., 2014) results: $\Omega_{\Lambda} = 0.693$, $\Omega_m = 0.307$, $\Omega_b = 0.04825$, $\sigma_8 = 0.8288$, $n_s = 0.9611$, $Y = 0.248$ and $H_0 = 67.77 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (i.e. $h = 0.6777$). The initial conditions were generated using second-order Lagrangian perturbation theory (Jenkins, 2010) and the phase information is taken from the public PANPHISA Gaussian white noise field (Jenkins, 2013). Full details of how the ICs were made are given in Appendix B of Schaye et al. (2015). The simulation suite was run with a modified version of the GADGET-3 Smoothed Particle Hydrodynamics (SPH) code (last described by Springel, 2005), and includes a full treatment of gravity and hydrodynamics. The modifications to the SPH method are collectively referred to as ANARCHY (Dalla Vecchia, (in prep.), see also Appendix A of Schaye et al., 2015; Schaller et al., 2015a), and use the C_2 kernel of Wendland (1995), the pressure–entropy formulation of SPH of Hopkins (2013), the time-step limiters introduced by Durier and Dalla Vecchia (2012), the artificial viscosity switch of Cullen and Dehnen (2010) and a weak thermal conduction term of the form proposed by Price (2008). The effects of this state-of-the-art formulation of SPH on the galaxy properties are explored in detail by Schaller et al. (2015a).

2.1. Subgrid model

Processes not resolved by the numerical scheme are implemented as subgrid source and sink terms in the differential equations. For each process, schemes were adopted that are as simple as possible and that only depend on the local hydrodynamic properties. This last requirement differentiates EAGLE from most other cosmological, hydrodynamical simulation projects

¹ <http://virgo.dur.ac.uk/>.

Table 1
Parameters describing the available simulations. From left-to-right the columns show: simulation name suffix; comoving box size; total number of particles; initial baryonic particle mass; dark matter particle mass; comoving Plummer-equivalent gravitational softening length; maximum physical softening length and the subgrid model parameters that vary: $n_{\text{H},0}$, n_{n} , C_{visc} and ΔT_{AGN} (see Section 4 of Schaye et al., 2015, for an explanation of their meaning).

Identifier	L (cMpc)	N	$m_{\text{g}} (M_{\odot})$	$m_{\text{dm}} (M_{\odot})$	$\epsilon_{\text{com}} (\text{ckpc})$	$\epsilon_{\text{phys}} (\text{pkpc})$	$n_{\text{H},0} (\text{cm}^{-3})$	n_{n}	C_{visc}	$\Delta T_{\text{AGN}} (\text{K})$
Ref-L0025N0376	25	2×376^3	1.81×10^6	9.70×10^6	2.66	0.70	0.67	$2/\ln 10$	2π	$10^{8.5}$
Ref-L0025N0752	25	2×752^3	2.26×10^5	1.21×10^6	1.33	0.35	0.67	$2/\ln 10$	2π	$10^{8.5}$
Recal-L0025N0752	25	2×752^3	2.26×10^5	1.21×10^6	1.33	0.35	0.25	$1/\ln 10$	$2\pi \times 10^3$	$10^{9.0}$
Ref-L0050N0752	50	2×752^3	1.81×10^6	9.70×10^6	2.66	0.70	0.67	$2/\ln 10$	2π	$10^{8.5}$
AGNdT9-L0050N0752	50	2×752^3	1.81×10^6	9.70×10^6	2.66	0.70	0.67	$2/\ln 10$	$2\pi \times 10^2$	$10^{9.0}$
Ref-L0100N1504	100	2×1504^3	1.81×10^6	9.70×10^6	2.66	0.70	0.67	$2/\ln 10$	2π	$10^{8.5}$

(e.g. Oppenheimer et al., 2010; Puchwein and Springel, 2013; Vogelsberger et al., 2014; Khandai et al., 2015) and ensures that galactic winds develop without pre-determined mass loading factors and directions, without any direct dependence on halo or dark matter properties.

The simulation tracks the time-dependent stellar mass loss due to winds from massive stars and AGB stars, core collapse supernovae and type Ia supernovae (Wiersma et al., 2009b). Radiative cooling and heating is implemented element-by-element following Wiersma et al. (2009a). Cold dense gas is prevented from artificial fragmentation by implementing an effective temperature pressure floor as described by Schaye and Dalla Vecchia (2008). Star formation is implemented stochastically following the pressure-dependent Kennicutt–Schmidt relation (Schaye and Dalla Vecchia, 2008), with the inclusion of a metal-dependent star formation threshold designed to track the transition from a warm, atomic to an unresolved, cold, molecular gas phase, as proposed by Schaye (2004). The initial stellar mass function is that given by Chabrier (2003). Feedback from star formation is implemented thermally and stochastically following the method of Dalla Vecchia and Schaye (2012). Seed black holes are placed in haloes greater than a threshold mass of $10^{10} M_{\odot}/h$ and tracked following the methodology of Springel et al. (2005a) and Booth and Schaye (2009). Gas accretion onto black holes follows a modified version of the Bondi–Hoyle accretion rate, described by (Rosas-Guevara et al., 2015, but modified as described by Schaye et al., 2015), and feedback is implemented following the stochastic AGN heating scheme described by Schaye et al. (2015) and making use of the energy threshold of Booth and Schaye (2009). The details of the implementation and parametrisation of these schemes are motivated and described in detail by Schaye et al. (2015).

Because of our limited understanding of these processes and because of the limited resolution of the simulations, the subgrid source and sink terms involve free parameters whose values must be determined by comparison of the simulation results to a subset of the observational data. In the case of EAGLE, the subgrid parameters were calibrated for feedback from star formation and AGN by using three properties of galaxies at redshift $z = 0$, specifically the galaxy stellar mass function, the galaxy size–stellar mass relation, and the black hole mass–stellar mass relation. The calibration strategy is described in detail by Crain et al. (2015) who also presented additional simulations to demonstrate the effect of parameter variations.

Once the simulations have been calibrated using a subset of the observational data, they can be validated by comparison to additional datasets. Studies have so far shown that the simulations broadly reproduce a variety of other observables such as the $z = 0$ Tully–Fisher relation, specific star formation rates and the column density distribution of intergalactic C IV and O VI (Schaye et al., 2015), the H I and H₂ properties of galaxies (Bahé et al., submitted,

Lagos et al., 2015), the column density distribution of intergalactic metals (Schaye et al., 2015), galaxy rotation curves (Schaller et al., 2015b), the $z = 0$ luminosity function and colour–magnitude diagram (Trayford et al., 2015), the evolution of the galaxy stellar mass function (Furlong et al., 2015) and the high-redshift H I column density distribution (Rahmati et al., 2015).

2.2. The simulations in the database

Table 1 summarises the simulations that have been incorporated into the database, including the comoving cubic box length, baryonic and non-baryonic particle masses and gravitational softening lengths. Together these parameters determine the dynamic range and resolution that can be achieved by the simulation. The simulation name includes a suffix to indicate the simulation box length in comoving megaparsecs (e.g. L0100) and the cube root of the initial number of particles per species (e.g. N1504). Simulations with the same subgrid model as the primary run (L0100N1504) are denoted with the prefix “Ref-”. As discussed in Schaye et al. (2015), the “Recal-” higher-resolution simulation uses values of the subgrid parameters that have been recalibrated following the same procedure used for the reference simulation to improve the fit to the $z = 0$ galaxy stellar mass function, allowing the user to test the weak convergence of the code.² See Schaye et al. (2015) for definitions and discussion of the concepts of weak and strong convergence. Note that Recal-L0025N0752 should be compared to the Ref-L0025N0376 calculation to ensure that the same range of halo mass is sampled in both cases, eliminating differences due to the simulation volume. To a similar end, the Ref-L0025N0752 model is provided to allow the user to test the strong convergence of the results. This simulation uses all the subgrid parameters of the reference model but at a higher mass resolution. All the 25 cMpc volumes share the same large-scale initial fluctuations, so that objects appear in (approximately) the same spatial locations in all three runs.

Finally, the database also includes the additional simulation AGNdT9-L0050N0752 that uses a higher AGN heating temperature and increased black hole accretion viscosity parameter, C_{visc} . As discussed by Schaye et al. (2015), this results in a better match to the properties of diffuse gas in galaxy group haloes, but has only a small effect on the properties of galaxies. This simulation uses the same initial phases as the Ref-L0050N0752 model, allowing objects to be matched.

2.3. Halo, subhalo and galaxy identification

The raw particle data themselves are not required for many comparisons with observations. In order to reduce the volume

² As discussed in Section 4, performing convergence tests is strongly encouraged.

of data to be downloaded and simplify analysis, we process the simulation outputs individually to locate bound structures which we identify with galaxies and their associated dark matter haloes. The processing steps are described in detail by Schaye et al. (2015). In brief, overdensities of dark matter are identified using the “Friends-of-Friends” (FoF) method (Davis et al., 1985) adopting a linking length of 0.2 times the average inter-particle spacing. Baryonic particles are then assigned to the same FoF-halo as their closest dark matter neighbour. Self-bound “subhaloes”, which can contain both baryonic and dark matter, are later identified using the SUBFIND algorithm (Springel et al., 2001; Dolag et al., 2009) using all particle species.

It is important to note that particles are not shared between subhaloes so that the correspondence between particles and subhaloes is unique. We identify the baryonic component of each subhalo with a galaxy and will refer to them as such from now on. Resolved subhaloes always have a clear central concentration and there is a clear identification between the galaxies in the simulations and galaxies that would be identified in observational studies. Note that small subhaloes, especially at high redshift, may not contain any stars or even gas but will still be present in the catalogues. A FoF halo may contain several subhaloes (or subgroups in the SUBFIND terminology); we define the subhalo that contains the particle with the lowest value of the gravitational potential to be the *central galaxy* while any remaining subhaloes are classified as *satellite galaxies* (denoted `SubGroupNumber` = 0 and `SubGroupNumber` > 0 respectively in the database nomenclature, see below).

The stellar mass assigned to a galaxy may include diffuse particles at a large distance. Such particles make up the intra-cluster/intra-group light and would not normally be included in a galaxy’s photometry. We therefore also include aperture-based measurements in the database.

Exceptionally, SUBFIND may identify an internal high-density component of a galaxy as a distinct subhalo. Such spurious identifications are discussed in Section 4 and are labelled in the main database table with the field `Spurious`.

For each simulation we release 29 snapshot outputs between redshift 20 and 0 (the full list of released output redshifts is given in the Appendix Table C.1). We later analyse the properties of each subhalo in post-processing in order to calculate galaxy and subhalo properties, such as stellar masses, galaxy sizes, star formation rates and luminosities. Each subhalo and hence each galaxy is assigned an index, its `GalaxyID`, that allows one to identify an object uniquely both in space and time. Note that since the `GalaxyID` is unique to a particular output redshift, a galaxy will change its `GalaxyID` over time. The 29 catalogues of galaxies are then linked through time via a galaxy merger tree, allowing one to track the evolution of a galaxy (through the evolution of its `GalaxyID`) with time. The construction and structure of these trees is presented in Section 2.6.

2.4. Integrated quantities

At each redshift the galaxies are processed one-by-one to produce integrated quantities from the raw particle information. These are the quantities stored in the different tables of the database.

For the simplest quantities, such as galaxy mass, metallicity or star formation rate, the post-processing only involves a simple summation over the particles but other quantities, such as luminosities in various filters, require much more involved calculations. The full list of quantities present in the database, together

with a description of the post-processing operations performed, is given in Appendix B. To allow for an easier comparison with observational measurements, masses, star formation rates and velocity dispersions are also computed within fixed spherical apertures.

2.5. Mock gri images

For visualisation purposes, images are provided for galaxies with 30 physical kpc (pkpc) aperture stellar masses $> 10^{10} M_{\odot}$. Images are generated from mock observations made using the SKIRT code (Baes et al., 2003; Camps and Baes, 2015), with GALAXEV (Bruzual and Charlot, 2003) and MAPPINGS III (Groves et al., 2008) spectra to represent star particles and young HII regions respectively, as described by Trayford et al. (in prep.). A square field of view of 60 pkpc on a side is used for observations in the SDSS *gri* bands (Doi et al., 2010), with the galaxy spectra red-shifted to $z = 0.1$ to approximate SDSS colours. No artificial seeing is added to the images. Each galaxy above the stellar mass threshold is observed *face-* and *edge-on* to the galactic plane, defined using the stellar angular momentum vector within 30 pkpc. A ‘box’ projection is also provided, with galaxies viewed down the simulation z -axis and the horizontal and vertical image axes corresponding to the simulation x and y axes respectively. The three-colour *gri* images are prepared from the virtual SKIRT observations adopting the method of Lupton et al. (2004). Fig. 1 shows these three images for the same example galaxy.

2.6. Merger trees

As galaxies rarely evolve in isolation, they are subject to mergers with neighbouring galaxies. This adds serious complexity to tracing the history of an individual galaxy from the present-day to its formation and as such we must construct a *merger tree* to connect galaxies across simulation output times. Descendant subhaloes and hence galaxies are identified using the D-TREES algorithm (Jiang et al., 2014), with a complete description of its adaptation to the EAGLE simulations provided in Qu et al. (in prep.). In essence, the algorithm traces subhaloes using the N_{link} most bound particles of any species, identifying the subhalo that contains the majority of these particles as a subhalo’s descendant at the next output time. We define $N_{\text{link}} = \min(100, \max(0.1N_{\text{galaxy}}, 10))$, where N_{galaxy} is the total number of particles in the parent subhalo. This allows the identification of descendants, even in the case where most particles have been stripped and it minimises the misprediction of mergers during flybys (Fakhouri and Ma, 2008; Genel et al., 2009).

The galaxy with the most N_{link} particles at the next output is identified as the single *descendant* of a galaxy, while a descendant galaxy can have multiple progenitors. The trees are stored in memory following the method introduced by Lemson and Springel (2006) for the *Millennium Simulation* (see also the supplementary material of Springel et al., 2005b, where the details of the tree ordering are summarised). However, the *main progenitor*, corresponding to the main branch of the tree, is defined as the progenitor with the largest ‘branch mass’, i.e., the mass summed across all earlier outputs as proposed by De Lucia and Blaizot (2007). This definition of the main progenitor, as opposed to the simple definition of the progenitor with the largest mass, is used to avoid main branch swapping in the case of similar-mass mergers, as explained by Qu et al. (in prep.). Note that because the progenitor with the largest branch mass determines the main branch of the




Fig. 1. Mock *gri* images of a galaxy at $z = 0.1$ as available in the database. The left, central and right panels correspond to the `Image_face`, `Image_edge` and `Image_box` views (in the database nomenclature) of the same simulated galaxy (`GalaxyID = 16116800` in the Ref-L0100N1504 simulation). The images are 60 pkpc on a side. Note the clear presence of a bulge, of dust absorption and of spiral arms.

tree, main branch galaxies do not necessarily correspond to the central galaxy (or `SubGroupNumber = 0` galaxy) of a given halo.

There are two further aspects of the merger trees that must be kept in mind when analysing the simulation:

- A galaxy can disappear from a snapshot but reappear at a later time (e.g. if one galaxy passes through another one). To account for this, descendants are identified using up to 5 snapshots at later times.
- Care must be taken when determining mass ratios, for example in the case of mergers, as galaxies can lose or gain mass due to the definition of the subhaloes.

Both of these relatively rare cases are considered further by Qu et al. (in prep.), who discuss their impact on the assembly of galaxy mass.

2.7. Technical aspects and infrastructure

Multiple layouts and frameworks are available for storing large datasets (such as MONGODB,³ SciDB,⁴ HADOOP,⁵...) each coming with advantages and shortcomings. In the case of galaxy catalogues extracted from cosmological simulations, the *Millennium simulation* used an SQL database for its public release and the wider astronomy community has since developed a familiarity with its structure and way to query the data. To allow users the simplest transition between databases, we have adopted the same framework and a similar table design as the *Millennium simulation* SQL database (Lemson and Virgo Consortium, 2006). More efficient ways of querying the data could exist, with differing database formats or table structures, however we decided that maintaining the familiar aspects of previous databases outweighs the potential performance gains.

The server hosting the front end web interface operates on Centos linux, running APACHE TOMCAT 6.0.24. This server interfaces with the database host, submitting queries and having their results streamed via a Java web application (originally written for the

Millennium simulation). The database itself is stored on a single physical Windows Server 2008 system with 128 GB of ram, 80 TB of disk storage and two Xeon E5-2670 CPUs which runs Microsoft SQL Server 2012. The main table for the largest simulation contains 65,996,151 rows, which corresponds to ≈ 300 GB of disk space.

Columns are indexed on disk as follows (see below for the description of the content of each table):

1. The **SubHalo** and **Sizes** tables have a clustered index on the `GalaxyID`. This allows joins between the tables and queries for progenitors and descendants to run efficiently. `GalaxyID` rows are assigned such that progenitors of each galaxy have a continuous range of `GalaxyID`.
2. The **SubHalo** tables have an additional index on (`Snapnum`, `GalaxyID`) due to the common nature of queries that request a particular time in the simulation.
3. The **Aperture** tables have a clustered index on the combination of (`GalaxyID`, `ApertureSize`) and (`ApertureSize`, `GalaxyID`) to aid queries searching for all information about a single galaxy or one aperture size for many galaxies respectively.
4. The **FOF** tables are clustered on the combination (`SnapNum`, `GroupID`), which uniquely identifies the FOF group and can be used to join to the **SubHalo** table.

Typical queries (such as the ones given as examples in Section 3) take a few milliseconds to complete on the server. More complex queries (i.e. joining multiple tables or navigating the merger trees for multiple galaxies at the same time) can take up to a few seconds. As the usage goes up, additional indexing of the columns could be added to improve the performance of common, more complex, queries.

The mock *gri* images have been processed once for the entire simulation and are stored on a separate server. When querying images, the SQL server generates valid HTML tags containing the links to the images. No caching has been put in place but such facility could easily be added in case of large demand.

3. Use of the database

This section provides an overview of the database interface and of the different tables available for each simulation.

³ <https://www.mongodb.org/>.

⁴ <http://www.paradigm4.com/>.

⁵ <http://hadoop.apache.org/>.

Simple examples of how to query and combine the tables are presented.

3.1. Database interface

The main interface to the EAGLE database is shown in Fig. 3. Users familiar with the *Millennium* database (Lemson and Virgo Consortium, 2006) and its clones will recognise the main features of the interface and should be able to adapt their scripts easily to the EAGLE database.

SQL queries can be typed in the main text box (number 1 in the Figure) and are submitted to the database by pressing either of the buttons to the right (number 2). Some help with SQL queries can be obtained by clicking on the corresponding button. The results of queries submitted to the *browser* are returned at the bottom of the page in the form of an HTML table⁶ (number 7). This allows users to submit small queries and quickly verify the syntax. If images are being queried, they will appear directly in the results table. Larger, more complex queries should be submitted to the *stream* and will be returned in Comma-Separated-Value (CSV) format in a new window. The number of rows returned by the *browser* queries can be specified via the drop-down menu (number 3). The *stream* queries always return all rows. Previous queries can be recovered using the drop-down menu (number 4).

The queries from this paper are available as examples (number 5). These can later be adapted to match the user's need. All the available simulations and their tables are listed in the left-hand panel (number 6) with links to the documentation describing each entry in the table. All registered users receive a private database (MyDB) in which they can store query results for further processing at a later date. A link to MyDB is provided (number 11). Examples of how to create and manage such private tables can be obtained by clicking on the buttons at the bottom of the screen (number 8). Finally, some documentation, a list of credits are given at the top of the page (numbers 9 & 10).

3.2. Galaxy merger-tree traversal

In order to simplify the navigation of the trees, the database is stored with depth-first ordering (see, Lemson and Springel, 2006, Qu et al., in prep.). The progenitors of a galaxy can then easily be identified. To allow simple traversing of the merger tree of a given galaxy (with its unique GalaxyID), three additional columns are assigned to each galaxy:

1. **TopLeafID:** This is the GalaxyID of the highest-redshift main branch progenitor.
All the galaxies on the *main progenitor branch* of a galaxy with GalaxyID i and TopLeafID j have a GalaxyID in the range $[i, j]$ in ascending redshift order.
2. **LastProgID:** This is the maximum GalaxyID of all progenitors irrespective of their branch.
All the galaxies on *any progenitor branch* of a galaxy with GalaxyID i and LastProgID k have a GalaxyID in the range $[i, k]$.
3. **DescendantID:** This is the GalaxyID of the unique descendant galaxy of i .
If no descendant galaxy is identified then the DescendantID of a galaxy is set to its own GalaxyID.

⁶ Note that the *browser* queries time out after 90 s. More substantial queries should be submitted via the *stream* queries option. These only time out after 30 min.

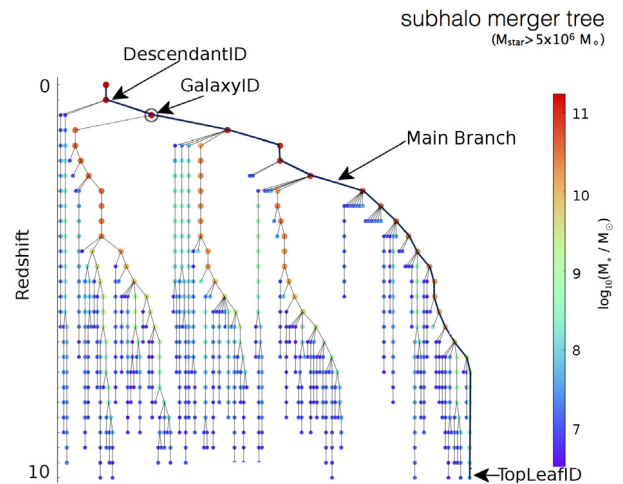


Fig. 2. Merger history of a galaxy with a $z = 0.18$ stellar mass $M_* \sim 10^{10} M_\odot$ indicated by the circled dot. Symbol colours and sizes are scaled with the logarithm of the stellar mass. The GalaxyID of this galaxy points towards it, as indicated by the arrow. The main progenitor branch is indicated with a thick black line, all other branches with a thin line. The TopLeaf ID gives the GalaxyID of the highest redshift galaxy on the main progenitor branch whilst the LastProgID (not shown) gives the maximum GalaxyID of all the progenitors of the galaxy considered. Querying all galaxies with an ID between GalaxyID and LastProgID will return all the progenitor galaxies in the tree. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2

SQL tables available for each simulation. The tables are prefixed with the name of the simulation to which they correspond. For example, the table of magnitudes for the 50 Mpc Ref-model is labelled **Ref10050N0752_Magnitudes** as can be seen on Fig. 3.

SQL table name	Contents
SubHalo	Main galaxy properties
FOF	Halo properties
Sizes	Galaxy sizes
Aperture	Galaxy properties in 3D apertures
Magnitudes	Galaxy photometry in the GAMA bands

In Fig. 2 we show a merger tree for a typical galaxy, indicated by its GalaxyID. The main branch is shown using a thicker blue line and the IDs required to navigate the tree are shown with arrows pointing towards the galaxy to which they correspond in the tree.⁷

Examples using the SQL language showing how to traverse the tree forwards and backwards in time are provided in Appendix A.

3.3. Content of the database

The EAGLE database for each simulation has information distributed across five SQL ‘tables’ listed in Table 2, whose contents are detailed in Appendix B.

SubHalo: This is the main table containing properties of galaxies, for example masses (of dark matter, gas, stars, and black holes), star formation rates, metallicities and angular momentum. The GalaxyID of a galaxy can be used to navigate through

⁷ Users familiar with the *Millennium* database can modify the queries by replacing HaloID with GalaxyID, mainLeafID or endMainBranchID with TopLeafID and lastProgenitorID with LastProgID. Note also that in the *Millennium* database, a galaxy with no descendant has its DescendantID set to -1 and not to GalaxyID as in the EAGLE database.

The screenshot shows the EAGLE Database web interface. At the top is a banner with the text "EAGLE Database" over a colorful galaxy background. Below the banner, the interface is divided into several sections:

- Documentation (9)**: A link to documentation.
- CREDITS/Acknowledgments (10)**: A link to credits and acknowledgments.
- Public Databases**: A section containing a list of available simulations and tables. A red box highlights this list, labeled "6) Available Simulations".
- Private (MyDB) Databases (11)**: A section for private databases, including logos for ICC, VIRG, and GAVO.
- Query Area (1)**: A large text input field for entering SQL queries. A red box highlights this area, labeled "1) Query area".
- Execute Query (2)**: Two buttons, "Query (stream)" and "Query (browser)", with a "Help" button below them. A red box highlights these buttons, labeled "2) Execute query".
- Query Results (7)**: A table showing query results in CSV format. A red box highlights this table, labeled "7) Query results (browser)".
- Previous queries (4)**: A section showing a list of previously executed queries with a "Show All" button. A red box highlights this section, labeled "4)".
- Demo queries (5)**: A section with buttons for "Subhalo", "Sizes", and "Magnitudes", each with sub-options like "SUB1", "SUB2", "SIZE1", "MAG1". A red box highlights this section, labeled "5) Demo queries".
- Private Database Examples (8)**: A section with buttons for "ShowTables", "Show Views", "Show Columns", "MyDB Size", "MyDB Table Size", "Create View", "Drop Table", and "Create Index". A red box highlights this section, labeled "8) Private Database Examples".
- Query time and rows (3)**: A section showing "Query time (in millisecond) = 2" and "Number of rows retrieved from database = 1000". A red box highlights this section, labeled "3)".

Fig. 3. The interface of the EAGLE database. SQL queries should be entered into the query area (1) and can be executed either via the 'browser' or 'stream' buttons (2). The *browser query* returns a limited number of results (3) at the bottom of the page (7), pressing the *Reformat* button will then return the full results in the selected format (default CSV) and *Plot(VOPlot)* is a simple way to visualise the data. This is the easiest method to test SQL scripts. The *stream query* returns all the results in a CSV format in a separate window to ease their download to a local device. Previous queries can be restored from the drop-down menu (4). The example query buttons (5) insert example SQL queries into the query area to help new users with the syntax and structure of the database. Similarly, examples creating and managing a private database are generated by clicking on the buttons (8). The list of available simulations and tables is given on the left hand side (6) with links to further documentation describing their contents. Users' own database tables are listed below (11). Further step-by-step documentation on how to use the web interface is provided (9) as well as links pointing to credits and acknowledgements (10).

its descendants and progenitors as well as to join the galaxy property table to other tables containing additional properties. The examples below demonstrate how to do this.

A full description of the contents of the **SubHalo** table is given in Table B.1.

Aperture: This table contains masses, star formation rates and velocity dispersions measured in a range of spherical apertures. Table B.4 gives a full list of the fields present in that SQL table. This table can be joined to the **SubHalo** table via the GalaxyID of the objects.

Magnitudes: This table contains non-dust-attenuated rest-frame broad-band magnitudes in the SDSS *ugriz* filters (Doi et al., 2010) and in the UKIRT *YJHK* filters (Hewett et al., 2006), computed in 30 pkpc spherical apertures for all galaxies with stellar mass greater than $10^{8.5} M_{\odot}$ as described in Trayford et al. (2015). See Table B.5 in the appendix for more details. This table can be joined to the **SubHalo** table via the GalaxyID of the objects.

Sizes: This table contains half-mass sizes of galaxies computed starting from apertures, as presented in Furlong et al., (in prep.).

See [Table B.3](#) in the appendix for a full list of available quantities. This table can be joined to the **SubHalo** table via the `GalaxyID` of the objects.

FOF: This table contains properties of *haloes*, for example mass and spherical overdensity radii. A full description of the contents of the FOF group table, including the units and dimensions of each variable, is given in [Table B.2](#). This table can be joined to the **SubHalo** table via the `GroupID` of the galaxies, given in the **SubHalo** table.

The **FOF** and **SubHalo** tables also contain a field with random number uniformly distributed in the range $[0, 1)$ allowing the users to generate unbiased sub-samples of galaxies or haloes.

3.4. Querying the database tables

In this section we will illustrate the use of the database by presenting simple example queries showing the basic usage of the different `SQL` tables.

The queries can be typed directly into the web interface or used in a `PYTHON` script, as described in [Appendix A](#), or using the `UNIX wget` command as described in the online documentation. The first example illustrates how to query the main galaxy table (**SubHalo**) in order to plot the relation between r_{\max} and v_{\max} at $z = 0$ (`Snapnum = 28`) for the `Ref-L0100N1504` simulation. In the database nomenclature, these quantities are `VmaxRadius` and `Vmax` (see [Table B.1](#)).

The `SQL` command to be typed in the input window is

Listing 1: Generate r_{\max} - v_{\max} table at $z = 0$

```
SELECT
  VmaxRadius as r_max, -- The two variables we
  Vmax as v_max        -- want to extract
FROM
  RefL0100N1504_SubHalo -- The simulation
WHERE
  SnapNum = 28          -- The snapshot
```

Clicking on the “Query (stream)” will open a new window containing the resulting two-column table with headers “`r_max`” and “`v_max`” in `CSV` format.

For many applications, multiple `SQL` tables have to be queried at the same time. The properties of a galaxy can be retrieved across the tables by joining their `GalaxyID`. A rest-frame colour–magnitude diagram using the `SDSS g` and `r` bands at $z = 0.1$ (`SnapNum = 27`) for central galaxies (`SubGroupNumber = 0`) with a stellar mass larger than $10^9 M_{\odot}$ (`Mass_Star > 1.0e9`) in a 30 pkpc aperture (`ApertureSize = 30`) can be constructed by joining the **SubHalo** table to the **Magnitudes** and **Aperture** ones.

This query reads

Listing 2: Generate table of $g - r$ vs. r colour-magnitude table for central galaxies with $M_* > 10^9$ at $z = 0.1$

```
-- Select the quantities we want
SELECT
  (MAG.g_nodust - MAG.r_nodust) as g_minus_r,
  MAG.r_nodust as r
-- Define aliases for the three tables
FROM
  RefL0100N1504_SubHalo as SH,
  RefL0100N1504_Magnitudes as MAG,
  RefL0100N1504_Aperture as AP
-- Apply the conditions
WHERE
  SH.SnapNum = 27 and          -- z=0.1
  SH.SubGroupNumber = 0 and   -- Centrals only
  AP.Mass_Star > 1.0e9 and    -- Mass limit
```

```
AP.ApertureSize = 30 and -- Aperture size
-- Join the objects in the 3 tables
SH.GalaxyID = MAG.GalaxyID and
SH.GalaxyID = AP.GalaxyID
```

and will return a two-column table with “`g_minus_r`” and “`r`” as headers containing the colours and r -band magnitudes of the selected galaxies.

Note that, as discussed in [Section 4](#), we recommend to always use quantities measured in apertures to avoid incorporating intra-cluster light into mass or star formation rate estimates.

Another common use of the database is to track one galaxy across time. To this end, one can navigate through the main progenitor branch. This final example tracks an interesting object (`GalaxyID = 1848116`) discovered at redshift $z = 1$ through time and constructs the stellar metallicity evolution accompanied by the mock *gri* face-on images of the object at all redshifts. One hence has to construct a query that returns all of the descendants (on the main branch) of the object by finding all galaxies that have the interesting object’s `GalaxyID` between their own `GalaxyID` and their `TopLeafID`. To get the progenitors one additionally requests all galaxies with `GalaxyID` between the interesting object’s `GalaxyID` and its `TopLeafID`. This demonstrates the merger tree navigation introduced in [Section 2.6](#). Note that adding conditions on the snapshot number (`SnapNum`) helps speed up the queries dramatically. This query reads

Listing 3: Returns the evolution along the main branch of stellar metallicity with redshift of a given galaxy with its images. To return the evolution along all branches replace `TopLeafID` with `LastProjID` in line 20.

```
-- Select the quantities we want
SELECT
  SH.Redshift as z,
  SH.Stars_Metallicity as Z,
  SH.Image_Face as face
-- Define two aliases for the main table
FROM
  -- Properties we want to extract
  RefL0025N0752_Subhalo as SH,
  -- Acts as a reference point
  RefL0025N0752_Subhalo as REF
-- Apply the conditions
WHERE
  REF.GalaxyID=1848116 and -- GalaxyID at z=1
  -- To find descendants
  ((SH.SnapNum > REF.SnapNum and REF.GalaxyID
  between SH.GalaxyID and SH.TopLeafID) or
  -- To find progenitors
  (SH.SnapNum <= REF.SnapNum and SH.GalaxyID
  between REF.GalaxyID and REF.TopLeafID))
-- Order the output by redshift
ORDER BY
  SH.Redshift
```

and will return a sorted table with a redshift and a metallicity column as well as a column containing the postage-stamp images of the galaxy at each redshift when using the “Query (browser)” button. These examples along with the more complex queries given in [Appendix A](#) are listed on the webpage documentation.

4. Recommendations, caveats and credits

4.1. Caveats regarding the usage of the data

In this section we list a series of recommendations and known limitations that the authors have uncovered while working on the analysis of the simulation and the preparation of the database. These points should be taken into consideration to exploit the

simulation outputs fully and to avoid mistakes in the interpretation of the results.

Finite resolution. When using the galaxy catalogues, it should be remembered that the properties of low-mass galaxies should be treated with caution. Large numbers of particles are required to adequately sample the formation history of a galaxy. In general, we find that many galaxy properties are unreliable below a stellar mass of $10^9 M_{\odot}$ for the intermediate resolution simulations (Schaye et al., 2015). For any given quantity, these effects can be assessed by comparing the Ref-L0025N0376 simulation with the higher-resolution Recal-L0025N0752 and Ref-L0025N0752 simulations.

Finite volume. Although the main simulation is one of the largest of its kind, its volume is still only 10^{-3} cGpc³, a volume much smaller than the volumes typically probed by surveys of the extragalactic Universe. This implies that rare objects are unlikely to be found in the simulation volume. Moreover, due to missing large-scale modes, the number density of rare objects will typically be underestimated. Only a handful of haloes with mass M_{200} (Group_M_Crit200 in the FOF table) above $10^{14} M_{\odot}$ are present in the main simulation, limiting the analysis of cluster-like objects. The convergence with box size can be assessed by comparing the main simulation to the smaller volumes that use the same resolution.

Aperture masses and SFRs. The stripping of satellite galaxies as they orbit within a halo generates a significant mass loss at large radii. The resulting diffuse light (and any diffuse star formation) is extremely difficult to observe and is not commonly included in observational galaxy catalogues. Furthermore, the total galaxy stellar masses and star formation rates can depend strongly on the precise assignment of particles to the main subhalo within each FOF group by the SUBFIND algorithm, which can lead to spurious total mass evolution. For these reasons, studies published by the EAGLE team use aperture masses and star formation rates, typically in an aperture of 30 pkpc. As discussed by Schaye et al. (2015), this corresponds roughly to an R_{80} Petrosian aperture and is hence particularly well-suited to comparison with observations. We recommend the use of aperture values when available.

Self-bound star clusters and black holes. As discussed by Schaye et al. (2015), small dense stellar regions within galaxies may occasionally be identified by SUBFIND as distinct subhaloes and hence ‘galaxies’. These appear in the catalogue as rather unusual objects with little stellar mass but anomalously high metallicity or black hole mass. These “spurious” galaxies are flagged in the database in the column Spurious (see the table in Appendix B). Such objects should not be considered as genuine galaxies and should be discarded from samples of simulated galaxies.

Black hole masses and accretion rates. The black hole masses given in the main table (table SubHalo, column BlackHoleMass) do not directly correspond to the mass of the central supermassive black hole of a galaxy, but to a summed value of all black holes assigned to that subhalo. For cases where BlackHoleMass > $10^6 M_{\odot}$ this closely approximates the mass of the most massive black hole. MassType_BH refers to the sum of the black hole particle masses (see Appendix D for details of particle and subgrid masses) and therefore should not be used for a galaxy’s black hole mass. Similarly, due to the coarse time sampling of the outputs, the high temporal variability of the black hole accretion rates cannot be captured in the database outputs and as such the quantity BlackHoleMassAccretionRate should be treated with great care.

Stellar velocity dispersion and morphology. The field StellarVelocityDispersion stored in the SubHalo table is a measure of the kinetic energy of the stars, $\sigma = \sqrt{3E_k/2M}$, and not a measure of the amount of stellar kinetic energy in dispersion as opposed to rotation. In particular, it cannot be used to distinguish rotationally supported galaxies (spirals) from dispersion supported galaxies (ellipticals).

Galaxy images and magnitude tables. The images provided in the database are generated using only the particles within a particular subhalo, in order to correspond with an entry in the database tables. As a result satellites or merging partners may not be visible in the images. While the images are observed as if redshifted to $z = 0.1$ to approximate typical SDSS colours, the magnitude tables are measured in the rest-frame. The inclusion of different population synthesis models, dust absorption and the relative scaling of images also implies that images are not reducible to magnitude table entries.

This simulation is not the real Universe. The papers presenting EAGLE have shown that the simulation broadly reproduces a wide set of observational properties of galaxies and the intergalactic medium. When using the database it should nevertheless be remembered that there are known discrepancies between the simulation results and observational data. In particular, we highlight the following points:

- Although the $z = 0.1$ stellar mass function was used in the calibration of the simulation, the stellar mass density is approximately 20% lower than inferred from observations (Schaye et al., 2015; Furlong et al., 2015). This missing mass can be related to the slight undershoot of the “knee” of the simulated galaxy stellar mass function.
- The evolution of specific star formation rates broadly follows the trends seen in observational data, but with a normalisation lower by, depending on redshift, 0.3–0.5 dex (Schaye et al., 2015; Furlong et al., 2015). Note, however, that the EAGLE galaxies are in good agreement with the recent recalibration of star formation indicators by Chang et al. (2015) (see Fig. 5 of Schaller et al., 2015a).
- The present-day stellar mass–metallicity relation in the intermediate-resolution Ref-model is flatter than the one inferred from observational data (Schaye et al., 2015). Note, however, that the relation becomes steeper in the higher-resolution Recal-L0025N0752 simulation, in agreement with the observations.
- The transition from active to passive galaxies occurs at too high a stellar mass at $z = 0$ (Schaye et al., 2015; Trayford et al., 2015).

This list of flaws is certainly not exhaustive. Future papers will undoubtedly uncover further deficiencies.

4.2. Acknowledgement of usage

To recognise the effort of the individuals involved in the design and execution of these simulations, in their post-processing and in the construction of the database, we kindly request the following:

- Publications making use of the EAGLE data extracted from the public database are kindly requested to cite the original papers introducing the project (Schaye et al., 2015; Crain et al., 2015) as well as this paper.
- Publications making use of the database should add the following line in their acknowledgement section: “We acknowledge the Virgo Consortium for making their simulation data available.”

The EAGLE simulations were performed using the DiRAC-2 facility at Durham, managed by the ICC, and the PRACE facility Curie based in France at TGCC, CEA, Bruyères-le-Châtel⁸.

- Furthermore, publications referring to specific aspects of the subgrid models, hydrodynamics solver, or post-processing steps (such as the construction of images or photometric quantities, and the construction of merger trees), are kindly requested to not only cite the above papers, but also the original papers describing these aspects. The appropriate references can be found in section 2 of this paper and in Schaye et al. (2015).

5. Conclusions

This paper introduces a public SQL relational database⁸ containing the integrated quantities and merger histories for more than 10^5 galaxies from the EAGLE suite of hydrodynamic simulations. The database contains all the galaxies from the largest EAGLE simulation as well as galaxies from smaller volumes where the resolution and AGN model were varied. The details of these simulations are presented by Schaye et al. (2015) and a list of published results using the simulation can be found on our websites.⁹

For each galaxy in the database and at each redshift, we provide a wide range of basic halo and galaxy properties such as stellar masses, gas masses, unextincted magnitudes, angular momenta, star formation rates and *gri* images, as well as extensive information on metal abundances. Three additional tables give the properties of galaxies measured in a series of apertures, more physically motivated galaxy sizes and galaxy photometry. Using their merger trees, galaxies can be tracked through time and their assembly history explored by analysing their progenitors.

By making the halo and galaxy data public we hope that our simulations will be helpful both for comparison with observational data, and as a tool for gaining physical insight into the physics of galaxy formation.

In Section 4 we presented some limitations of the simulations that should be borne in mind when using the database. In particular, caution should be exercised because of the finite resolution of the simulations. Over time we intend to make additional data products available as the relevant papers are accepted for publication. These will include, among other quantities, photometry including dust extinction and information on the morphology of the galaxies. At later stages, we may also release merger trees with higher time resolution, more simulations models from Crain et al. (2015) as well as the raw particle data.

The EAGLE database will hopefully be a powerful resource for the community to explore the physics of galaxy formation, and to help interpret observational data.

Acknowledgements

This work would have not been possible without Lydia Heck and Peter Draper’s technical support and expertise. We are grateful to all members of the Virgo Consortium and the EAGLE collaboration who have contributed to the development of the codes and simulations used here, as well as to the people who helped with the analysis. We thank Jaime Salcido for his help producing




Fig. A.1. Figure created by the PYTHON script.

Fig. 3, Violeta Gonzalez-Perez, Qi Guo and Claudia Lagos for useful comments on early drafts as well as Chris Barber, Bart Clauwens and Sean McGee for testing earlier versions of the EAGLE database.

This work was supported by the Science and Technology Facilities Council (grant number ST/F001166/1); European Research Council (grant numbers GA 267291 “Cosmiway” and GA 278594 “GasAroundGalaxies”) and by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office (AP P7/08 CHARM). RAC is a Royal Society University Research Fellow.

This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grant ST/H008519/1, and STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure. We acknowledge PRACE for awarding us access to the Curie machine based in France at TGCC, CEA, Bruyères-le-Châtel.

The web site described in this paper was based on the one build for the *Millennium Simulation* as part of the activities of the German Astrophysical Virtual Observatory (GAVO).

Appendix A. Examples of more complex queries

Python—Galaxy stellar mass function. This example¹⁰ replicates Fig. 4 from Schaye et al. (2015) comparing the galaxy stellar mass function at $z = 0.1$ (`SnapNum = 27`) in 30 pkpc apertures for three of the EAGLE simulations. The link to the database is created with the module `EAGLESQTOOLS` available from the release website¹¹ (this module serves as an interface to access the EAGLE database). After the connection is established (on line 9), the module can submit queries to the database. Each of the chosen table properties (in this case we have only chosen the galaxy stellar masses) is returned in a dictionary that can be then manipulated like any other PYTHON dictionary. We use the `GROUP BY` SQL keyword to bin the data directly on the server and reduce the amount of data being downloaded. The output created by this script is shown in Fig. A.1.

⁸ Available at the address <http://www.eaglesim.org/database.php>.

⁹ <http://eagle.strw.leidenuniv.nl/> and <http://www.eaglesim.org>.

¹⁰ Which can also be downloaded here: <http://icc.dur.ac.uk/Eagle/Database/GSMF.py>.

¹¹ Or directly here: <http://icc.dur.ac.uk/Eagle/Database/eagleSqlTools.py>.

```

import eagleSqlTools as sql
import numpy as np
import matplotlib.pyplot as plt

# Array of chosen simulations. Entries refer to the simulation name and comoving box length.
mySims = np.array([('RefL0100N1504', 100.), ('AGNdT9L0050N0752', 50.), ('Recall0025N0752', 25.)])

# This uses the eagleSqlTools module to connect to the database with your username and password.
# If the password is not given, the module will prompt for it.
con = sql.connect("<username>", password="<password>")

for sim_name, sim_size in mySims:

    print sim_name

    # Construct and execute query for each simulation. This query returns the number of galaxies
    # for a given 30 pkpc aperture stellar mass bin (centered with 0.2 dex width).
    myQuery = "SELECT \
                0.1+floor(log10(AP.Mass_Star)/0.2)*0.2 as mass, \
                count(*) as num \
            FROM \
                %s_SubHalo as SH, \
                %s_Aperture as AP \
            WHERE \
                SH.GalaxyID = AP.GalaxyID and \
                AP.ApertureSize = 30 and \
                AP.Mass_Star > 1e8 and \
                SH.SnapNum = 27 \
            GROUP BY \
                0.1+floor(log10(AP.Mass_Star)/0.2)*0.2 \
            ORDER BY \
                mass"%(sim_name, sim_name)

    # Execute query.
    myData = sql.execute_query(con, myQuery)

    # Normalize by volume and bin width.
    hist = myData['num'][:] / float(sim_size)**3.
    hist = hist / 0.2

    plt.plot(myData['mass'], np.log10(hist), label=sim_name, linewidth=2)

# Label plot.
plt.xlabel(r'log$_{10}$ M$_{*}$ [M$_{\odot}$]', fontsize=20)
plt.ylabel(r'log$_{10}$ dn/dlog$_{10}$ (M$_{*}$) [cMpc$^{-3}$]', fontsize=20)
plt.tight_layout()
plt.legend()

plt.savefig('GSMF.png')
plt.close()

```

SQL—Black hole mass vs. stellar mass. This example replicates Fig. 10 from [Schaye et al. \(2015\)](#) showing the black hole mass as a function of stellar mass at redshift $z = 0.1$ ($\text{SnapNum} = 27$) for the reference (L0100N1504) run. As mentioned in Section 4 we use the black hole subgrid mass, `BlackHoleMass`, and treat the stellar mass of a galaxy as the mass contained within a 30 pkpc aperture. `SubHalo` table properties are connected to the `Aperture` table via each galaxy's unique `GalaxyID`.

```

-- Select the quantities we want
SELECT
    AP_Star.Mass_Star as sm,
    SH.BlackHoleMass as bhm
-- Define aliases for the two tables
FROM
    RefL0100N1504_Subhalo as SH,
    RefL0100N1504_Aperture as AP_Star
-- Apply the conditions
WHERE
    SH.SnapNum = 27                -- z=0.1
    and SH.GalaxyID = AP_Star.GalaxyID -- Match galaxies to Aperture table
    and AP_Star.ApertureSize = 30      -- Select aperture size to be 30 pkpc
    and AP_Star.Mass_Star > 0         -- Only return stellar masses > 0
    and SH.BlackHoleMass > 0         -- Only return black hole masses > 0

```

SQL—Galaxy size vs. stellar mass. This example is similar to Fig. 9 from [Schaye et al. \(2015\)](#) comparing galaxy size as a function of stellar mass at redshift $z = 0.1$ ($\text{SnapNum} = 27$) for each galaxy in the reference (L0100N1504) run. For galaxy sizes, we use the half mass radius of the galaxies from the `Sizes` table. We connect them to the galaxy's stellar mass via the unique `GalaxyID` identifier. As with the previous example, we must connect to the `SubHalo` table to retrieve the `SnapNum` value.

```

-- Select the quantities we want
SELECT
    AP.Mass_Star as sm,
    SIZES.R_halfmass100 as size

```

```
-- Define aliases for the three tables
FROM
    RefL0100N1504_Subhalo as SH,
    RefL0100N1504_Aperture as AP,
    RefL0100N1504_Sizes as SIZES
-- Apply the conditions
WHERE
    SH.SnapNum = 27 -- z=0.1
    and SH.GalaxyID = AP.GalaxyID -- Match galaxies to Aperture table
    and SH.GalaxyID = SIZES.GalaxyID -- Match galaxies to Sizes table
    and AP.ApertureSize = 30 -- Select aperture size to be 30 pkpc
    and AP.Mass_Star > 0 -- Only return stellar masses > 0
```

SQL—Joining FOF and SubHalo tables. This example shows how to join the properties of galaxies to their parent FoF halo. In this case, we compute the offset between the centre of the potential of the galaxy and the FoF halo. When dealing with positions within these volumes, remember to account for box periodicity. In principle, it is not necessary to match the SnapNum of the tables as well as the GroupID, but this speeds up the query.

```
-- Select the quantities we want
SELECT
    SH.CentreOfPotential_x as sh_x,
    SH.CentreOfPotential_y as sh_y,
    SH.CentreOfPotential_z as sh_z,
    FOF.GroupCentreOfPotential_x as fof_x,
    FOF.GroupCentreOfPotential_y as fof_y,
    FOF.GroupCentreOfPotential_z as fof_z,
    SH.MassType_Star as mstar,
    FOF.GroupMass as fof_mass,
    square(SH.CentreOfPotential_x-FOF.GroupCentreOfPotential_x)
    + square(SH.CentreOfPotential_y-FOF.GroupCentreOfPotential_y)
    + square(SH.CentreOfPotential_z-FOF.GroupCentreOfPotential_z) as dist
-- Define aliases for the two tables
FROM
    RefL0050N0752_Subhalo as SH,
    RefL0050N0752_FOF as FOF
-- Apply the conditions
WHERE
    SH.MassType_Star > 1.0e11 -- Only return stellar masses > 1.0e11
    and SH.SnapNum = 27 -- z=0.1
    and FOF.SnapNum = SH.SnapNum -- Join SnapNum to speed up query
    and FOF.GroupID = SH.GroupID -- Join GroupID to speed up query
```

SQL—Linking a progenitor to its descendants. This example shows how to select a random subset of Milky Way like galaxies, and extract information about the location and specific star formation rates (within a 30 pkpc aperture) of all their progenitors above a stellar mass of $10^9 M_{\odot}$.

```
-- Select the quantities we want
SELECT
    DES.GalaxyID,
    PROG.Redshift,
    PROG.MassType_DM,
    PROG.MassType_Star,
    AP.SFR / (AP.Mass_Star+0.0001) as ssfr,
    PROG.CentreOfPotential_x,
    PROG.CentreOfPotential_y,
    PROG.CentreOfPotential_z
-- Define aliases for the three tables
FROM
    RefL0100N1504_Subhalo as PROG,
    RefL0100N1504_Subhalo as DES,
    RefL0100N1504_Aperture as AP
-- Apply the conditions
WHERE
    DES.MassType_Star between 1.0e10 and 6e10 -- Select Milky Way like stellar mass
    and DES.MassType_DM between 5.0e11 and 2.0e12 -- Select Milky Way like halo mass
    and DES.RandomNumber < 0.1 -- Take a random subset of these
    and DES.SnapNum = 28 -- At redshift z=0
    and PROG.GalaxyID between DES.GalaxyID and DES.LastProgID -- Then find galaxy progenitors
    and AP.ApertureSize = 30 -- Select aperture size to be 30 pkpc
    and AP.GalaxyID = DES.GalaxyID -- Match galaxies to Aperture table
    and AP.Mass_Star > 1.0e9 -- Only return galaxies with stellar mass > 1e9
-- Order the output
ORDER BY
    DES.MassType_Star desc,
    PROG.Redshift asc,
    PROG.MassType_Star desc
```

Appendix B. Description of all fields contained in the database

See [Table B.1](#).

Appendix C. List of snapshot output times

See [Table C.1](#).

Appendix D. Detailed expressions for quantities in the database

In order to remove any ambiguity in the quantities provided in the database, this appendix summarises the fundamental equations that are being solved and the co-ordinate system used in the numerical code.

The equations that describe the evolution of a gravitating fluid are the continuity, Euler, energy and Poisson equations ([Peebles, 1980](#)). In order to provide a precise definition of the symbols used to describe database entries, we write these equations as

$$\frac{\partial \rho}{\partial t} + (\mathbf{v} \cdot \nabla) \rho \equiv \frac{d\rho}{dt} = -\rho \nabla \cdot \mathbf{v} \quad (\text{D.1})$$

$$\frac{d\mathbf{v}}{dt} = -\frac{1}{\rho} \nabla p - \nabla \Phi \quad (\text{D.2})$$

$$\frac{du}{dt} = -\frac{p}{\rho} \nabla \cdot \mathbf{v} - \mathcal{C} \rho \quad (\text{D.3})$$

$$\nabla^2 \Phi = 4\pi G(\rho + \rho_{\text{col}}) - \Lambda, \quad (\text{D.4})$$

where ρ is gas density, ρ_{col} the density due to collisionless matter (i.e., stars, dark matter and black holes), p the effective gas pressure, Φ the (Newtonian) gravitational potential and Λ the cosmological constant. The variable

$$u = \frac{p}{(\gamma - 1)\rho} = \frac{k_{\text{B}}T}{(\gamma - 1)\mu m_{\text{H}}}, \quad (\text{D.5})$$

is the thermal energy per unit mass, with $\gamma = 5/3$ the ratio of specific heats for a mono-atomic gas, and μ the mean molecular weight in units of the Hydrogen mass, m_{H} . The term $\mathcal{C}(T)\rho$ describes radiative cooling and heating. The database also reports the value of the pseudo-entropy S , defined as

$$S \equiv \frac{p}{\rho^\gamma}. \quad (\text{D.6})$$

We use the standard notation for proper time, t , position, \mathbf{r} , and velocity $\mathbf{v} \equiv d\mathbf{r}/dt \equiv \dot{\mathbf{r}}$. Partial derivatives are defined so that $\partial/\partial t$ is the time derivative at constant position \mathbf{r} , $d/dt \equiv \partial/\partial t + (\mathbf{v} \cdot \nabla)$ is the Lagrangian time derivative, and the spatial derivative $\nabla \equiv \partial/\partial \mathbf{r}$ is computed at constant time.

We solve these equations in an expanding coordinates described by the scale factor $a(t)$ which satisfies the Friedmann equations,

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3} \bar{\rho}_t + \frac{\Lambda}{3} \quad (\text{D.7})$$

$$\ddot{a} = -\frac{4\pi G}{3} \bar{\rho}_t a + \frac{\Lambda}{3} a, \quad (\text{D.8})$$

where $\Lambda \equiv 3H_0^2 \Omega_\Lambda$ (with $H_0 \equiv H(a=1)$), and $\bar{\rho}_t$ is the mean total density, $\bar{\rho}_t = \bar{\rho} + \bar{\rho}_{\text{col}}$. We apply periodic boundary conditions in this expanding reference frame.

The simulation uses comoving coordinates to simplify the integration of Eqs. (D.1)–(D.4). These are defined as

$$\mathbf{x} \equiv \frac{\mathbf{r}}{a} \quad (\text{D.9})$$

$$\hat{\rho} \equiv a^3 \rho \quad (\text{D.10})$$

$$\hat{u} \equiv a^{-2} u \quad (\text{D.11})$$

$$\hat{\Phi} = a \left(\Phi - \frac{2\pi}{3} G \bar{\rho}_t r^2 + \frac{1}{6} \Lambda r^2 \right) \quad (\text{D.12})$$

$$\hat{\nabla} \equiv a \nabla \quad (\text{D.13})$$

$$\hat{p} = (\gamma - 1) \hat{\rho} \hat{u} \quad (\text{D.14})$$

$$\hat{S} = S. \quad (\text{D.15})$$

In these variables, the velocity

$$\mathbf{v} = \dot{a}\mathbf{x} + \mathbf{v}_p \quad (\text{D.16})$$

$$\mathbf{v}_p \equiv a\dot{\mathbf{x}}, \quad (\text{D.17})$$

where \mathbf{v}_p is referred to as the peculiar velocity. We will use the term ‘comoving variable’ when a quantity is expressed in comoving variables (i.e. \mathbf{x} , $\dot{\mathbf{x}}$ and hatted variables), and ‘physical’ or ‘proper’ otherwise. In particular we will express comoving distances in cMpc (for comoving mega parsecs) and physical distances in pMpc (for proper or physical mega parsecs), and similarly for ckpc and pkpc.

The equations are solved by representing the collisionless mass as well as the gas by particles. We denote particle masses as m_i , for particle i . In EAGLE, star particles lose mass to gas particles to represent mass loss from stars. Each star particle therefore has two mass variables, its current particle mass, m , and its birth mass \tilde{m} :

$$\begin{aligned} m &= \text{current particle mass of star} \\ \tilde{m} &= \text{birth mass of star}, \end{aligned} \quad (\text{D.18})$$

see [Wiersma et al. \(2009b\)](#) for more details. Black holes also have two mass variables associated with them, a particle mass m , and a subgrid mass \tilde{m} . It is the subgrid mass that enters the equations describing the accretion rate of a black hole. In short,

$$\begin{aligned} m &= \text{particle mass of black hole} \\ \tilde{m} &= \text{subgrid mass of black hole}. \end{aligned} \quad (\text{D.19})$$

Once a black hole is significantly more massive than the seed mass, particle and subgrid mass trace each other closely, see [Booth and Schaye \(2009\)](#) and [Rosas-Guevara et al. \(2015\)](#) for details.

Having defined comoving variables, the comoving energy \hat{E} of a collisionless halo is

$$\hat{E} = \frac{1}{2} \sum_i m_i (a^2 \dot{\mathbf{x}}_i)^2 + \frac{1}{2} \sum_i m_i \hat{\Phi}_i, \quad (\text{D.20})$$

and is conserved for an isolated halo, as is its comoving spin $\hat{\mathbf{L}}$,

$$\hat{\mathbf{L}} = \sum_i m_i (\mathbf{x} - \mathbf{x}_{\text{com}}) \times (a^2 \dot{\mathbf{x}}_i - a^2 \dot{\mathbf{x}}_{\text{com}}). \quad (\text{D.21})$$

Here

$$\mathbf{x}_{\text{com}} = \frac{\sum_i m_i \mathbf{x}_i}{\sum_i m_i}, \quad (\text{D.22})$$

is the comoving position of the centre of mass (taking into account periodic boundary conditions), and $\dot{\mathbf{x}}_{\text{com}}$ its time derivative.

The database uses comoving co-ordinates to record the locations of haloes and galaxies. For example, the position of

Table B.1

Full listing of the content of the main galaxy properties table and description of the columns. These properties are contained in tables denoted [modelname]_SubHalo. The first five lines of the table give the indices used to navigate between the side tables and through the merger trees. Particle types are dark matter, gas, stars and black holes and collective properties such as Mass sum over all of these particles unless otherwise stated. Columns in this table exist for each of three different components: star-forming gas (SF), non-star-forming gas (NSF) and stars (Stars). As these properties are repeated for each of these components, we only describe them once. In the database each property will be preceded with either [SF/NSF/Stars]_ before its name. For instance, the metallicity field will exist in three variants: SF_Metallicity, NSF_Metallicity and Stars_Metallicity for the metallicity of the star-forming gas, of the non star-forming gas and of the stars, respectively. Any sum used to describe a property is the sum of all particles for that component only.

SubHalo		
Field	Units	Description
GalaxyID	–	Unique identifier of a galaxy. This identifier enables linking the SubHalo table to the Aperture , Magnitudes and Sizes tables.
LastProgID	–	Used for merger tree traversal, Section 2.6.
TopLeafID	–	Used for merger tree traversal, Section 2.6.
DescendantID	–	GalaxyID of the descendant of this galaxy, Section 2.6.
GroupID	–	Unique identifier of the FoF halo hosting this galaxy. This identifier enables linking the SubHalo table to the FOF table.
Redshift	–	Redshift at which these properties are computed.
Snapnum	–	Snapshot number at which these properties are computed.
GroupNumber	–	Integer identifier of the FoF halo hosting this galaxy. GroupNumber is only unique to a given snapshot and hence cannot be used to identify the same halo across multiple snapshots.
SubGroupNumber	–	Integer identifier of this galaxy within its FoF halo. SubGroupNumber is only unique to a given FoF halo in a given snapshot and hence cannot be used to identify the same galaxy across multiple outputs. The condition “SubGroupNumber = 0” selects central galaxies.
Spurious	–	Value is 1 if the galaxy is an artefact of the SUBFIND algorithm and 0 if the galaxy is a genuine object, see Section 4.
Image_face	–	Weblink to the mock <i>gri</i> image of the galaxy in the three different orientations (face-on, edge-on and along the simulation <i>z</i> -axis). When querying the database via the browser, the image appears in the column of the results table.
Image_edge	–	
Image_box	–	
BlackHoleMass	M_{\odot}	Sum of all the black hole subgrid masses in this galaxy. See Eq. (D.19) for a description of the subgrid mass \tilde{m} of a black hole, and Section 4 for cautions using black hole masses.
BlackHoleMassAccretionRate	$M_{\odot} \text{ yr}^{-1}$	Total instantaneous accretion rate of all black holes, see Section 4 for cautions.
CentreOfMass_x	cMpc	Co-moving position of the centre of mass, Eq. (D.23).
CentreOfMass_y		
CentreOfMass_z		
CentreOfPotential_x	cMpc	Co-moving position of the minimum of the gravitational potential defined by the position of the most bound particle.
CentreOfPotential_y		
CentreOfPotential_z		
GasSpin_x	pkpc km s ⁻¹	Spin per unit mass of all gas particles, L/M , with L given by Eq. (D.21).
GasSpin_y		
GasSpin_z		
HalfMassRad_DM	pkpc	Physical radius enclosing half of the dark matter mass.
HalfMassRad_Gas	pkpc	Physical radius enclosing half of the gas mass.
HalfMassRad_Star	pkpc	Physical radius enclosing half of the stellar mass.
HalfMassRad_BH	pkpc	Physical radius enclosing half of the black hole particle mass, m , as defined in Eq. (D.19).
HalfMassProjRad_DM	pkpc	Projected physical radius enclosing half of the dark matter mass, averaged over three orthogonal projections.
HalfMassProjRad_Gas	pkpc	Projected physical radius enclosing half of the gas mass, averaged over three orthogonal projections.
HalfMassProjRad_Star	pkpc	Projected physical radius enclosing half of the stellar mass, averaged over three orthogonal projections.
HalfMassProjRad_BH	pkpc	Projected physical radius enclosing half of the black hole particle mass, m (Eq. (D.19)), averaged over three orthogonal projections.
InitialMassWeightedBirthZ	z	Mean redshift of formation of stars, weighted by birth mass \tilde{m} (Eq. (D.18)). Calculated via $\sum_i \tilde{m}_i \bar{z}_i / \sum_i \tilde{m}_i$ where \bar{z}_i is the redshift the star particle i was formed and \tilde{m}_i its birth mass.
InitialMassWeightedStellarAge	Gyr	Mean age of stars, weighted by birth mass. Calculated via $\sum_i \tilde{m}_i (t - \bar{t}_i) / \sum_i \tilde{m}_i$ where t is cosmic time, and \bar{t}_i and \tilde{m}_i formation time and birth mass of the star particle i , respectively.
KineticEnergy	$M_{\odot} (\text{km/s})^2$	Total kinetic energy E_k , Eq. (D.26).
Mass	M_{\odot}	Total current mass of all particles (i.e. $\sum_i m_i$ where m_i is the mass of the particle).
MassType_DM	M_{\odot}	Total dark matter mass.
MassType_Gas	M_{\odot}	Total gas mass.
MassType_Star	M_{\odot}	Total stellar mass, $\sum_i m_i$, where m_i is the stellar particle mass from Eq. (D.18).
MassType_BH	M_{\odot}	Total black hole mass, $\sum_i m_i$, where m_i is the black hole particle mass from Eq. (D.19).
RandomNumber	–	Random number uniform in the range [0, 1).
StarFormationRate	$M_{\odot} \text{ yr}^{-1}$	Total star formation rate, $\sum_i \dot{m}_{*,i}$, where $\dot{m}_{*,i}$ is the star formation rate of gas particle i .
StellarInitialMass	M_{\odot}	Sum of birth masses of all stars, $\sum_i \tilde{m}_i$, where \tilde{m}_i is the birth mass of star particle i from Eq. (D.18).
StellarVelDisp	km s ⁻¹	One dimensional velocity dispersion of stars, $((2 \text{Star_KineticEnergy}) / (3 \text{MassType_Star}))^{1/2}$, where $\text{Star_KineticEnergy}$ is the kinetic energy of stars.
ThermalEnergy	$M_{\odot} (\text{km/s})^2$	Total thermal energy E_{th} , Eq. (D.28).
TotalEnergy	$M_{\odot} (\text{km/s})^2$	Total energy E_{tot} , Eq. (D.29).
Velocity_x	km s ⁻¹	Peculiar velocity, Eq. (D.24).
Velocity_y		

(continued on next page)

Table B.1 (continued)

Field	Units	Description
SubHalo		
Velocity_z		
Vmax	km s ⁻¹	Maximum value of the circular velocity, $(G(M(<r)/r))^{1/2}$, where $M(<r)$ is the total mass enclosed in a sphere of physical radius r .
VmaxRadius	pkpc	Physical radius where the circular velocity equals Vmax.
Hydrogen	–	
Helium	–	
Carbon	–	
Nitrogen	–	
Oxygen	–	
Neon	–	
Magnesium	–	Total mass in this element divided by the total mass (both for a given component). These are therefore absolute abundances which do not depend on the solar abundance.
Silicon	–	
Sulphur	–	
Calcium	–	
Iron	–	
IronFromSNIA	–	Total mass in Iron contributed by ejecta from Type Ia supernovae, divided by the total mass.
KineticEnergy	M _⊙ (km/s) ²	Total kinetic energy E_K , Eq. (D.26).
Mass	M _⊙	Total mass, $\sum_i m_i$, where m_i is the particle mass.
MassFromAGB	M _⊙	Total mass contributed by ejecta of AGB stars.
MassFromSNII	M _⊙	Total mass contributed by ejecta from massive stars and type II supernovae.
MassFromSNIA	M _⊙	Total mass contributed by ejecta from type Ia SN supernovae.
MassWeightedEntropy	km s ² $(\frac{10^{10} M_{\odot}}{\text{Mpc}})^{2/3}$	Mass-weighted pseudo entropy of all particles, $\sum_i m_i S_i / \sum_i m_i$, where S_i is the pseudo-entropy (Eq. (D.6)) and m_i is the mass of the particle i . (Entry present for gas components only.)
MassWeightedTemperature	K	Mass-weighted temperature, $\sum_i m_i T_i / \sum_i m_i$, where T_i is the temperature and m_i is the mass of the particle i . (Entry present for gas components only.)
Metallicity	–	Metal mass fraction, $\sum_i m_i Z_i / \sum_i m_i$, where Z_i is the metallicity and m_i is the mass of the particle i .
MetalsFromAGB	M _⊙	Total metal mass contributed by ejecta from AGB stars.
MetalsFromSNII	M _⊙	Total metal mass contributed by ejecta from massive stars and SN Type II supernovae.
MetalsFromSNIA	M _⊙	Total metal mass contributed by ejecta from Type Ia supernovae.
Spin_x		
Spin_y	pkpc km s ⁻¹	Spin per unit mass, L/M , from Eq. (D.21).
Spin_z		
ThermalEnergy	M _⊙ (km/s) ²	Total thermal energy E_u , Eq. (D.28). (Entry present for gas components only.)
TotalEnergy	M _⊙ (km/s) ²	Total energy E_{tot} , Eq. (D.29). The potential energy contribution does include the other components as well, $E_{\phi} = \frac{1}{2} \sum_i m_i \frac{\hat{\phi}_i}{a}$.

Table B.2

Full listing of the content of the halo table and description of the columns. These properties are contained in tables denoted [**modelname**]_{FOF}. This table can be linked to the [**modelname**]_{SubHalo} table using the unique GroupID identifier.

Field	Units	Description
FOF		
GroupID	–	Unique identifier of a halo (i.e. a FoF group). This identifier enables linking a halo to all its galaxies and their properties in the SubHalo table.
Redshift	–	Redshift at which these properties are computed.
SnapNum	–	Snapshot number containing that halo.
GroupCentreOfPotential_x		
GroupCentreOfPotential_y	cMpc	Co-moving position of the minimum of the gravitational potential of the halo.
GroupCentreOfPotential_z		
GroupMass	M _⊙	Total Friends-of-Friends mass of this halo.
Group_M_Crit200		
Group_M_Crit500	M _⊙	Total mass within the corresponding Group_R_Critxxx radius, where xxx = 200, 500 or 2500, respectively
Group_M_Crit2500		
Group_M_Mean200		
Group_M_Mean500	M _⊙	Total mass within the corresponding Group_R_Meanxxx radius, where xxx = 200, 500 or 2500, respectively
Group_M_Mean2500		
Group_M_TopHat200	M _⊙	Total mass within radius Group_R_TopHat200.
Group_R_Crit200		
Group_R_Crit500	pkpc	Physical radius within which the mean density is xxx times the <i>critical</i> density of the Universe, where xxx = 200, 500 or 2500, respectively.
Group_R_Crit2500		
Group_R_Mean200		
Group_R_Mean500	pkpc	Physical radius within which the mean density is xxx times the <i>mean</i> density of the Universe, where xxx = 200, 500 or 2500, respectively.
Group_R_Mean2500		
Group_R_TopHat200	pkpc	Physical radius within which the mean density is $18\pi^2 + 82(\Omega_m(z) - 1) - 39(\Omega_m(z) - 1)^2$ times the critical density of the Universe. This is based on the spherical top-hat collapse model of Bryan and Norman (1998).
NumOfSubhalos	–	Number of subhaloes (galaxies) identified as belonging to this halo.
RandomNumber	–	Random number uniform in the range [0, 1).

Table B.3

Full listing of the content of the galaxy sizes table and description of the columns. These properties are contained in tables denoted **[modelname]_Sizes**. This table contains half-mass sizes of the *stellar* component of galaxies using spherical apertures (Furlong et al. in prep.). The `GalaxyID` column can be used to join this table to the corresponding **[modelname]_SubHalo** table. Only galaxies with total stellar mass $M_* > 10^8 M_\odot$ have entries in this table.

Sizes		
Field	Units	Description
GalaxyID	–	Unique identifier of a galaxy as per SubHalo table.
R_halfmass30	pkpc	Half mass radius of stellar component within a spherical (3D) 30 pkpc aperture.
R_halfmass100	pkpc	Half mass radius of stellar component within a spherical (3D) 100 pkpc aperture.
R_halfmass30_projected	pkpc	Projected half mass radius of stellar component within a circular (2D) 30 pkpc aperture (averaged over three orthogonal projections).
R_halfmass100_projected	pkpc	Projected half mass radius of stellar component within a circular (2D) 100 pkpc aperture (averaged over three orthogonal projections).

Table B.4

Full listing of the content of the aperture table and description of the columns. These properties are contained in tables denoted **[modelname]_Aperture**. This table contains measurements within spherical apertures centred on the minimum of the gravitational potential of a given galaxy. Each row represents a set of measurements for a single galaxy using a single aperture size in physical kpc. The `GalaxyID` column can be used to join this table to the corresponding **[modelname]_SubHalo** table.

Aperture		
Field	Units	Description
GalaxyID	–	Unique identifier of a galaxy as per SubHalo table.
ApertureSize	pkpc	Spherical (3D) aperture radius used for this measurement. Quantities are measured in a sphere centred at the centre of the potential, <i>i.e.</i> at the location of the most bound particle. Available aperture sizes are: 1, 3, 5, 10, 20, 30, 40, 50, 70 and 100 pkpc.
VelDisp	km s ⁻¹	One dimensional velocity dispersion of stars, $((2 \text{ KineticEnergy_Star}) / (3 \text{ Mass_Star}))^{1/2}$, where <code>KineticEnergy_Star</code> is the kinetic energy of stars, and the sum is over stars within the aperture.
SFR	M _⊙ yr ⁻¹	Star formation rate within the aperture.
Mass_BH	M _⊙	Total particle mass, $\sum_i m_i$ (Eq. (D.19)), of all black holes within the aperture.
Mass_DM	M _⊙	Total dark matter mass within the aperture.
Mass_Gas	M _⊙	Total gas mass within the aperture.
Mass_Star	M _⊙	Total stellar mass, $\sum_i m_i$ (Eq. (D.18)), within the aperture.

Table B.5

Full listing of the content of the magnitudes table and description of the columns. These properties are contained in tables denoted **[modelname]_Magnitudes**. This table contains absolute rest-frame magnitudes without dust attenuation for all galaxies with $M_* > 10^{8.5} M_\odot$ contained in the **SubHalo** table. This table can be joined to the **SubHalo** table using the `GalaxyID` field. The magnitudes in the different SDSS (Doi et al., 2010) and UKIRT (Hewett et al., 2006) filters have been computed in 30 pkpc spherical apertures following the procedure described by Trayford et al. (2015).

Magnitudes			
Field	Units	Description	
GalaxyID	–	Unique identifier of a galaxy as per SubHalo table.	
u_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>u</i> band without dust attenuation.	
g_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>g</i> band without dust attenuation.	
r_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>r</i> band without dust attenuation.	
i_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>i</i> band without dust attenuation.	
z_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>z</i> band without dust attenuation.	
Y_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>Y</i> band without dust attenuation.	
J_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>J</i> band without dust attenuation.	
H_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>H</i> band without dust attenuation.	
K_nodust	mag	Rest-frame absolute magnitude (AB) in the <i>K</i> band without dust attenuation.	

the centre of a galaxy or halo (stored as `CentreOfMass` in the database) is

$$\text{CentreOfMass} = \mathbf{x}_{\text{com}} = \frac{\sum_i m_i \mathbf{x}_i}{\sum_i m_i}, \quad (\text{D.23})$$

where the sum runs over all particles that belong to the object taking into account periodic boundary conditions. Similarly, the centre of potential of an object (database variable `CentreOfPotential`) is given in comoving coordinates.

The velocity of a halo or galaxy (database variable `Velocity`) refers to its peculiar velocity,

$$\text{Velocity} = a \dot{\mathbf{x}}_{\text{com}} = a \frac{\sum_i m_i \dot{\mathbf{x}}_i}{\sum_i m_i}. \quad (\text{D.24})$$

All other variables are expressed in physical coordinates, for example the spin vector of a galaxy is computed as

$$\text{Spin} = \frac{\sum_i m_i (\mathbf{r}_i - \mathbf{r}_{\text{com}}) \times (\mathbf{v}_i - \mathbf{v}_{\text{com}})}{\sum_i m_i}, \quad (\text{D.25})$$

where \mathbf{r}_{com} and \mathbf{v}_{com} are the physical position and velocity of the centre of mass. The expressions for physical kinetic, potential, thermal, and total energy are, respectively

$$E_K = \frac{1}{2} \sum_i m_i (\mathbf{v} - \mathbf{v}_{\text{com}})^2 \quad (\text{D.26})$$

$$E_\phi = \frac{1}{2} \sum_i m_i \frac{\hat{\phi}_i}{a} \quad (\text{D.27})$$

Table C.1

List of all output redshifts for all the simulations present in the database. Note that SQL queries are made easier by the use of the snapshot number rather than the redshift. Lookback times are given in Gigayears.

SnapNum	Redshift	Lookback time	Expansion factor
28	0.00	0.00	1.000
27	0.10	1.34	0.909
26	0.18	2.29	0.846
25	0.27	3.23	0.787
24	0.37	4.16	0.732
23	0.50	5.19	0.665
22	0.62	6.01	0.619
21	0.74	6.71	0.576
20	0.87	7.37	0.536
19	1.00	7.93	0.499
18	1.26	8.86	0.443
17	1.49	9.49	0.402
16	1.74	10.05	0.365
15	2.01	10.53	0.332
14	2.24	10.86	0.309
13	2.48	11.16	0.287
12	3.02	11.66	0.249
11	3.53	12.01	0.221
10	3.98	12.25	0.201
9	4.49	12.46	0.182
8	5.04	12.63	0.166
7	5.49	12.75	0.154
6	5.97	12.86	0.143
5	7.05	13.04	0.124
4	8.07	13.16	0.110
3	8.99	13.25	0.100
2	9.99	13.32	0.091
1	15.13	13.53	0.062
0	20.00	13.62	0.047

$$E_u = \sum_i m_i (a^2 \hat{u}_i) \quad (\text{D.28})$$

$$E_{\text{tot}} = E_K + E_\phi + E_u. \quad (\text{D.29})$$

References

- Baes, M., et al., 2003. Radiative transfer in disc galaxies—III. The observed kinematics of dusty disc galaxies. *Mon. Not. R. Astron. Soc.* 343, 1081–1094. doi:10.1046/j.1365-8711.2003.06770.x. arXiv:astro-ph/0304501.
- Booth, C.M., Schaye, J., 2009. Cosmological simulations of the growth of supermassive black holes and feedback from active galactic nuclei: method and tests. *Mon. Not. R. Astron. Soc.* 398, 53–74. doi:10.1111/j.1365-2966.2009.15043.x. arXiv:0904.2572.
- Bruzual, G., Charlot, S., 2003. Stellar population synthesis at the resolution of 2003. *Mon. Not. R. Astron. Soc.* 344, 1000–1028.
- Bryan, G.L., Norman, M.L., 1998. Statistical properties of x-ray clusters: Analytic and numerical comparisons. *Astrophys. J.* 495, 80–99. doi:10.1086/305262. arXiv:astro-ph/9710107.
- Camps, P., Baes, M., 2015. SKIRT: An advanced dust radiative transfer code with a user-friendly architecture. *Astron. Comput.* 9, 20–33. doi:10.1016/j.ascom.2014.10.004. arXiv:1410.1629.
- Carlberg, R.G., Couchman, H.M.P., Thomas, P.A., 1990. Cosmological velocity bias. *Astrophys. J.* 352, L29–L32. doi:10.1086/185686.
- Chabrier, G., 2003. Galactic stellar and substellar initial mass function. *Publ. Astron. Soc. Pac.* 115, 763–795. doi:10.1086/376392. arXiv:astro-ph/0304382.
- Chang, Y.Y., van der Wel, A., da Cunha, E., Rix, H.W., 2015. Stellar masses and star formation rates for 1M galaxies from SDSS+WISE. *Astrophys. J. Suppl.* 219, 8. doi:10.1088/0067-0049/219/1/8. arXiv:1506.00648.
- Crain, R.A., et al., 2015. The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations. *Mon. Not. R. Astron. Soc.* 450, 1937–1961. doi:10.1093/mnras/stv725. arXiv:1501.01311.
- Cullen, L., Dehnen, W., 2010. Inviscid smoothed particle hydrodynamics. *Mon. Not. R. Astron. Soc.* 408, 669–683. doi:10.1111/j.1365-2966.2010.17158.x. arXiv:1006.1524.
- Dalla Vecchia, C., Schaye, J., 2012. Simulating galactic outflows with thermal supernova feedback. *Mon. Not. R. Astron. Soc.* 426, 140–158. doi:10.1111/j.1365-2966.2012.21704.x. arXiv:1203.5667.
- Davis, M., Efstathiou, G., Frenk, C.S., White, S.D.M., 1985. The evolution of large-scale structure in a universe dominated by cold dark matter. *Astrophys. J.* 292, 371–394. doi:10.1086/163168.
- De Lucia, G., Blaizot, J., 2007. The hierarchical formation of the brightest cluster galaxies. *Mon. Not. R. Astron. Soc.* 375, 2–14. doi:10.1111/j.1365-2966.2006.11287.x. arXiv:astro-ph/0606519.

- Doi, M., et al., 2010. Photometric response functions of the Sloan Digital Sky Survey imager. *Astron. J.* 139, 1628–1648. doi:10.1088/0004-6256/139/4/1628. arXiv:1002.3701.
- Dolag, K., Borgani, S., Murante, G., Springel, V., 2009. Substructures in hydrodynamical cluster simulations. *Mon. Not. R. Astron. Soc.* 399, 497–514. doi:10.1111/j.1365-2966.2009.15034.x. arXiv:0808.3401.
- Dubois, Y., et al., 2014. Dancing in the dark: galactic properties trace spin swings along the cosmic web. *Mon. Not. R. Astron. Soc.* 444, 1453–1468. doi:10.1093/mnras/stu1227. arXiv:1402.1165.
- Durier, F., Dalla Vecchia, C., 2012. Implementation of feedback in smoothed particle hydrodynamics: towards concordance of methods. *Mon. Not. R. Astron. Soc.* 419, 465–478. doi:10.1111/j.1365-2966.2011.19712.x. arXiv:1105.3729.
- Fakhouri, O., Ma, C.P., 2008. The nearly universal merger rate of dark matter haloes in Λ CDM cosmology. *Mon. Not. R. Astron. Soc.* 386, 577–592. doi:10.1111/j.1365-2966.2008.13075.x. arXiv:0710.4567.
- Furlong, M., et al., 2015. Evolution of galaxy stellar masses and star formation rates in the EAGLE simulations. *Mon. Not. R. Astron. Soc.* 450, 4486–4504. doi:10.1093/mnras/stv852. arXiv:1410.3485.
- Genel, S., Genzel, R., Bouché, N., Naab, T., Sternberg, A., 2009. The halo merger rate in the millennium simulation and implications for observed galaxy merger fractions. *Astrophys. J.* 701, 2002–2018. doi:10.1088/0004-637X/701/2/2002. arXiv:0812.3154.
- Groves, B., Dopita, M.A., Sutherland, R.S., Kewley, L.J., Fischera, J., Leitherer, C., Brandl, B., van Breugel, W., 2008. Modeling the pan-spectral energy distribution of starburst galaxies. IV. The controlling parameters of the starburst SED. *Astrophys. J. Suppl.* 176, 438–456. doi:10.1086/528711. arXiv:0712.1824.
- Henriques, B.M.B., White, S.D.M., Thomas, P.A., Angulo, R., Guo, Q., Lemson, G., Springel, V., Overzier, R., 2015. Galaxy formation in the Planck cosmology—I. Matching the observed evolution of star formation rates, colours and stellar masses. *Mon. Not. R. Astron. Soc.* 451, 2663–2680. doi:10.1093/mnras/stv705. arXiv:1410.0365.
- Hewett, P.C., Warren, S.J., Leggett, S.K., Hodgkin, S.T., 2006. The UKIRT infrared deep sky survey ZY JHK photometric system: passbands and synthetic colours. *Mon. Not. R. Astron. Soc.* 367, 454–468. doi:10.1111/j.1365-2966.2005.09969.x. arXiv:astro-ph/0601592.
- Hopkins, P.F., 2013. A general class of Lagrangian smoothed particle hydrodynamics methods and implications for fluid mixing problems. *Mon. Not. R. Astron. Soc.* 428, 2840–2856. doi:10.1093/mnras/sts210. arXiv:1206.5006.
- Jenkins, A., 2010. Second-order Lagrangian perturbation theory initial conditions for resimulations. *Mon. Not. R. Astron. Soc.* 403, 1859–1872. doi:10.1111/j.1365-2966.2010.16259.x. arXiv:0910.0258.
- Jenkins, A., 2013. A new way of setting the phases for cosmological multiscale Gaussian initial conditions. *Mon. Not. R. Astron. Soc.* 434, 2094–2120. doi:10.1093/mnras/stt1154. arXiv:1306.5968.
- Jiang, L., Helly, J.C., Cole, S., Frenk, C.S., 2014. N-body dark matter haloes with simple hierarchical histories. *Mon. Not. R. Astron. Soc.* 440, 2115–2135. doi:10.1093/mnras/stu390. arXiv:1311.6649.
- Katz, N., Hernquist, L., Weinberg, D.H., 1992. Galaxies and gas in a cold dark matter universe. *Astrophys. J.* 399, L109–L112. doi:10.1086/186619.
- Khandai, N., Di Matteo, T., Croft, R., Wilkins, S., Feng, Y., Tucker, E., DeGraf, C., Liu, M.S., 2015. The MassiveBlack-II simulation: the evolution of haloes and galaxies to $z = 0$. *Mon. Not. R. Astron. Soc.* 450, 1349–1374. doi:10.1093/mnras/stv627. arXiv:1402.0888.
- Lacey, C.G., et al., 2015. A unified multi-wavelength model of galaxy formation. *ArXiv e-prints* arXiv:1509.08473.
- Lagos, C.d.P., et al., 2015. Molecular hydrogen abundances of galaxies in the EAGLE simulations. *Mon. Not. R. Astron. Soc.* 452, 3815–3837. doi:10.1093/mnras/stv1488. arXiv:1503.04807.
- Lemson, G., Springel, V., 2006. Cosmological simulations in a relational database: Modelling and storing merger trees. In: Gabriel, C., Arviset, C., Ponz, D., Enrique, S. (Eds.), *Astronomical Data Analysis Software and Systems XV*. In: *Astronomical Society of the Pacific Conference Series*, vol. 351, p. 212.
- Lemson, G., Virgo Consortium, T., 2006. Halo and galaxy formation histories from the millennium simulation: Public release of a VO-oriented and SQL-queryable database for studying the evolution of galaxies in the Λ CDM cosmogony. *ArXiv Astrophysics e-prints* arXiv:astro-ph/0608019.
- Lupton, R., Blanton, M.R., Fekete, G., Hogg, D.W., O’Mullane, W., Szalay, A., Wherry, N., 2004. Preparing red–green–blue images from CCD data. *Publ. Astron. Soc. Pac.* 116, 133–137. doi:10.1086/382245. arXiv:astro-ph/0312483.
- Nelson, D., et al., 2015. The Illustris simulation: Public data release. *ArXiv e-prints*. doi:10.1016/j.ascom.2015.09.003. arXiv:1504.00362.
- Okamoto, T., Shimizu, I., Yoshida, N., 2014. Reproducing cosmic evolution of galaxy population from $z = 4$ to 0. *Publ. Astron. Soc. Japan* 66, 70. doi:10.1093/pasj/psu046. arXiv:1404.7579.
- Oppenheimer, B.D., Davé, R., Kereš, D., Fardal, M., Katz, N., Kollmeier, J.A., Weinberg, D.H., 2010. Feedback and recycled wind accretion: assembling the $z = 0$ galaxy mass function. *Mon. Not. R. Astron. Soc.* 406, 2325–2338. doi:10.1111/j.1365-2966.2010.16872.x. arXiv:0912.0519.
- Peebles, P.J.E., 1980. The large-scale structure of the universe.
- Planck Collaboration, et al., 2014. Planck 2013 results. I. Overview of products and scientific results. *Astron. Astrophys.* 571, A1. doi:10.1051/0004-6361/201321529. arXiv:1303.5062.
- Porter, L.A., Somerville, R.S., Primack, J.R., Johansson, P.H., 2014. Understanding the structural scaling relations of early-type galaxies. *Mon. Not. R. Astron. Soc.* 444, 942–960. doi:10.1093/mnras/stu1434. arXiv:1407.0594.

- Price, D.J., 2008. Modelling discontinuities and Kelvin Helmholtz instabilities in SPH. *J. Comput. Phys.* 227, 10040–10057. doi:10.1016/j.jcp.2008.08.011. arXiv:0709.2772.
- Puchwein, E., Springel, V., 2013. Shaping the galaxy stellar mass function with supernova- and AGN-driven winds. *Mon. Not. R. Astron. Soc.* 428, 2966–2979. doi:10.1093/mnras/sts243. arXiv:1205.2694.
- Rahmati, A., Schaye, J., Bower, R.G., Crain, R.A., Furlong, M., Schaller, M., Theuns, T., 2015. The distribution of neutral hydrogen around high-redshift galaxies and quasars in the EAGLE simulation. *Mon. Not. R. Astron. Soc.* 452, 2034–2056. doi:10.1093/mnras/stv1414. arXiv:1503.05553.
- Rosas-Guevara, Y.M., et al., 2015. The impact of angular momentum on black hole accretion rates in simulations of galaxy formation. *Mon. Not. R. Astron. Soc.* 454, 1038–1057. doi:10.1093/mnras/stv2056. arXiv:1312.0598.
- Schaller, M., Dalla Vecchia, C., Schaye, J., Bower, R.G., Theuns, T., Crain, R.A., Furlong, M., McCarthy, I.G., 2015a. The EAGLE simulations of galaxy formation: the importance of the hydrodynamics scheme. *Mon. Not. R. Astron. Soc.* 454, 2277–2291. doi:10.1093/mnras/stv2169. arXiv:1509.05056.
- Schaller, M., et al., 2015b. Baryon effects on the internal structure of Λ CDM haloes in the EAGLE simulations. *Mon. Not. R. Astron. Soc.* 451, 1247–1267. doi:10.1093/mnras/stv1067. arXiv:1409.8617.
- Schaye, J., 2004. Star formation thresholds and galaxy edges: Why and where. *Astrophys. J.* 609, 667–682. doi:10.1086/421232. arXiv:astro-ph/0205125.
- Schaye, J., et al., 2015. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Mon. Not. R. Astron. Soc.* 446, 521–554. doi:10.1093/mnras/stu2058. arXiv:1407.7040.
- Schaye, J., Dalla Vecchia, C., 2008. On the relation between the Schmidt and Kennicutt–Schmidt star formation laws and its implications for numerical simulations. *Mon. Not. R. Astron. Soc.* 383, 1210–1222. doi:10.1111/j.1365-2966.2007.12639.x. arXiv:0709.0292.
- Springel, V., 2005. The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.* 364, 1105–1134. doi:10.1111/j.1365-2966.2005.09655.x. arXiv:astro-ph/0505010.
- Springel, V., Di Matteo, T., Hernquist, L., 2005a. Modelling feedback from stars and black holes in galaxy mergers. *Mon. Not. R. Astron. Soc.* 361, 776–794. doi:10.1111/j.1365-2966.2005.09238.x. arXiv:astro-ph/0411108.
- Springel, V., et al., 2005b. Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature* 435, 629–636. doi:10.1038/nature03597. arXiv:astro-ph/0504097.
- Springel, V., White, S.D.M., Tormen, G., Kauffmann, G., 2001. Populating a cluster of galaxies—I. Results at $[z = 0]$. *Mon. Not. R. Astron. Soc.* 328, 726–750. doi:10.1046/j.1365-8711.2001.04912.x. arXiv:astro-ph/0012055.
- Szalay, A.S., Kunszt, P.Z., Thakar, A.R., Gray, J., Slutz, D., 2000. The sloan digital sky survey and its archive. In: Manset, N., Veillet, C., Crabtree, D. (Eds.), *Astronomical Data Analysis Software and Systems IX*. In: *Astronomical Society of the Pacific Conference Series*, vol. 216, p. 405. arXiv:astro-ph/9912382.
- Trayford, J.W., et al., 2015. Colours and luminosities of $z = 0.1$ galaxies in the EAGLE simulation. *Mon. Not. R. Astron. Soc.* 452, 2879–2896. doi:10.1093/mnras/stv1461. arXiv:1504.04374.
- Vogelsberger, M., et al., 2014. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *Mon. Not. R. Astron. Soc.* 444, 1518–1547. doi:10.1093/mnras/stu1536. arXiv:1405.2921.
- Wendland, H., 1995. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv. Comput. Math.* 4, 389–396. doi:10.1007/BF02123482. URL <http://www.springerlink.com/index/10.1007/BF02123482>.
- White, S.D.M., Frenk, C.S., 1991. Galaxy formation through hierarchical clustering. *Astrophys. J.* 379, 52–79. doi:10.1086/170483.
- Wiersma, R.P.C., Schaye, J., Smith, B.D., 2009a. The effect of photoionization on the cooling rates of enriched, astrophysical plasmas. *Mon. Not. R. Astron. Soc.* 393, 99–107. doi:10.1111/j.1365-2966.2008.14191.x. arXiv:0807.3748.
- Wiersma, R.P.C., Schaye, J., Theuns, T., Dalla Vecchia, C., Tornatore, L., 2009b. Chemical enrichment in cosmological, smoothed particle hydrodynamics simulations. *Mon. Not. R. Astron. Soc.* 399, 574–600. doi:10.1111/j.1365-2966.2009.15331.x. arXiv:0902.1535.