

Annual Review of Astronomy and Astrophysics

Hydrodynamical Simulations of the Galaxy Population: Enduring Successes and Outstanding Challenges

Robert A. Crain^{1,*} and Freeke van de Voort^{2,*}

¹Astrophysics Research Institute, Liverpool John Moores University, Liverpool, United Kingdom; email: r.a.crain@ljmu.ac.uk

²Cardiff Hub for Astrophysics Research and Technology, School of Physics and Astronomy, Cardiff University, Cardiff, United Kingdom; email: freeke@astro.cf.ac.uk

Annu. Rev. Astron. Astrophys. 2023. 61:473–515

First published as a Review in Advance on
June 20, 2023

The *Annual Review of Astronomy and Astrophysics* is
online at astro.annualreviews.org

<https://doi.org/10.1146/annurev-astro-041923-043618>

Copyright © 2023 by the author(s). This work is
licensed under a Creative Commons Attribution 4.0
International License, which permits unrestricted
use, distribution, and reproduction in any medium,
provided the original author and source are credited.
See credit lines of images or other third-party
material in this article for license information.

*The authors contributed equally to this manuscript.

Keywords

galaxies, galaxy formation, galaxy halos, gas, numerical simulations

Abstract

We review the progress in modeling the galaxy population in hydrodynamical simulations of the Λ CDM cosmogony. State-of-the-art simulations now broadly reproduce the observed spatial clustering of galaxies; the distributions of key characteristics, such as mass, size, and SFR; and scaling relations connecting diverse properties to mass. Such improvements engender confidence in the insight drawn from simulations. Many important outcomes, however, particularly the properties of circumgalactic gas, are sensitive to the details of the subgrid models used to approximate the macroscopic effects of unresolved physics, such as feedback processes. We compare the outcomes of leading simulation suites with observations, and with each other, to identify the enduring successes they have cultivated and the outstanding challenges to be tackled with the next generation of models. Our key conclusions include the following:

- Realistic galaxies can be reproduced by calibrating the ill-constrained parameters of subgrid feedback models. Feedback is dominated by stars and black holes in low-mass and high-mass galaxies, respectively.
- Adjusting or disabling the processes implemented in simulations can elucidate their impact on observables, but outcomes can be degenerate.

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

- Similar galaxy populations can emerge in simulations with dissimilar feedback implementations. However, these models generally predict markedly different gas flow rates into, and out of, galaxies and their halos. CGM observations are thus a promising means of breaking this degeneracy and guiding the development of new feedback models.

Contents

1. INTRODUCTION	474
2. METHODS	478
2.1. Initial and Boundary Conditions	478
2.2. Gravitational and Hydrodynamical Evolution	479
2.3. Subgrid Methods	480
2.4. Calibration of Subgrid Feedback Models	487
2.5. Verification and Convergence	489
3. KEY PROPERTIES OF SIMULATED GALAXY POPULATIONS	490
3.1. Galaxy Stellar Mass Function	490
3.2. Size and Morphology	491
3.3. Galaxy Clustering	493
3.4. Star-Formation Histories	494
3.5. Galaxy Colors	495
4. GALAXY SCALING RELATIONS	496
4.1. Supermassive Black Holes	496
4.2. The Star-Forming Main Sequence	497
4.3. The Tully–Fisher Relation	497
4.4. The Mass–Metallicity Relations	498
4.5. Cold Gas in Galaxies	498
5. COSMIC GAS	499
5.1. Absorption System Statistics	499
5.2. Physical Properties of the Circumgalactic Medium	500
5.3. Gas Inflows and Outflows	500
5.4. Halo Baryon Fractions	501
6. THE INFLUENCE OF ENVIRONMENT	503
6.1. Satellite Galaxies: Stripping and Starvation	503
6.2. Environmental Effects: Beyond Halo Mass	504
7. FUTURE OUTLOOK	505
8. SUMMARY	506

1. INTRODUCTION

The present-day galaxy population exhibits a remarkable diversity of characteristics, such as masses, star-formation rates (SFRs), morphologies, nuclear activity, and gas and dust content. The population’s spatial distribution is highly structured and heterogeneous: Galaxies can be found in isolation in low-density environments or can comprise the population of rich clusters. Observations of the distant cosmos reveal that both the properties and spatial distribution of the galaxy population have evolved markedly over nearly 14 billion years of cosmic history. Reconciliation

of this panoply of observed characteristics with a comprehensive theory of galaxy formation is a challenge at the frontier of the natural sciences.

Owing to the nonlinearity of the collapse, hierarchical assembly, and relaxation of protogalactic structure, and the complexity of the myriad physical processes that influence galaxies, direct numerical simulation of the evolution of representative cosmological volumes is, in principle, the most desirable method of approaching the challenge. In recent years, this approach has emerged as the foremost means of interpreting observations, particularly those from highly multiplexed galaxy surveys, and of seeking a clearer understanding of the origin of galaxy properties. The widespread adoption of hydrodynamical simulations of representative cosmological volumes has been driven primarily by major improvements in their correspondence with the observed galaxy population, coupled with improved access to the simulation data, a development catalyzed by the public release of data from major simulation campaigns (e.g., Nelson et al. 2015, 2019; McAlpine et al. 2016; Villaescusa-Navarro et al. 2023).

It is easy to now take for granted the availability of these versatile and realistic models. However, although simulations of representative volumes have yielded good agreement with observations of diffuse intergalactic gas (e.g., as traced by the Lyman- α forest) for over two decades (Theuns et al. 1998, Davé et al. 1999), reproducing the properties and spatial distribution of galaxies and their star-forming gas reservoirs has proven a more stubborn challenge. For many years, simulated galaxies formed far too many stars, particularly at early cosmic epochs when the cosmic inflow rate is high, thus ending up too massive, too compact, and with too little angular momentum (e.g., Navarro & Benz 1991, Navarro et al. 1995, Sommer-Larsen et al. 1999, Navarro & Steinmetz 2000). Unsurprisingly then, they also exhibited unrealistic surface density and rotation velocity profiles (Abadi et al. 2003a,b). Simulations of representative cosmic volumes yielded galaxy stellar mass functions (GSMFs) with the wrong shape and normalization, generally yielding too many galaxies at fixed stellar mass (e.g., Crain et al. 2009, Lackner et al. 2012, Khandai et al. 2015).

Mitigation of this overcooling of gas into stars, and the concomitant spurious transfer of angular momentum from gas to dark matter (DM), was demonstrated early in the history of hydrodynamical simulations of galaxies via the inclusion of gas heating mechanisms (e.g., Katz & Gunn 1991, Mihos & Hernquist 1994, Navarro & Steinmetz 1997, Weil et al. 1998). The inclusion of energetic feedback mechanisms, as a means to regulate gas cooling and star formation, has fostered major improvements in the realism of simulated galaxies, for example, enabling the realization of individual disk galaxies with encouragingly realistic surface density and rotation profiles (e.g., Governato et al. 2004, Okamoto et al. 2005, Governato et al. 2007, Guedes et al. 2011).

However, the macroscopic efficiencies of feedback mechanisms are governed by microphysics acting on spatial scales several orders of magnitude below the resolution scale of galaxy population simulations (e.g., Orlando et al. 2005), precluding their calculation *ab initio*. These processes (and others) are therefore treated with simplified subgrid models (see Section 2.3), which approximate the effects of unresolved processes and couple them to numerically resolved scales, thereby producing an effective model of galaxy formation. In the absence of authoritative empirical constraints on how microphysics influences macroscopic scales (though see Lopez et al. 2011, Rosen et al. 2014), the subgrid implementations of feedback in popular simulation codes that have emerged are diverse and can produce conspicuously dissimilar outcomes when applied to identical initial conditions (see Scannapieco et al. 2012, and references therein). The dramatic variation of outcomes that emerge from controlled suites of simulations in which (only) the subgrid implementation of feedback is changed, and/or their parameters are varied systematically over plausible ranges, highlights that the influence of feedback is the most important systematic uncertainty in galaxy-formation modeling (Oppenheimer et al. 2010, Schaye et al. 2010, Vogelsberger et al. 2013, Kim et al. 2014).

GSMF: galaxy stellar mass function

The most productive strategy that has emerged in response to this uncertainty is to calibrate the parameters of subgrid feedback models (see Section 2.4), with the aim of reproducing judiciously chosen observable characteristics of the galaxy population. Clearly, the characteristics used for the calibration cannot be considered as predictions of the simulation, but other properties can be considered as outcomes stemming from the implemented physical processes, as long as they are not compromised by the simplifications of the subgrid modeling. This approach has been used, to varying degrees, by several flagship-scale simulations of the galaxy population, including *Illustris* (Genel et al. 2014, Vogelsberger et al. 2014, Nelson et al. 2015), *Magneticum* (Hirschmann et al. 2014), *EAGLE* (Crain et al. 2015, Schaye et al. 2015, McAlpine et al. 2016), *BlueTides* (Feng et al. 2016), *Romulus* (Tremmel et al. 2017), *IllustrisTNG* (Nelson et al. 2018b, Pillepich et al. 2018b, Springel et al. 2018), *FABLE* (Henden et al. 2018), and *SIMBA* (Davé et al. 2019). An example of the DM, gas, and stellar light distributions that emerge from simulating a Milky Way–mass DM halo with a modern galaxy-formation model is shown in **Figure 1**. The central and satellite galaxies visible have a clear gaseous and DM component as well. Further substructure is present in the DM and gas without a counterpart in either of the other components.

The key advantage of the calibration approach is that it steers simulations toward the production of a broadly realistic simulated galaxy population. The galaxies interact self-consistently with the intergalactic medium (IGM), enabling the address of diverse lines of inquiry. Several of the aforementioned projects have demonstrated, for example, that it is possible to reproduce the present-day GSMF, from the mass scale of dwarf galaxies to that of central group galaxies, with an accuracy comparable with the systematic uncertainty on the observational measurement. A subset of the models also consider diagnostics inferred from X-ray observations of galaxy groups/clusters, in order to ensure reasonable reproduction of the properties of intragroup/intracluster gas. Simulations of realistic galaxies also enable additional model components to be “bolted on” either on the fly or in postprocessing, for example, to follow the formation and evolution of globular clusters (Pfeffer et al. 2018) or to predict rates of gamma ray bursts (Metha & Trenti 2020). The realism

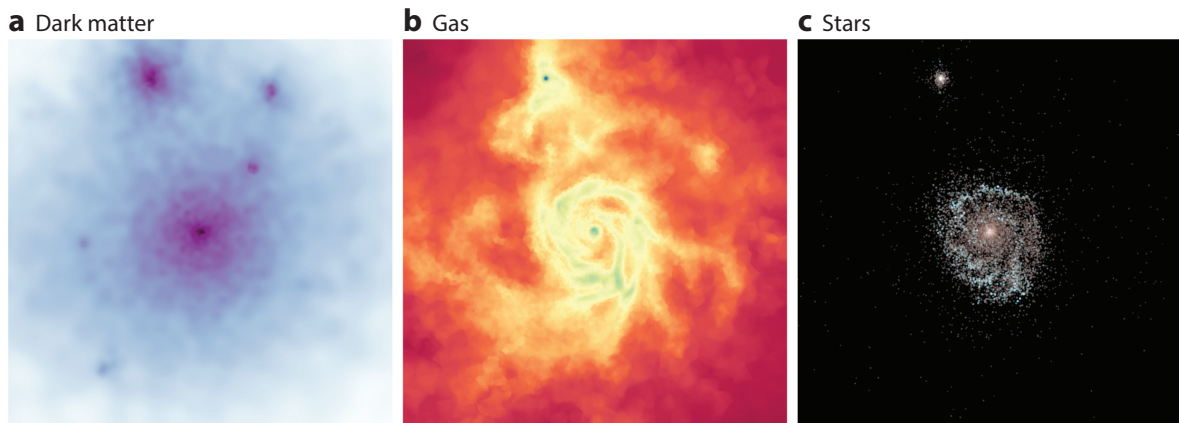


Figure 1

An illustrative example of the level of detail with which modern cosmological simulations are able to reproduce the structure of galaxies similar in mass ($M_{\text{halo}} \sim 10^{12} M_{\odot}$) to the Milky Way, based on *IllustrisTNG*. From left to right, the panels show the dark matter surface density, the gas surface density, and a three-color image of the stellar surface brightness from the *U*, *B*, and *K* bands. Each image has a 140-kpc field of view. The central galaxy is gas rich, star forming, and kinematically disk dominated, with a significant bulge component. It is accompanied by a gas-rich satellite galaxy and dark matter substructure without an associated baryonic component. The gas density shows clumpy and filamentary structure not present in the dark matter and stars.

and versatility of the simulations has resulted in a remarkable scientific impact: A conservative estimate based on searches using the NASA Astrophysics Data System is that data and/or data products from these models feature in over one thousand astrophysics research articles published since 2014.

It is imperative to remember that much of galaxy-formation and evolution modeling remains distantly removed from fundamental, *ab initio* theory. There is no guarantee that the reproduction of a particular observable represents a unique solution; i.e., different subgrid models may yield similarly successful outcomes. It is possible that the approximations necessitated by subgrid models may deliver success only at a particular resolution or only at the expense of a failure elsewhere. Simulations are therefore often better suited to offering qualitative insight rather than quantitative predictions: Some of the most instructive outcomes from simulation suites have stemmed from varying the efficiency of physical mechanisms to isolate their influence (and hence indicate why a simulation reproduces a particular observable), or from a failure to reproduce particular observational measurements and so illuminate a fundamental shortcoming of the implemented physics or even the adopted cosmogony.

The chief objective of this review is therefore to offer a critical assessment of the improved understanding of the formation and evolution of the galaxy population fostered by the current generation of state-of-the-art hydrodynamical (or magnetohydrodynamical) simulations of the Λ -cold dark matter (Λ CDM) cosmogony. To this end, we highlight the key successes of the simulations that we believe will endure and the shortcomings or absences of consensus that present outstanding challenges to be addressed by future models. Where possible, we offer candid explanations for the origin of shortcomings, which we hope will be particularly helpful to nonspecialists and new practitioners. We focus on simulations that follow reasonably representative cosmological volumes ($L \simeq 100$ comoving Mpc, hereafter cMpc) at fixed resolution, and thus yield simulated populations of galaxies. These simulations are complementary to zoom simulations of individual galaxies, but we focus on the former because in general it is simpler to compare their outcomes with observations and to characterize their convergence behavior (Section 2.5), and because one cannot confuse trends due to resolution with those due to galaxy or halo mass, which is a danger when analyzing suites of zoom simulations that maximize the resolution at each mass scale. A major advantage of this type of simulation is the diversity of the lines of inquiry that they enable, so we necessarily restrict ourselves to reviewing their more fundamental outcomes.

We explore results primarily (but not exclusively) from the EAGLE, Horizon-AGN (Dubois et al. 2012, Kaviraj et al. 2017), IllustrisTNG, and SIMBA projects, which are now somewhat mature simulation suites that have been studied in detail (key numerical details of these simulations are given in Table 1). The flagship simulation of each project follows the evolution of

Λ CDM: Λ -cold dark matter

Table 1 Details of Simulations used for Illustrative Examples

Simulation	Hydro method	Box size (cMpc)	DM particles per dimension	Baryonic mass resolution (M_{\odot})	Baryonic spatial resolution ^a (pkpc)
EAGLE ^b	Modern SPH	100.0	1504	1.9×10^6	0.70
Horizon-AGN	AMR	142.0	1024	1.0×10^7	1.00
IllustrisTNG ^c	Moving mesh	110.7	1820	1.4×10^6	0.19
SIMBA ^d	MFM	147.0	1024	1.8×10^7	0.74

Abbreviations: AMR, adaptive mesh refinement; DM, dark matter; MFM, meshless finite mass; SPH, smoothed particle hydrodynamics.

^aMaximum size of the gravitational softening for gas.

^bExcept where explicitly stated, we consider the Ref-L100N1504 simulation from the EAGLE suite.

^cExcept where explicitly stated, we consider the TNG-100 simulation from the IllustrisTNG suite.

^dExcept where explicitly stated, we consider the m100n1024 simulation from the SIMBA suite.

SPH: smoothed
particle
hydrodynamics

AMR: adaptive mesh
refinement

a cosmological volume of $L \sim 100$ cMpc with baryonic mass (spatial) resolution of $\sim 10^{6-7} M_{\odot}$ (~ 1 proper kpc, hereafter pkpc, in the ISM), and yields a galaxy population that is, to a greater or lesser degree, broadly realistic. They span the range of hydrodynamics solvers widely in use [modern smoothed particle hydrodynamics (SPH), adaptive mesh refinement (AMR), and hybrid Lagrangian–Eulerian approaches; see Section 2], and adopt a diverse range of subgrid implementations of baryonic physics, each differing from the others in at least several major aspects. Of these, only *IllustrisTNG* solves the equations of magnetohydrodynamics; the others use pure hydrodynamics. *EAGLE* and *IllustrisTNG* are the most well-studied and readily accessible of these simulations and, consequently, feature more frequently in this review. Our focus on these simulation suites is to provide illustrative examples and should not be misinterpreted as implicit criticism of other simulation campaigns nor dismissal of their successes.

The review is structured as follows. We discuss the methods used in the execution, calibration, and analysis of galaxy population simulations in Section 2. We review in Sections 3 and 4 the key properties of simulated galaxy populations and scaling relations connecting diverse properties of simulated galaxies to their stellar mass, respectively. Section 5 focuses on the properties of the gaseous environments of simulated galaxies, and Section 6 examines environmental influences. We discuss likely future directions for the discipline in Section 7 and provide a brief summary of our conclusions in Section 8.

2. METHODS

2.1. Initial and Boundary Conditions

Cosmological simulations begin from initial conditions that specify the fluctuations in the density field at an epoch at which there is no significant nonlinearity, generally $z \gtrsim 100$. From this initial state, the equations of motion can be integrated forward into the nonlinear regime numerically. The initial matter power spectrum (which specifies the density contrast relative to the mean density as a function of spatial scale) is usually expressed as the product of an initial spectrum (with random phases) resulting from inflation, and a transfer function (which can be calculated with a Boltzmann solver such as *CAMB*; Lewis & Challinor 2011) that represents the subsequent linear evolution of each mode. Cosmological simulators are therefore in the fortunate position of having initial conditions that are well constrained by observations of the cosmic microwave background (CMB) radiation at $z \approx 1,100$.

Generating Gaussian random fields with a specified power spectrum, a process first detailed by Efstathiou et al. (1985), has evolved into a specialist discipline. We only briefly summarize the process here and encourage those with a particular interest to read the work of, for example, Jenkins (2010, 2013), Hahn & Abel (2011), and Hahn et al. (2021). The process comprises two key stages to realize the appropriate density fluctuations: the creation of uniform particle distribution throughout the simulation volume, followed by the application of displacements to positions and velocities of the particles, or, alternatively, small adjustments to the masses of particles. The unperturbed particle distribution is usually constructed by tiling a cubic grid or a “glass” distribution (White 1994). Random particle distributions are unsuitable because they exhibit a white noise power spectrum that, even in the absence of the intended displacements, fosters rapidly growing nonlinear structure. The displacements are calculated with linear or low order perturbation theory. Periodic boundary conditions are applied to opposing faces of the volume, ensuring that its mean density remains fixed and that no artificial boundaries are imposed.

The finite size of the simulation domain imposes an upper limit to the spatial scale of the fluctuations that can be sampled within it (correspondingly, there is a minimum wavenumber of $k = 2\pi/L$). As large-scale fluctuations seed the formation of rare features in the cosmic large-scale

structure, such as galaxy clusters and large voids, the emergent space density of such features is underestimated (Bagla & Prasad 2006, Reed et al. 2007). Volumes of $L \simeq 100$ cMpc are too small to realize galaxy clusters (whose present-day space density is $<10^{-6} \text{ Mpc}^{-3}$), such that detailed examination with (magneto)hydrodynamical simulations of the galaxy population within them requires the use of zoomed initial conditions (e.g., Bahé et al. 2017, Cui et al. 2018, Tremmel et al. 2019). The scarcity of rare features in finite simulation volumes also makes them sensitive to cosmic variance, as different realizations of the same initial density fluctuations can result in statistically significant differences in the emergent large-scale structure.

2.2. Gravitational and Hydrodynamical Evolution

Modeling the formation of the galaxy population entails solving the partial differential equations that govern the temporal evolution of the cosmic matter and radiation fields. Simulations of the galaxy population apply numerical techniques to solve the equations governing the gravitational evolution of matter, the hydrodynamical evolution of gas, and in some cases the interaction of gas with evolving radiation and magnetic fields. Other physical processes are treated with subgrid models. Since the key numerical techniques are documented, and their advantages and shortcomings discussed, in detail elsewhere (see, e.g., Springel 2010a, Price 2012, Teyssier 2015), we only briefly discuss them. We focus primarily on a discussion of subgrid methods (Section 2.3), because these dominate systematic uncertainties at the resolution of the simulations examined here.

DM is most commonly treated as a collisionless fluid (but for self-interacting treatments see, e.g., Davé et al. 2001, Robertson et al. 2017), whose evolution in the continuum limit is described by the collisionless Boltzmann equation (CBE), under the influence of the gravitational potential given by Poisson’s equation. The potential is assumed to be Newtonian, because velocities on resolved scales are nonrelativistic. The high dimensionality of the CBE necessitates solving the coupled equations using a finite set of N -body tracer particles that sample the fluid’s phase space distribution. The $\mathcal{O}(N^2)$ scaling of the computational cost of solving Poisson’s equation (stemming from the long-range nature of gravity that requires consideration of $N - 1$ contributions to the potential at each particle) can be reduced to a scaling on the order of $\mathcal{O}(N \log N)$ by approximating the contribution from distant particles. This is achieved via multipole expansion (e.g., Barnes & Hut 1986, Carrier et al. 1988) and/or by mapping the tracer distribution to a mesh and solving in Fourier space using fast transform methods (e.g., Hockney & Eastwood 1981). To prevent the unphysical scattering of close particle pairs, the gravitational force is softened on small scales using a kernel function, for which forms of varying complexity have been proposed (e.g., Monaghan 1985, Wendland 1995). It is common to adopt a softening scale that is fixed in comoving units (e.g., $\simeq 1/25$ of the mean interparticle separation) and limited to a maximum proper size to ensure that the internal structure of DM halos can be resolved at late cosmic epochs (Power et al. 2003).

Cosmic gas is assumed to be ideal, collisional, inviscid, and nonconducting, enabling its dynamics to be described by the Euler equations rather than the more general Navier–Stokes equations. Traditionally, solution of the equations has been via two distinct approaches for discretizing the fluid: in either volume (the Eulerian approach used by mesh-based schemes) or mass (the Lagrangian approach used by particle-based schemes). Each approach has well-advertised shortcomings that contribute to inconsistent results when applied to relatively simple cosmological structure formation problems involving nonradiative hydrodynamics (e.g., Frenk et al. 1999, Sembolini et al. 2016). Eulerian methods are the de facto standard for many computational fluid dynamics problems, but the dynamic range needed to resolve galaxies in cosmological volumes demands the added complexity of adaptively refined meshes (AMR; e.g., Abel et al. 2002), which may still fail to adequately capture the gravitational collapse of low-contrast fluctuations at early

SMBH: supermassive
black hole

CGM: circumgalactic
medium

times (O’Shea et al. 2005) and cause overmixing of gas with differing entropies, e.g., in the cores of galaxy clusters (Wadsley et al. 2008, Mitchell et al. 2009). The inherent Galilean noninvariance of mesh methods (Tasker et al. 2008) is particularly undesirable because the relative velocities of galaxies are generally much greater than the sound speed of the gas bound to them.

SPH (Gingold & Monaghan 1977, Lucy 1977) samples the fluid with tracer particles, naturally adapting the resolution within overdense structures. SPH is, hence, well suited to following the hierarchical growth of structure and enables the self-gravity of the gas to be treated in a fashion identical to that of the DM but requires the inclusion of an artificial viscosity term to capture shocks (Monaghan 1997), which are resolved poorly. Traditional SPH implementations suffer from multivalued particle pressure at contact discontinuities that create unphysical surface tension and inhibit phase mixing (e.g., Agertz et al. 2007), though this problem can be mitigated via the use of various adaptations (Ritchie & Thomas 2001, Price 2008, Hopkins 2013, Wadsley et al. 2017, Rosswog 2020, Borrow et al. 2022c), which have become collectively known colloquially as modern or corrected SPH.

In the past decade, several software packages have emerged that treat cosmic gas with sophisticated schemes that seek to capture the advantages of both Lagrangian and Eulerian approaches, i.e., continuous adaptability of resolution and geometry of the (magneto)hydrodynamics calculation, Galilean invariance, accurate treatment of fluid mixing, and high-fidelity shock capturing without the use of artificial viscosity. These schemes (sometimes referred to as arbitrary Lagrangian–Eulerian, or ALE) typically use a mesh that deforms and moves with the fluid flow (moving mesh; e.g., Springel 2010b, Vandenbroucke & De Rijcke 2016) or solve the Riemann problem without a mesh (e.g., meshless finite volume, MFV, or meshless finite mass, MFM; Hopkins 2015) (though note that MFM is not strictly an ALE method because it uses resolution elements with fixed mass). These approaches have proven successful for particular problems, but simultaneously realizing the benefits of Eulerian and Lagrangian methods generally incurs a marked increase in computational cost and memory footprint.

Nonstandard refinement criteria have been used with both AMR and moving mesh schemes, basing control of the resolution of the hydrodynamics calculation on conditions other than the fluid density. This has enabled examination of fluid flows in greater detail, for example, around supermassive black holes (SMBHs; Curtis & Sijacki 2015) or in the gaseous halos around galaxies (i.e., the circumgalactic medium or CGM; Hummels et al. 2019, Peebles et al. 2019, van de Voort et al. 2019).

2.3. Subgrid Methods

The need to uniformly sample representative cosmic volumes restricts the feasible resolution of galaxy population simulations. Let us consider a brief example: Assuming $\Omega_b = 0.05$ and $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, uniformly sampling the gas fluid in a cubic volume of $L = 100 \text{ cMpc}$ with tracers of mass of $10^5 M_\odot$ (comparable with the stellar mass of ultrafaint dwarf galaxies) requires $\simeq 68$ billion fluid elements ($N = 4, 082^3$). Simulations including a broad suite of baryonic physics can require $\sim 1 \text{ kB}$ per fluid element, incurring a total memory footprint for the baryonic part of the calculation of $\simeq 68 \text{ TB}$. Assuming an equal number of N -body particles for the DM, at 100 bytes per particle, brings the footprint to $\simeq 75 \text{ TB}$. With current high-performance computing facilities typically having 2–4 GB of memory per core, our example would require execution on between 20,000 and 40,000 cores. Load balancing relatively high-resolution simulations at this scale remains extremely challenging despite the development of sophisticated schemes for this purpose (see, e.g., Menon et al. 2015, Schaller et al. 2016, Weinberger et al. 2020), largely because of the extreme dynamic range of the time step hierarchy. As a result, calculations of this

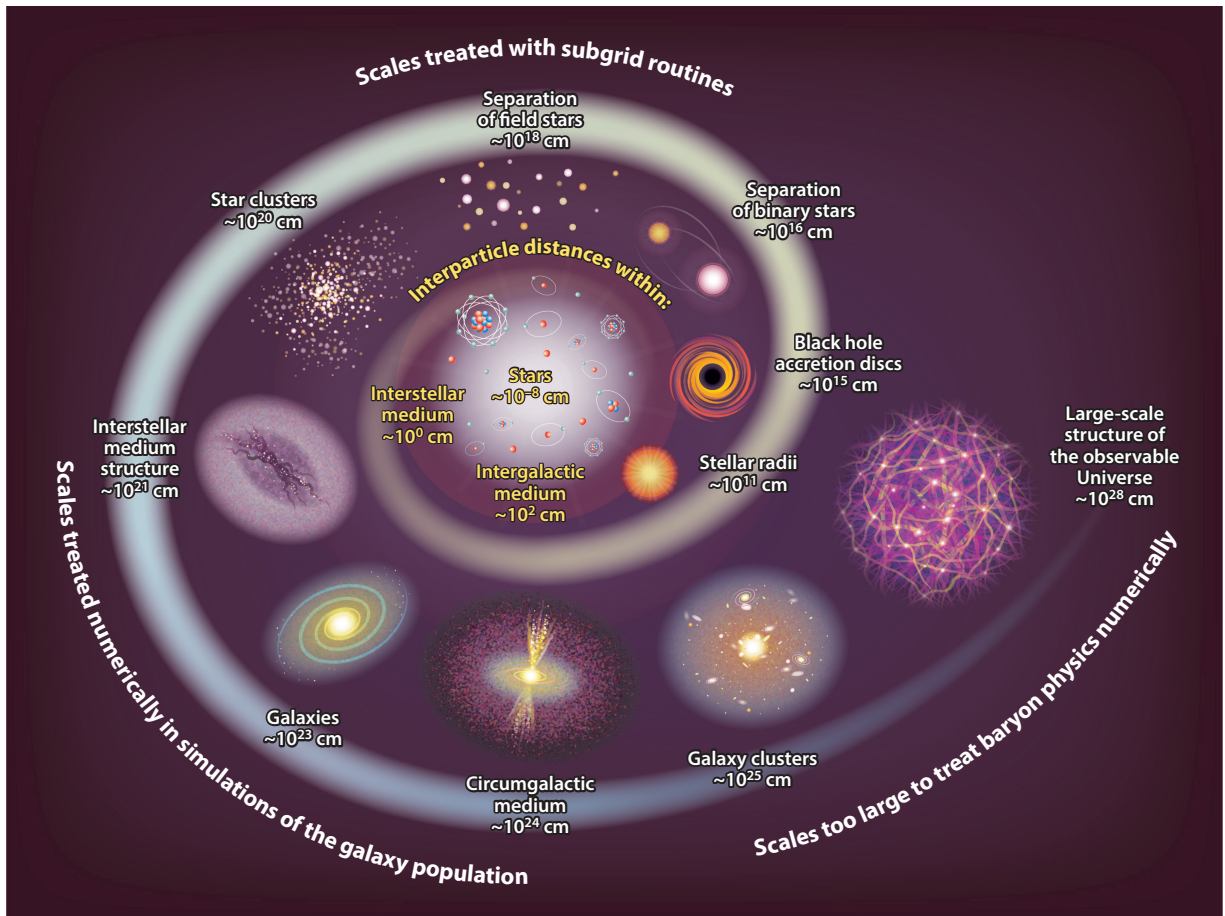


Figure 2

Schematic illustration of the extreme dynamic range of the length scales of the physical processes influencing the formation and evolution of galaxies. Cosmological, hydrodynamical simulations are able to follow cosmic structure formation up to $\mathcal{O}(100 \text{ Mpc})$ scales and have to resort to subgrid models on small scales. Future simulations of the galaxy population will push both of these boundaries, increasing their dynamic range, but the need for subgrid models will remain for years to come.

scale need to occupy these large core counts for prohibitively long periods (often several months). The challenge of accommodating the germane physical processes into simulations is illustrated schematically by **Figure 2**, which shows their characteristic length scales and highlights that subgrid models remain a critical component of simulations of the galaxy population for the foreseeable future. We review the key physical processes treated with subgrid methods below.

2.3.1. Radiative processes in cosmic gas. Radiative processes enable cosmic gas to dissipate its internal energy. In the absence of an incident radiation field, the ionization balance and cooling rate of diffuse cosmic gas is governed by two-body processes (e.g., collisional excitation and ionization, collisional recombination, and free-free emission). The cooling rate due to these processes scales as the square of the gas density, and the contribution of each elemental species scales linearly with its abundance, enabling convenient tabulation of collisional ionization equilibrium (CIE) cooling rates as a function of temperature and composition (e.g., Dalgarno & McCray 1972).

CIE: collisional
ionization equilibrium

AGN: active galactic nucleus

UVB: UV/X-ray background

ISM: interstellar medium

CIE rates (plus inverse Compton cooling due to the CMB) for gas of primordial composition were adopted by the first generation of hydrodynamical simulations of galaxies (e.g., Katz et al. 1996) and remain in use in a limited number of modern simulations [e.g., **MassiveBlack-II** (Khandai et al. 2015), **Romulus** (Tremmel et al. 2015)]. CIE rates for metal-enriched gases can differ from primordial rates by an order of magnitude (e.g., Boehringer & Hensler 1989, Sutherland & Dopita 1993) and are more widely used. However, it should be noted that uncertainties in nucleosynthetic yields (see Section 2.3.2) propagate into uncertainties in gas cooling rates. Metal cooling allows enriched ionized gas to cool below 10^4 K, though in many simulations such cooling is effectively disabled (at least for dense gas, $n_{\text{H}} \gtrsim 0.1 \text{ cm}^{-3}$) by the use of an effective equation of state or pressure floor (see Section 2.3.3). CIE rates for enriched gases (with abundance ratios similar to those of the Solar Neighborhood) are used by, for example, the **Horizon-AGN** and **BlueTides** (Feng et al. 2016) simulations.

Cosmic radiation fields can overionize gas relative to CIE, reducing the cooling rate or even yielding radiative heating (e.g., Efstathiou 1992, Gnedin & Hollon 2012). The most dramatic example is the epoch of reionization (EoR), when intergalactic neutral hydrogen (H I) is (re-)ionized by radiation from the first galaxies and active galactic nuclei (AGNs). The accurate treatment of this process in simulations, which is necessary to interpret a diverse range of observable properties of the first generations of galaxies and intervening IGM, requires the inclusion of explicit, and computationally expensive, radiative transfer calculations. Simulations of the galaxy population instead generally adopt a temporally evolving metagalactic UV/X-ray background (UVB) radiation field (for popular parameterizations, see Haardt & Madau 1996, 2001, 2012; Faucher-Giguère et al. 2009), switched on at the instantaneous reionization redshift inferred from the Thomson scattering optical depth of the CMB. The field is assumed to be spatially uniform, because the clustering scale of the photon sources is shorter than the mean free path of the photons in the diffuse IGM (e.g., Zuo 1992), though this approximation is less applicable close to galaxies. Cooling rates are then computed assuming the gas to be optically thin and in ionization equilibrium, using spectral synthesis codes such as **CLOUDY** (Ferland et al. 2017).

The photoionization rate scales linearly with density, so radiative cooling rates considering photoionization and CIE must be tabulated as a function of density, temperature, redshift, and composition (e.g., Smith et al. 2008). The contribution of each species is additive, enabling element-by-element tabulation and affording the flexibility to model cooling in gases with abundance ratios that deviate from those of the Solar Neighborhood. Wiersma et al. (2009a) provide such a tabulation, for the set of 11 most-important radiative coolants (H, He, C, N, O, Ne, Mg, Si, S, Ca, and Fe) exposed to the Haardt & Madau (2001) UVB model. These tables are used by the **EAGLE**, **Illustris**/**IllustrisTNG**, **FABLE**, and **Magneticum** simulation suites, though only **EAGLE** and **Magneticum** exploit the element-by-element nature of the tables; the other suites scale the cooling rate as a function of metallicity, assuming Solar Neighborhood abundance ratios.

As the column density of cosmic gas approaches a value characteristic of Lyman-limit systems ($N_{\text{H I}} \gtrsim 10^{17} \text{ cm}^{-2}$), it becomes optically thick, attenuating the incident radiation flux and increasing the cooling rate relative to the optically thin case. The corresponding volume density of gas for which self-shielding is important is $n_{\text{H}} \gtrsim 10^{-2} \text{ cm}^{-3}$ (e.g., Furlanetto et al. 2005), which is characteristic of the interstellar medium (ISM). The reduction of the UVB photoionization rate in self-shielded gas can be approximated using fitting functions calibrated against full radiation transport calculations (e.g., Gnedin et al. 2009, Rahmati et al. 2013). This approximation is applied on the fly (for $z < 6$) by the **Illustris** and **IllustrisTNG** simulations. The self-shielding approximation by Rahmati et al. (2013) is also used by the **Grackle** library (last described by Smith et al. 2017), which computes cooling rates for interstellar gas using a primordial chemistry reaction network, in tandem with tabulated rates for metal species in the optically thin regime. The

primordial network also accounts for deviations from ionization equilibrium. This library is used by the SIMBA simulations.

Local radiation fields in the ISM can dominate over the UVB, which further complicates the calculation of cooling rates (e.g., Schirber et al. 2004, Miralda-Escudé 2005). Molecules and dust grains within the ISM attenuate the incident radiation spectrum (in dissimilar ways), and also contribute to the cooling rate (e.g., Omukai et al. 2005). Ploeckinger & Schaye (2020) provide cooling tables for gas, shielded by itself and dust grains, in ionization equilibrium with radiation from both a spatially uniform, temporally evolving UVB and an interstellar radiation field of variable intensity, but with spectral shape fixed to that of the Milky Way (Black 1987).

2.3.2. Element abundance evolution. Besides influencing radiative cooling rates, element abundance patterns in cosmic gas and stars encode a wealth of astrophysical information. Elements heavier than lithium (metals) were produced by nucleosynthetic processes within stars and supernovae (SNe; and in some cases, kilonovae), with some fraction of the synthesized elements returned to the local ISM as ejecta. The enriched gas can then be incorporated into future generations of stars or transported into the CGM or IGM by galaxy-scale outflows. When galaxies fall into more massive halos, such as groups and clusters of galaxies, their enriched ISM can be stripped by the ram pressure it experiences from the CGM or intracluster medium (ICM).

Modern schemes for distributing and transporting the elements synthesized by stellar populations (chemodynamics) remain similar to the first generation of implementations (e.g., Theis et al. 1992, Steinmetz & Mueller 1994, Carraro et al. 1998). Stellar particles, each assumed to represent a mono-age simple stellar population (SSP), donate enriched mass to neighboring fluid tracers. Elements then share the destiny of the tracer to which they are donated, enabling their transport and dispersal by gas flows, and incorporation into subsequent generations of stars, in a fashion consistent with the adopted hydrodynamics scheme. Early schemes focused on enrichment by type II supernovae (SNeII), whose progenitors have short lifetimes ($\lesssim 30$ Myr, which is comparable with the sound-crossing time of fluid elements representing the diffuse ISM) and whose nucleosynthetic products can be distributed immediately upon the formation of the SSP (the instantaneous recycling approximation). Although the total metallicity is generally dominated by elements synthesized by SNeII, modeling the abundance of elements such as iron, carbon, and nitrogen requires consideration of element release by type Ia supernovae (SNeIa) and long-lived stars that experience an asymptotic giant branch (AGB) phase. Modern simulations therefore generally adopt subgrid models that follow the timed release of individual elements from multiple nucleosynthetic channels, e.g., Oppenheimer & Davé (2008), used by SIMBA; Tornatore et al. (2007), used by *Magneticum*; and Wiersma et al. (2009b), used by EAGLE and *Illustris*/*IllustrisTNG*.

The progenitors of AGB stars and SNeII are, respectively, intermediate-mass ($0.8 \lesssim M \lesssim 8 M_{\odot}$) and high-mass ($M \gtrsim 8 M_{\odot}$) stars. An initial mass function (IMF) must therefore be specified in order to set the relative number of these progenitors (it is also required to set the energetics of stellar feedback; see Section 2.3.5). It is typically assumed to follow the form deduced from Solar Neighborhood number counts (e.g., Chabrier 2003), with the lower mass cutoff being the hydrogen mass burning limit ($\simeq 0.07\text{--}0.09 M_{\odot}$) and the upper mass cutoff reflecting the (more uncertain) maximum observed mass of stars. Typically limits of $0.1 M_{\odot}$ and $100 M_{\odot}$ are adopted. Uncertainties in the form and limits of the IMF translate directly into uncertainty on the cosmic metal budget. Mass loss during the AGB phase occurs in a brief phase at the end of the main sequence lifetime. Lifetimes are challenging to constrain observationally and are generally inferred from stellar evolution models (e.g., Romano et al. 2005). The models concur that lifetimes are a strongly decreasing function of stellar mass, with some inferring a weak dependence on metallicity. The intermediate-mass progenitors of AGB stars have lifetimes of $10^8\text{--}10^{10}$ years, which are

SN(e): supernova(e)

ICM: intracluster medium

SSP: simple stellar population

SN(e)II: type II supernova(e)

SN(e)Ia: type Ia supernova(e)

AGB: asymptotic giant branch

long compared to the sound crossing time of ISM fluid elements (and, hence, the time steps on which dense gas is advanced), but are in general shorter than, or comparable with, the dynamical timescales of galaxies. The mass dependence of stellar lifetimes therefore has a tangible influence on element release.

SNeIa are thermonuclear explosions resulting from binary star evolution. Their progenitors remain ill constrained, with the most plausible cases being the accretion of mass onto a Chandrasekhar-mass white dwarf from a nondegenerate companion or the merger of two white dwarfs (e.g., Maoz et al. 2014). The SNIa rate resulting from either scenario is thus complicated by the ill-constrained properties of binaries, e.g., their mass fraction, mass function, and initial separation; however, the involvement of at least one white dwarf dictates that the binary progenitor is typically old ($\gtrsim 10^9$ years). Forward modeling the SNIa rate as a function of an SSP’s age is therefore subject to many uncertainties (e.g., Greggio & Renzini 1983, Greggio 2005), fostering the alternative approach of assuming an empirical delay function for the SSP’s SNIa rate, which integrates to unity at $t = \infty$ and whose free parameters can be calibrated against cosmic SNIa observations (see, e.g., Mannucci et al. 2006, Wiersma et al. 2009b).

The nucleosynthetic yields of SNeII, SNeIa, and AGB stars represent a significant systematic uncertainty for chemodynamical modeling. The uncertainties stem from the complexity of the physics involved, such as rotation in SNIi progenitors, or the efficiency of convective envelope burning in AGB stars, the postexplosion time at which the yields are quoted (synthesized isotopes can have radioactive decay timescales shorter than the expansion timescale of the ejecta), or inconsistencies in the assumed mass ranges of progenitors. Nucleosynthetic calculations remain challenging to reconcile with the observed abundances of stars in the Galaxy, often motivating ad hoc and element-specific rescaling factors (at the factor of $\simeq 2$ level) to the resulting theoretical yields (e.g., François et al. 2004, Portinari et al. 2004). As such, the absolute element abundances predicted by simulations must be considered uncertain by at least a similar factor.

Mixing and diffusion are further significant sources of uncertainty for chemodynamical evolution. Some degree of overmixing is generally expected in mesh-based simulations, because fluids are implicitly mixed on scales smaller than that of the smallest cells, but excellent numerical convergence behavior can be demonstrated with moving mesh treatments of hydrodynamics (van de Voort et al. 2020). Conversely, SPH (particularly traditional schemes) tends to underestimate mixing because of its inability to resolve dynamical instabilities (e.g., Agertz et al. 2007). In the absence of an explicit diffusion treatment, metal mixing in SPH simulations is further underestimated because metals are “stuck” to particles, resulting in poor sampling of fluids with a small but nonzero metallicity. The problem can be mitigated by the inclusion of a subgrid diffusion treatment (e.g., Wadsley et al. 2008, Greif et al. 2009), but the appropriate coefficients are unclear (e.g., Garnier et al. 2009). Wiersma et al. (2009b) argue that a simpler strategy to mitigate the metal sampling problem in SPH simulations is to use kernel-smoothed abundances. Irrespective of the hydro solver, undermixing can be exacerbated if, as is commonplace, sources of small-scale turbulence are neglected.

2.3.3. Subgrid models of the interstellar medium and star formation. Perturbations in a self-gravitating medium can grow only if their wavelength exceeds the Jeans (1928) length. Accurate numerical modeling of gravitationally collapsing systems therefore requires that the Jeans scales are adequately resolved: in hydrodynamical settings, this translates into the need to resolve pressure forces on the scale at which self-gravity begins to dominate. In Lagrangian simulations the condition is most naturally expressed in terms of the ratio of the particle mass and the Jeans mass (Bate & Burkert 1997), and in Eulerian counterparts in terms of the ratio of the cell size and the Jeans length (Truelove et al. 1997).

The thermogravitational collapse of the warm, diffuse phase of the ISM ($T \sim 10^4$ K, $n_{\text{H}} \sim 10^{-1} \text{ cm}^{-3}$) to the cold, dense phase ($T \lesssim 10^2$ K, $n \gtrsim 10^2 \text{ cm}^{-3}$) occurs quickly (McKee & Ostriker 1977). Simulations that seek to resolve the molecular ISM are therefore faced with demanding resolution requirements. The Jeans scales depend on both the density and temperature of the gas ($M_{\text{J}} \propto n_{\text{H}}^{-1/2} T^{3/2}$, $L_{\text{J}} \propto n_{\text{H}}^{-1/2} T^{1/2}$), so an increase of (at least) three decades in density and a decrease of (at least) two decades in temperature corresponds to reductions in the Jeans mass from $M_{\text{J}} \sim 10^7 M_{\odot}$ to $M_{\text{J}} \sim 10^2 M_{\odot}$, and the Jeans length from $L_{\text{J}} \sim 1 \text{ kpc}$ to $L_{\text{J}} \sim 1 \text{ pc}$. As noted earlier, appealing to such high resolution in representative cosmological volumes requires the use of very short timesteps and incurs a very large memory footprint, necessitating execution with large core counts that cannot be effectively exploited by the current generation of simulation codes. The pressure structure of the multiphase ISM therefore cannot be accurately resolved, typically requiring that a subgrid model is either used to predict the pressure structure as a function of the gas density (e.g., Yepes et al. 1997, Springel & Hernquist 2003), or to impose the structure “by hand” as a pressure floor (e.g., Machacek et al. 2001, Robertson & Kravtsov 2008) or as an equation of state (e.g., Schaye & Dalla Vecchia 2008).

In the absence of a numerical treatment of the cold ISM, star formation must be implemented with models whose efficiency is calibrated to reproduce observed scaling relations averaged over suitably large spatial scales ($\sim 1 \text{ kpc}$). The most commonly adopted form is that of a Schmidt (1959) law, for which the SFR density ($\dot{\rho}_{\star}$) scales linearly with gas density (ρ_{g}) over a dynamical time (t_{dyn}): $\dot{\rho}_{\star} \propto \epsilon_{\star} \rho_{\text{g}} / t_{\text{dyn}}$. Here, ϵ_{\star} is the efficiency per dynamical time, which is calibrated so that, operating in tandem with the density threshold above which star formation is allowed to proceed, the simulation reproduces the observed Kennicutt–Schmidt (KS) star-formation law (Kennicutt 1998). Because the latter is based on surface densities rather than volume densities, the calibration depends on the assumed equation of state. Schaye & Dalla Vecchia (2008) therefore introduced an alternative scheme that, under the assumption of vertical hydrostatic equilibrium in the ISM, enables the KS law to be expressed as a pressure law, eliminating the dependence of the star formation law on the equation of state and replacing the calibrated parameters with observables. Alternatively, Feldmann et al. (2023) presented results from the FIREbox simulation of a (small, $L = 22.1 \text{ cMpc}$) cosmological volume at $z = 0$, whose star-formation model (Hopkins et al. 2018) does not impose an efficiency per free-fall time, instead aiming to recover the observed effective efficiency via the influence of the model’s implemented ISM and feedback physics.

2.3.4. Supermassive black hole seeding and growth. SMBHs are represented by collisionless sink particles, subject only to gravity. They are usually seeded at the centers of halos (which do not already have a seed) identified on the fly by periodically applying a group-finder algorithm to the DM particle distribution (e.g., Di Matteo et al. 2008) or in high-density gas whose properties satisfy a number of criteria (e.g., Dubois et al. 2012). There is no consensus concerning the formation mechanism of SMBHs (see, e.g., Volonteri et al. 2021), but popular theories posit that their formation mass is below the resolution limit of simulations of the galaxy population, requiring that sink particles carry a subgrid SMBH mass (used by subgrid routines concerning SMBH evolution) in addition to their dynamical mass. Once the former exceeds the latter, the subgrid and dynamical masses grow in concert. Sink particles with masses less than or comparable with the mass of neighboring resolution elements do not experience a realistic dynamical friction force, so it is common to “pin” them to the halo’s center of potential or migrate them toward it. This practice has been shown to have significant consequences for the evolution of simulated galaxies (Bahé et al. 2022).

SMBHs grow via gas accretion and mergers with other SMBHs. The ambient gas accretion rate, \dot{m}_{accr} , is generally assumed to be at the minimum of the Eddington and Bondi–Hoyle–Lyttleton rates (e.g., Bondi & Hoyle 1944), or a modified version of the latter, though recent

models also explicitly consider the influence of the angular momentum of the gas (e.g., Anglés-Alcázar et al. 2017). The Bondi rate is estimated from the sink particle’s ambient gas properties, necessarily on spatial scales that are much larger than accretion disks, such that the simulations treat the physics of accretion onto SMBHs in a very simplistic fashion (see, e.g., Shlosman et al. 1990). The SMBH grows at a rate $\dot{m}_{\text{BH}} = \dot{m}_{\text{accr}}(1 - \epsilon_r)$, where ϵ_r is the SMBH’s radiative efficiency, which is generally assumed to follow the mean value for radiatively efficient accretion onto a Schwarzschild black hole, $\epsilon_r = 0.1$ (Shakura & Sunyaev 1973). Older simulations based on the popular model of Springel et al. (2005) boosted the Bondi rate (usually by a factor of $\alpha = 100$) to compensate the underestimate of the gas density near the Bondi radius, though modern simulations tend to apply either a density-dependent correction as advocated by Booth & Schaye (2009) or no correction because it is generally only important early in the SMBH’s growth history, making it degenerate with the choice of seed mass. A more sophisticated approach possible in mesh-based simulations is to refine the resolution of the simulation in the vicinity of SMBHs to ensure the Bondi radius is resolved (Curtis & Sijacki 2015).

2.3.5. Feedback processes. The ratio of galaxies’ stellar mass to halo mass, as inferred for example by subhalo abundance matching (e.g., Conroy et al. 2006), exhibits a characteristic peak at a halo mass of $\sim 10^{12} M_\odot$. The efficiency with which halos convert their cosmic “share” of gas into stars decreases toward lesser and greater halo mass scales. This is widely interpreted as a signal that galaxy growth is primarily regulated (or even quenched) by separate feedback mechanisms in these two regimes: the formation and evolution of stellar populations in low-mass galaxies, and accretion onto SMBHs in massive galaxies. The injection of energy by these mechanisms occurs on numerically unresolved scales, so simulations must approximate their macroscopic effects with subgrid models.

Stellar populations inject energy and momentum into the ISM via stellar winds, radiation, and SNe and, thus, potentially disrupt star-forming molecular clouds, create turbulence, and drive gas out of galaxy disks. The simplest subgrid treatment sums the energy liberated at each time step by a stellar particle (representing an SSP) and injects it thermally by raising the internal energy of the particle’s neighboring fluid elements. However, because these elements are orders of magnitude more massive than SN ejecta, the injected energy (canonically $E_{\text{SN}} \sim 10^{51}$ erg/SN) is only sufficient to heat individual fluid elements to $\sim 10^5$ K. In this regime the gas cooling time is short compared to its sound crossing time, leading to catastrophic artificial losses that preclude the formation of an adiabatic, energy-conserving blast wave (e.g., Katz et al. 1996, Dalla Vecchia & Schaye 2008, Creasey et al. 2011).

Popular solutions are to temporarily disable the cooling of gas, either by hand (e.g., Gerritsen & Icke 1997, Stinson et al. 2006) or by using an additional subgrid model to account for unresolved ISM physics (Keller et al. 2014); to heat gas particles stochastically by temperature increments $\Delta T \gg 10^5$ K (Kay et al. 2003, Dalla Vecchia & Schaye 2012); or to inject (some of) the energy in kinetic form (e.g., Navarro & White 1993, Springel & Hernquist 2003, Dubois & Teyssier 2008). Each approach has pros and cons but, for coupling efficiencies of order unity, they all in principle enable (low-mass) galaxies to drive outflows and achieve self-regulated growth. The kinetic approach affords the freedom to specify explicitly both the initial velocity and mass loading ($\eta = \dot{M}_{\text{outflow}}/\dot{M}_*$) of the wind, enabling calibration of the wind model against the outcomes of higher-resolution zoom simulations with more sophisticated ISM models, an approach adopted by SIMBA. The popular Springel & Hernquist (2003) implementation, which has been used in adapted form by many simulation suites including *Illustris* and *IllustrisTNG*, temporarily decouples “kicked” particles from the hydrodynamics scheme, in principle aiding numerical

convergence (see Section 2.5) but precluding interactions between winds and the ISM (Dalla Vecchia & Schaye 2008).

The physical mechanism by which AGN feedback couples to the ISM remains poorly understood, with a number of channels being plausible, such as radiation pressure on free electrons and/or dust grains or very high-velocity jets. Despite this absence of consensus, there is ample observational evidence to indicate that AGNs drive large-scale, high-velocity outflows of ionized and molecular gas (e.g., Maiolino et al. 2012, Ciccone et al. 2014). Subgrid AGN feedback models generally assume that some fraction, ϵ_f , of the radiated luminosity of an SMBH accretion disk couples to the surrounding ISM: The AGN feedback energy $E_{\text{AGN}} = \epsilon_f \epsilon_r \dot{M}_{\text{BH}} c^2$, where \dot{M}_{BH} is the accretion rate and c is the speed of light (e.g., Springel et al. 2005). As with stellar feedback, the injected energy can be coupled to the numerically modeled gas fluid in thermal and/or kinetic form. The simplest approach assumes that the coupling efficiency is fixed, such that the energy injection rate is proportional to the accretion rate. However, motivated by the observation that high-velocity jets are more typically associated with accretion onto SMBHs at small fractions of the Eddington rate (Churazov et al. 2005, Merloni & Heinz 2008, Best & Heckman 2012), some models adopt a relatively low (high) value of ϵ_f when the SMBH is accreting at a high (low) fraction of the Eddington rate (e.g., Sijacki et al. 2007, Dubois et al. 2012, Weinberger et al. 2017, Davé et al. 2019). Because the Eddington rate depends linearly on the mass of the SMBH, this approach can act like a switch, whereby efficient feedback is triggered once the SMBH reaches a threshold mass.

2.4. Calibration of Subgrid Feedback Models

The numerical values of the parameters governing subgrid models can, in some cases, be chosen by reference to observables or theoretical arguments. In general though, the appropriate values are not known a priori and may be resolution dependent. The efficiencies of feedback processes are the most salient example: The microphysics are not well understood theoretically, and observations do not authoritatively characterize outflow properties on the scales at which subgrid models recouple to the hydrodynamics scheme (see, e.g., Chisholm et al. 2016a,b). Furthermore, changing the resolution of a simulation often changes the frequency of individual feedback events, and the energy, mass, and momentum they each inject. Even if over some timescale the same total energy (for example) is injected, the macroscopic efficiency of feedback can depend on this intermittency because the numerical losses differ.

It is therefore increasingly common practice to calibrate the parameters of subgrid feedback models, so that simulations broadly reproduce key properties of the galaxy population. Depending on the sensitivity of the subgrid models to resolution, the parameters may require recalibration if the resolution is changed. The choice of calibration diagnostics is somewhat arbitrary, but clearly they should be well-characterized by observations on scales resolved by the simulations and should be sensitive to the parameter(s) requiring calibration. Most often, the stellar mass of galaxies is used (via either the observed GSMF or the inferred stellar mass-to-halo mass ratio) to calibrate the efficiency of feedback associated with stellar evolution, and the mass of central SMBHs at fixed galaxy stellar mass to calibrate the efficiency of feedback associated with their growth. Reproduction of these quantities alone does not guarantee a realistic galaxy population, so complementary observables may also be used. For example, by also considering the sizes of disk galaxies, the EAGLE simulations more accurately reproduced many important galaxy scaling relations (Crain et al. 2015). IllustrisTNG further considers the cosmic SFR density and halo gas fractions. SIMBA's feedback efficiencies were tuned only against the GSMF and the $M_{\text{BH}}-M_*$ relation, the latter via the accretion efficiency rather than the AGN feedback efficiency. Horizon-AGN's stellar feedback efficiency was not explicitly calibrated but inferred from the Starburst99 (Leitherer

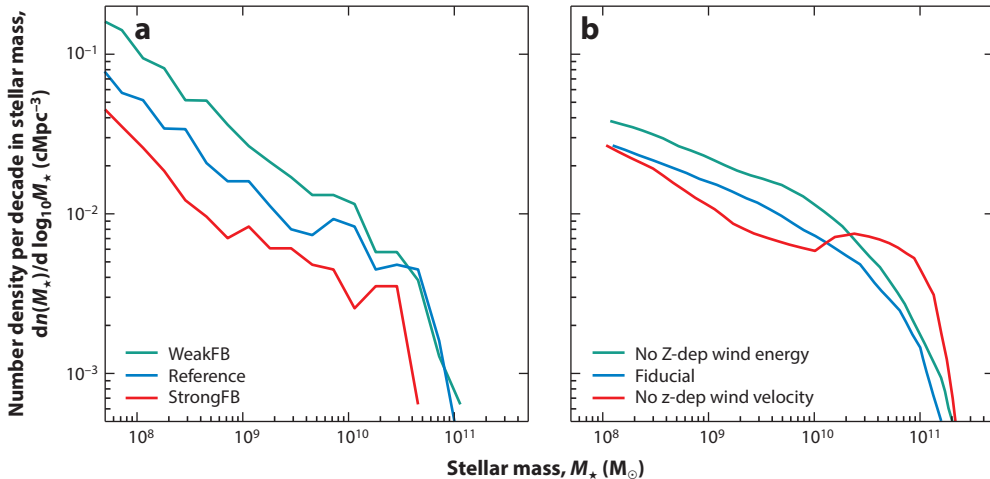


Figure 3

An illustration of the influence on the GSMF of adjusting the subgrid stellar feedback model in calibrated simulations. (a) The response of the present-day GSMF of EAGLE, realized in an $L = 25$ cMpc volume, to changing the adopted stellar feedback efficiency (which scales linearly with the injected energy) by factors of 0.5 (WeakFB, green curve) and 2.0 (StrongFB, red curve). (b) The response of the present-day GSMF of IllustrisTNG, realized in an $L = 37$ cMpc volume, to disabling the metallicity dependence of the injected energy (No Z-dep wind energy, green curve) and to disabling the redshift dependence of the wind injection velocity (No z-dep wind velocity, red curve). Panel a adapted from Crain et al. (2015); panel b adapted from Pillepich et al. (2018b). Abbreviations: GSMF, galaxy stellar mass function; StrongFB, strong feedback; WeakFB, weak feedback.

et al. 1999) spectrophotometric model. Its AGN feedback model was, in common with EAGLE and IllustrisTNG, calibrated with reference to the $M_{\text{BH}}-M_*$ relation. As noted in Section 1, care must be taken not to interpret, as predictions, emergent properties that are closely related to the calibration diagnostics nor those properties significantly compromised by the simplifications of the subgrid modeling.

Figure 3 shows the influence on the GSMF of adjusting the subgrid stellar feedback model in calibrated simulations. **Figure 3a** shows the influence of using stellar feedback efficiencies (effectively a coefficient specifying the fraction of E_{SN} that couples to the ISM) for EAGLE’s stochastic thermal heating treatment, which differ from the reference model by factors of 0.5 (weak feedback, WeakFB) and 2 (strong feedback, StrongFB). The shifts follow from galaxies of a fixed stellar mass becoming associated with less (more) massive DM halos in the WeakFB (StrongFB) model. Because the DM halo mass function is steep, even a small change to the stellar mass–halo mass relation significantly alters the space density of galaxies at fixed stellar mass. **Figure 3b** shows the influence in IllustrisTNG of disabling the scaling of the energy of its kinetically driven stellar winds as a function of metallicity (No Z-dep wind energy), and separately, the scaling of the velocity of these winds as a function of redshift (No z-dep wind velocity). The energy scaling (effectively a mass loading scaling at fixed velocity) suppresses the growth of low-mass galaxies and acts in a fashion similar to that of the EAGLE feedback efficiency scaling in **Figure 3a**. The velocity scaling prevents kinetically driven winds from stalling in relatively massive galaxies (see also Crain et al. 2009, Oppenheimer et al. 2010).

In analogy to calibrating the stellar feedback efficiency, a similar procedure is often used to calibrate the efficiency of AGN feedback, whereby the value of ϵ_f is adjusted to achieve broad reproduction of the observed relation between the stellar mass of galaxies and the mass of their central SMBH. Booth & Schaye (2010) show that so long as AGN feedback is numerically efficient

and well sampled, stellar masses are insensitive to extreme variations of ϵ_f , as SMBHs compensate by growing to the mass that enables them to inject the feedback energy needed to (self-)regulate gas inflow. As such, ϵ_f primarily sets the normalization of the SMBH–galaxy mass scaling relation.

2.5. Verification and Convergence

Numerical simulations, irrespective of their purpose or application, are simplified approximations of real phenomena or systems. It is therefore necessary to verify their outcomes to ensure they are fit for purpose. For many terrestrial applications, simulations can be validated against experimental data (e.g., computational fluid dynamics simulations are often confronted with wind tunnel measurements). Clearly, cosmological simulations can appeal to no such testbed. The performance of hydrodynamics solvers can be examined using a set of idealized tests (e.g., shock tubes, point explosions, vortex problems, fluid instability tests), but being a strongly nonlinear process resulting from the interplay of complex physical processes, galaxy formation offers no simple tests with analytic or known solutions that can be used for validation. A laborious but illuminating strategy is therefore to apply multiple models to the same problem, enabling an assessment of the consensus between the models (e.g., Scannapieco et al. 2012, Kim et al. 2014, Cui et al. 2018).

Predictive power demands that the outcomes of simulations are robust to changes of various aspects, such as changes in resolution and the size of the simulation domain, a characteristic often referred to as convergence. Unhelpfully, the physics of galaxy formation does not readily lend itself to achieving converged results. As noted in Section 2.1, changes to the size of the simulation volume influence the diversity of the environments and halo population that can be realized by the simulations and the mass and spatial scales for which cosmic variance becomes important. They can also impact volume-integrated properties, such as the global SFR density.

Resolution convergence is particularly challenging to achieve, owing to the dominant role played by physical processes taking place at, or below, the numerical resolution scale (see **Figure 2**). In simple terms, a higher-resolution simulation can resolve smaller galaxy/halo progenitors, enabling them to be followed at earlier cosmic epochs. They can also resolve higher gas densities, influencing (net) radiative cooling rates and, in some cases, interfacing with subgrid models at a different spatial scale. As remarked in Section 2.3.5, subgrid models can be designed to mitigate this resolution sensitivity, for instance, by temporarily decoupling gas from hydrodynamical forces or radiative cooling and/or by using the generally better-converged properties of the local DM distribution, rather than those of the gas, as inputs. However, these choices introduce their own drawbacks, moving the modeling philosophy closer to the phenomenological approach of semi-analytic modeling. As noted in Section 2.4, changing the resolution can also incur more subtle consequences, such as altering the intermittency and numerical efficiency of feedback events.

It is therefore good practice to compare the properties of galaxy populations that emerge from simulations when the volume (or box size) and resolution are varied (ideally individually). Such tests are straightforward and inexpensive to conduct when using smaller volumes and lower resolutions. To mitigate the computational cost of testing for convergence at higher resolution, it is common to conduct high-resolution simulations in a smaller simulation volume and to run a partner simulation of the same volume at the fiducial resolution to control for box size effects. With such a suite of simulations, it is possible to make authoritative assessments of the degree to which the properties of the galaxy population are robust to volume and resolution changes. However, it is important to remain vigilant to the possibility that the appearance of a result being converged may be contingent on the particular implementation of one or more subgrid models.

Many subgrid models employ stochastic treatments of continuous processes, such as the conversion of gas particles or cells into stellar particles, or injecting energy kinetically and/or thermally into the gas through feedback processes. A number of recent studies have examined the

impact of this subgrid stochasticity on emergent macroscopic properties, in simulations of individual galaxies (Keller et al. 2019; Davies et al. 2021, 2022) and periodic volumes (Genel et al. 2019, Borrow et al. 2022b), finding that small differences due to stochasticity can propagate into significant differences in present-day properties, even for galaxies resolved with many thousands of particles. The systematic uncertainty due to stochasticity is more significant in smaller samples of galaxies, i.e., zooms and small periodic volumes, and needs to be considered when interpreting results, including when analyzing small simulations for the purposes of model calibration. In such cases, the uncertainty can be mitigated with brute force, by running multiple realizations of the same simulation using different random number seeds.

3. KEY PROPERTIES OF SIMULATED GALAXY POPULATIONS

In this section, we examine the degree to which state-of-the-art hydrodynamical simulations reproduce key properties of the galaxy population. Clearly the definition of ‘key’ is somewhat arbitrary, and opinions will differ between practitioners, but we focus here on diagnostics that are sufficiently important to be widely considered as validation tests. Marked failure of a simulation to reproduce these diagnostics will undermine confidence in the broader conclusions one may draw from it. In some cases, these properties are used as calibration diagnostics, so they should not be considered predictions.

3.1. Galaxy Stellar Mass Function

The $z \simeq 0$ GSMF is the observational diagnostic most frequently used to calibrate the parameters of subgrid feedback models. Simulations of $L \sim 100$ cMpc and baryonic mass resolution $m_g \sim 10^6 M_\odot$ adequately resolve and sample the GSMF over the range $8 \lesssim \log_{10}(M_\star/M_\odot) \lesssim 11$, thus demanding realistic stellar masses for galaxies with space densities that differ by at least two decades. The present-day GSMFs of EAGLE, Horizon-AGN, IllustrisTNG, and SIMBA are shown in **Figure 4** together with measurements from the Sloan Digital Sky Survey (SDSS) and Galaxy and Mass Assembly (GAMA) surveys. This highlights that it has proven possible to reproduce the GSMF over much of this stellar mass range in simulations of the galaxy population, in some cases with an accuracy better than or comparable with the ~ 0.3 dex systematic uncertainty on stellar masses inferred using population synthesis models (e.g., Conroy et al. 2009). Although the EAGLE, IllustrisTNG, and SIMBA simulations were calibrated to achieve this reproduction, success was not guaranteed because their subgrid models afford only limited freedom. For comparison, the figure also shows the GSMF of the flagship $L = 22.1$ cMpc FIREbox simulation, in which the parameters governing stellar feedback processes (this simulation does not model AGN feedback) were not calibrated to this observable. FIREbox and Horizon-AGN overproduce the space density of galaxies, especially for those with stellar mass of $9 \lesssim \log_{10}(M_\star/M_\odot) \lesssim 10$.

A common aspect of models with realistic GSMFs is the operation of efficient feedback at all mass scales. Simulations that do not reproduce the observed GSMF commonly (but not exclusively; e.g., Pakmor et al. 2022) form too many galaxies at fixed stellar mass. As noted in Section 2.4, this is a consequence of feedback failing to adequately regulate galaxy growth and thus allowing galaxies to form in low-mass DM halos with too high space density. Failure to regulate galaxy growth can be due to the implemented feedback being unintentionally inefficient for numerical reasons (e.g., poor sampling of energy injection events or injection of too little energy per feedback event), because a feedback mechanism (e.g., AGN) is not included or because a comprehensive search of the model’s plausible many-dimensional parameter space is too computationally expensive. The latter is a major barrier to exhaustive calibration and has motivated the adoption of Gaussian process emulation to accelerate the procedure (Kugel & Borrow 2022).

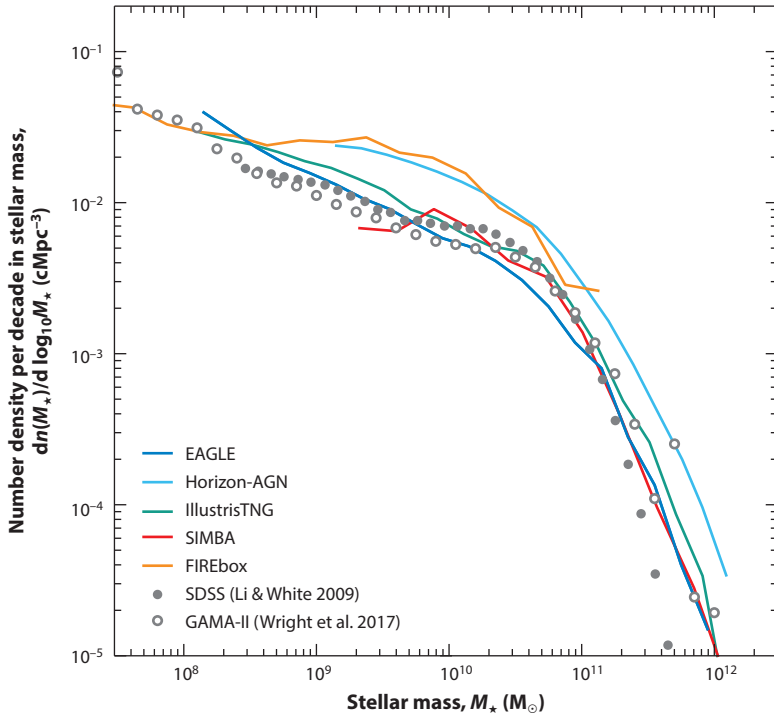


Figure 4

The present-day GSMF of the four simulations listed in **Table 1**: EAGLE (data from Schaye et al. 2015), Horizon-AGN (data from Kaviraj et al. 2017), IllustrisTNG (data from Pakmor et al. 2022), SIMBA (data from Davé et al. 2019). We also show the GSMF of the FIREbox simulation (data from Feldmann et al. 2023), whose resolution is higher, and simulation volume smaller, than the other simulations. Filled and open circles show GSMF measurements from the SDSS and GAMA surveys (Li & White 2009, Wright et al. 2017); for clarity we omit their error bars. EAGLE, IllustrisTNG, and SIMBA recover the observed space density of galaxies at fixed mass (over the mass range for which the observational measurements concur) to within 0.3 dex. Due to overcooling, Horizon-AGN and FIREbox yield a space density for galaxies of 10^9 – $10^{10} M_\odot$ that is too high by up to 0.5 dex. Abbreviation: GSMF, galaxy stellar mass function.

Reproduction of the GSMF at $z \simeq 0$ alone is insufficient to ensure that simulations also recover its evolution with cosmic time, because plausible feedback models can yield unrealistic star-formation histories (SFHs; Crain et al. 2015). However, the use of complementary calibration diagnostics (e.g., galaxy sizes or the cosmic SFR density) enables SFHs to be constrained sufficiently to reproduce GSMFs as early as $z \simeq 7$ to an accuracy consistent with the systematic uncertainties on observationally inferred masses (e.g., Furlong et al. 2015, Pillepich et al. 2018a, Davé et al. 2019). Simulations of $L \sim 100$ cMpc are ill suited for confrontation with the demographics of very high-redshift galaxies derived from early *James Webb Space Telescope* (JWST) imaging (e.g., Labbé et al. 2023), as they do not sample the rare fluctuations in the density field that likely seed the formation of the bright, rare sources preferentially detected by these observations (Kannan et al. 2022).

3.2. Size and Morphology

Reproduction of galaxies with realistic sizes has long been viewed as a prominent challenge for cosmological simulations, owing primarily to the recognition that overcooling can lead to spurious

SFH: star-formation history

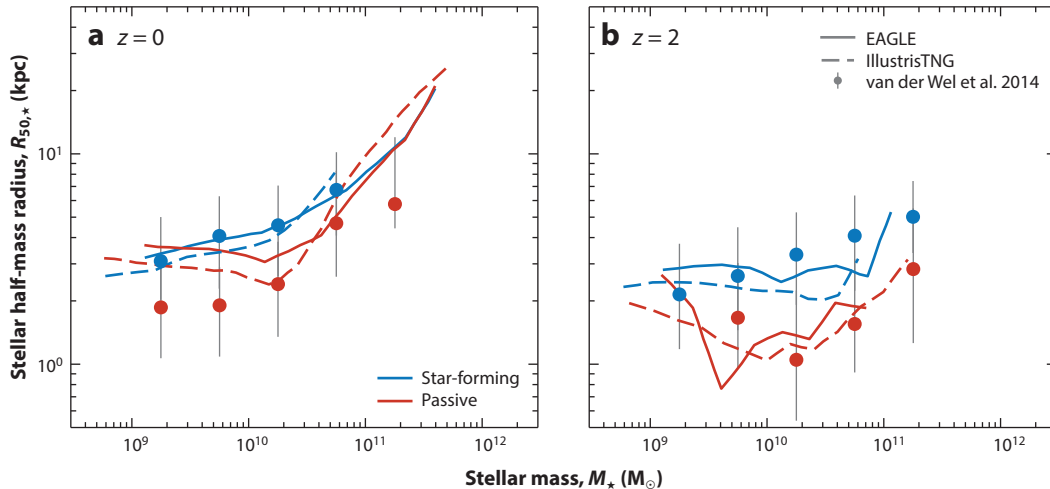


Figure 5

Present-day galaxy sizes as a function of stellar mass for star-forming (*blue curves*) and passive (*red curves*) galaxies in EAGLE (*solid curves*; data from Furlong et al. 2017) and IllustrisTNG (*dashed curves*; data from Genel et al. 2018), at (a) $z = 0$ and (b) $z = 2$. Filled symbols with error bars denote the observational size measurements of van der Wel et al. (2014). Both simulations broadly reproduce the observed size–mass relations of both galaxy types and their evolution over $\simeq 10$ Gyr, though at late times both simulations overestimate the size of low-mass passive galaxies.

angular momentum transfer from cold gas and stars to DM (Katz & Gunn 1991) and, thus, to the formation of galaxies that are too compact as well as too massive. Clearly, ensuring that galaxies form with broadly realistic stellar masses is then a necessary condition for reproduction of realistic sizes.

However, the original Illustris simulation yielded present-day galaxies too large by a factor of $\simeq 2$ (Snyder et al. 2015), highlighting that realistic masses are not a sufficient condition for realistic sizes. Using EAGLE, Crain et al. (2015) demonstrated that the overcooling and angular momentum problems are somewhat separable: Not only must simulated galaxies form a realistic mass of stars but their SFHs must also be broadly realistic, so that stars form from natal gas with the correct angular momentum distribution. For this reason, the sizes of present-day disk galaxies were explicitly considered when calibrating the parameters of EAGLE’s subgrid stellar feedback model. Sizes were also considered by IllustrisTNG during the development of their model of stellar winds.

Figure 5 shows the galaxy size–stellar mass relations of, separately, star-forming and passive galaxies in EAGLE (data from Furlong et al. 2017) and IllustrisTNG (data from Genel et al. 2018), at $z = 0$ (**Figure 5a**) and $z = 2$ (**Figure 5b**), respectively. In both simulations, galaxy size is defined as the stellar half-mass radius, R_{50} . Filled symbols with error bars are the observational size measurements of galaxies at the corresponding redshifts, from van der Wel et al. (2014). EAGLE and IllustrisTNG broadly reproduce the observed size–mass relations of star-forming and passive galaxies, and the evolution of these relations from smaller sizes at fixed stellar mass, from $z = 2$. The simulations differ from the observations by $\simeq 0.1$ dex at low redshift and $\simeq 0.3$ dex at early times and, thus, compare much more favorably with the observations than was the case for prior simulations (e.g., McCarthy et al. 2012).

Interestingly, both Horizon-AGN (see Dubois et al. 2016) and SIMBA (Davé et al. 2019) exhibit size–mass relations for which, at fixed mass, passive galaxies are larger than star-forming

counterparts. In both cases, the relation for star-forming galaxies is in broad agreement with observational measurements, indicating that their passive galaxies are too large. Although in *EAGLE* and *IllustrisTNG* star-forming galaxies tend to be larger than passive counterparts, as is observed, the sizes of both galaxy types are similar at low mass (likely due to the poor sampling of SFHs), indicating that passive galaxies in these simulations are also too extended. This may be a consequence of spurious and prolonged dynamical heating of their stellar particles by more massive DM particles (Ludlow et al. 2019); if so, this problem can be mitigated via the adoption of baryonic and DM particles of similar mass (e.g., as per Tremmel et al. 2015), at the expense of a greater computational cost and memory footprint.

An outcome seen in both *EAGLE* and *IllustrisTNG* is a correlation, at fixed stellar mass, between the sizes and the specific SFRs of galaxies. Interestingly, in *IllustrisTNG* the trend is not seen when one considers only star-forming galaxies, and the overall trend reflects the changing relative fractions of (extended) star-forming and (compact) passive galaxies, from low to high stellar mass. In contrast, massive ($M_\star \gtrsim 10^{9.5} M_\odot$) star-forming galaxies in *EAGLE* exhibit a clear correlation between size and specific SFR.

A related issue is whether simulated galaxy populations reproduce the diversity of observed morphologies (and kinematics). *EAGLE*, *Horizon-AGN*, and *IllustrisTNG* reproduce galaxies with the characteristic early- and late-type morphologies of the Hubble sequence (see Dubois et al. 2014, Schaye et al. 2015, Huertas-Company et al. 2019), and the correspondence of morphology (and/or kinematics) with the position of galaxies in the color–magnitude diagram or its proxies (Correa et al. 2017). This is not an outcome that the simulations were (or could meaningfully be) calibrated to reproduce and, in common with sizes, largely follows from the build up of angular momentum in intergalactic gas due to large-scale tidal torques, the dynamical coherence of gas as it is accreted onto halos and galaxies, the preferential removal of low angular momentum gas in feedback-driven outflows, and the transfer of angular momentum from stars to the DM as disks are disrupted by gravitational instabilities and mergers. Detailed analysis of the simulations reveals correlations between the morphology of galaxies and properties of their host halos, such as angular momentum (e.g., Zavala et al. 2016, Yang et al. 2021), flattening (Thob et al. 2019), and assembly time (Davies et al. 2020, 2021). However, it should be borne in mind that the simulations can reproduce observed morphologies only in a relatively broad sense, owing to their internal structure being too smooth and their vertical scale heights too large (see Section 2.3.3). It will be fascinating to assess the degree to which future simulations, incorporating detailed models of the multiphase ISM, are able to recover the detailed internal structure and kinematics of the galaxy population.

3.3. Galaxy Clustering

Accurate reproduction of observed clustering statistics, which connect galaxies to the properties of their DM environment, represents an important validation test of galaxy population simulations, particularly because clustering properties are not used as calibration diagnostics. Authoritative predictions also provide a means to stress-test (semi)analytic methodologies (Chaves-Montero et al. 2016, Guo et al. 2016), which are widely used to generate predictions for the large, mildly nonlinear scales probed by cosmological surveys.

Springel et al. (2018) showed that *IllustrisTNG* (both TNG100 and its lower-resolution, $L = 302.6$ cMpc counterpart TNG300) broadly reproduces, on scales of $1 \lesssim r \lesssim 10 h^{-1}$ Mpc and as a function of stellar mass and color, the clustering of galaxies revealed in the low-redshift cosmos by SDSS observations. Similarly, Artale et al. (2017) examined the clustering of *EAGLE* galaxies on scales of $1 \lesssim r \lesssim 6$ Mpc as a function of mass and color and found it to be broadly consistent with that measured from the GAMA survey, with the exception of low-mass (and poorly resolved) red galaxies, which cluster too strongly. Crain et al. (2017) showed that the clustering of gas-rich

galaxies in EAGLE is in excellent agreement with that inferred from 21-cm surveys and that the dependence of clustering on galaxy color (at fixed stellar mass) also manifests as a dependence on atomic hydrogen fraction.

3.4. Star-Formation Histories

The evolution of the cosmic SFR density (i.e., the volumetric SFR) is a fundamental observable whose measurement is a forefront goal in observational astronomy. Its precise characterization is challenging owing to selection effects and observational systematic uncertainties (see, e.g., Madau & Dickinson 2014, Behroozi et al. 2019), but it remains a natural validation benchmark for simulations of representative cosmic volumes.

The curves in **Figure 6**, in correspondence with the left y-axis, show the evolution of the cosmic SFR density of EAGLE, Horizon-AGN, IllustrisTNG, and SIMBA. The simulations that broadly reproduce the present-day GSMF yield similar SFR density histories (at the factor of $\simeq 2$ level). The integral of the cosmic SFR density (over cosmic time) differs from the integral of the GSMF (over mass) only due to mass loss from stellar evolution and a small correction due to the use of apertures to measure the stellar mass of galaxies. In the context of observational measurements, a factor of 2 is nonnegligible, being comparable with the long-standing tension between the observed cosmic stellar mass density and the mass loss–corrected integral of the cosmic SFR density,

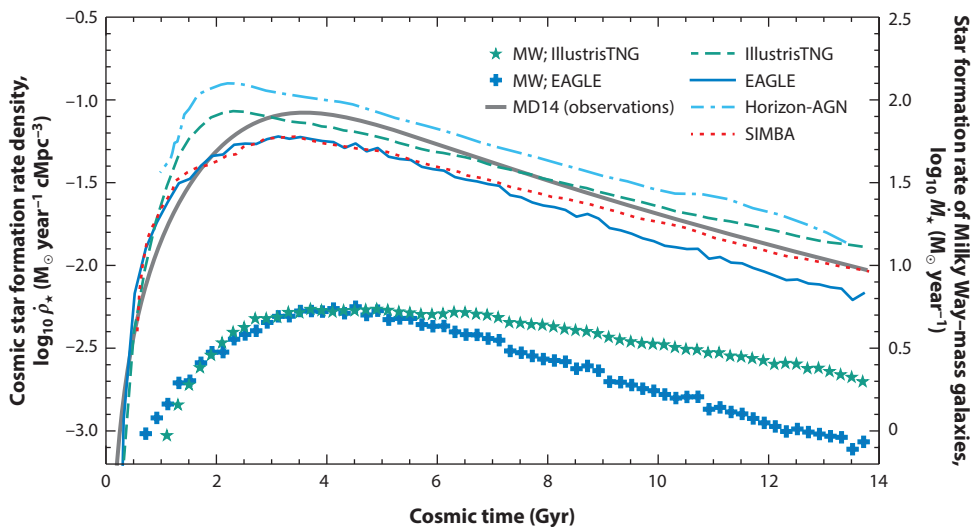


Figure 6

The colored curves show the cosmic SFR density history of EAGLE, Horizon-AGN (adapted from Kaviraj et al. 2017), IllustrisTNG, and SIMBA (adapted from Davé et al. 2019) (colors per legend). The gray curve denotes the fit to observations compiled by Madau & Dickinson (2014), converted to the Chabrier (2003) IMF. Colored symbols show the typical SFH (per the *right yaxis*) in EAGLE (*crosses*) and IllustrisTNG (*stars*) of central galaxies hosted by present-day halos with mass $M_{\text{halo}} \simeq 10^{12} M_{\odot}$, which is broadly similar to that of the Milky Way. The cosmic SFR densities in modern simulations agree with each other to within a factor of a few and the shapes are fairly similar, decreasing substantially over the past 10 Gyr. However, the SFH of a subset of galaxies can show larger variations as shown by the average SFHs of central galaxies in Milky Way-mass halos, which diverge toward the present day for EAGLE and IllustrisTNG. Abbreviations: IMF, initial mass function; MW, Milky Way; SFH, star-formation history; SFR, star-formation rate; SFRD, star-formation rate density.

though this apparent discrepancy may be resolved by the use of more sophisticated panchromatic spectral energy distribution models (Leja et al. 2019).

The symbols in **Figure 6**, relative to the right y -axis, show the average SFH of central galaxies hosted by present-day halos with mass $M_{\text{halo}} \simeq 10^{12} M_{\odot}$, which is broadly similar to the halo mass of the Milky Way, in **EAGLE** and **IllustrisTNG**. These histories peak at similar epochs with similar SFRs, but they diverge for $t \gtrsim 5$ Gyr ($z \lesssim 1.3$), declining more in **EAGLE** than in **IllustrisTNG**. Clearly, model-to-model differences at a particular halo mass scale are not simple rescalings of the difference in cosmic SFR density and instead reflect the complicated mass dependence of the physical processes treated by the simulations. This mass dependence is also manifest as subtle but significant differences in the shape of the GSMF realized by the simulations, such that the decline of the typical Milky Way–like SFH in **EAGLE** is likely related to the simulation “undershooting” the space density of galaxies at the knee of the GSMF. The SFHs of galaxies in simulations of the galaxy population are also impacted by the simplified subgrid treatment of the ISM (Section 2.3.3), which yields artificially smooth gas distributions and, by extension, artificially smooth SFHs. This hinders the use of SFHs from the simulation to interpret observations (e.g., Sparre et al. 2017) and may have significant consequences for galaxy evolution (Pontzen & Governato 2012).

3.5. Galaxy Colors

Galaxy surveys have revealed a remarkable bimodality in the color distribution of galaxies. Star-forming galaxies appear blue, whereas passive (or quiescent) galaxies are redder, owing to their lack of short-lived blue stars. There remains much ongoing debate concerning which processes regulate and quench star formation, with such processes usually categorized as internal or external/environmental. We focus here on color bimodality in relatively massive (and predominantly central) galaxies and defer discussion of the evolution of satellite galaxies embedded in larger host halos to Section 6.1.

Semianalytic models first highlighted that the emergence of color bimodality is a natural consequence of the use of energetic feedback in massive galaxies as a means to shape the high-mass end of the present-day GSMF away from the power-law distribution of DM halos (e.g., Kang et al. 2005). To date, this shaping is achieved in all realistic state-of-the-art simulations via the onset of efficient AGN feedback as the dominant regulation mechanism in halos of mass $M_{\text{halo}} \gtrsim 10^{12} M_{\odot}$. Because these simulations are usually calibrated to reproduce the GSMF, it is not unreasonable to expect that they yield bimodal galaxy colors. However, accurate reproduction of the observed distribution of galaxies in the color–mass plane also relies on the simulations yielding realistic SFHs and metallicities, because these properties also influence the observable properties of galaxies and are, hence, inputs to the population synthesis models used to translate physical properties into observables. One must also apply reasonable corrections for obscuration and attenuation by dust.

EAGLE (Trayford et al. 2015, 2017) and **IllustrisTNG** (Nelson et al. 2018b) have been shown to reproduce the key features of the color–mass plane recovered from galaxy surveys. Having satisfied this validation step, the simulations enable the origin of bimodality to be examined in detail and the sensitivity of the color–mass distribution to details of the models to be explored. **EAGLE** and **IllustrisTNG** present the consensus that bimodality is a direct consequence of the onset of AGN feedback. In **EAGLE**, the presence of blue galaxies that are brighter than those observed likely indicates that its AGN feedback does not adequately quench some fraction of massive galaxies (Trayford et al. 2015). Similarly, the greater fraction of red disks and blue spheroids in **IllustrisTNG** compared to observations may indicate that efficient AGN feedback is triggered in the wrong objects or at the wrong time (Rodríguez-Gómez et al. 2019); this is a problem whose resolution may require that AGN triggering is more sensitive to mergers (Bustamante & Springel

2019). Donnari et al. (2021) compare the closely related fraction of quenched galaxies as a function of stellar mass for a number of state-of-the-art simulations, finding qualitative but not quantitative agreement for central galaxies.

4. GALAXY SCALING RELATIONS

We next examine the degree to which well-known galaxy scaling relations are reproduced by state-of-the-art simulations of the galaxy population. The form of, and scatter about, these scaling relations encodes valuable information concerning the physics of galaxy evolution, but extracting insight from observed relations often requires the interpretive power of simulations. Confidence in the insight obtained is naturally greater if multiple simulations present a consensus for the origin of particular relations. As with Section 3, we caution that some of the scalings presented here are popularly used as calibration diagnostics and cannot be treated as predictions.

4.1. Supermassive Black Holes

The primacy of AGN feedback as the mechanism by which state-of-the-art simulations regulate and quench star formation in massive galaxies renders the relationship between the mass of galaxies and that of their central SMBH of particular importance. Although AGN feedback is not thought to dominate in low-mass galaxies, the latter are the sites in which SMBHs are seeded. SMBHs grow in concert with their host galaxy, yielding a power-law scaling between their masses. Whether this relation is causal is the subject of energetic debate, as it may, for example, reflect a natural outcome of hierarchical assembly (Jahnke & Macciò 2011), though it has been argued that this explanation is difficult to reconcile with the masses of SMBHs in dwarf galaxies (King & Nealon 2021).

Figure 7a, adapted from Habouzit et al. (2021), shows the relation between the mass of the SMBHs and their host galaxies in EAGLE, Horizon-AGN, IllustrisTNG, and SIMBA. The four simulations are broadly consistent with observed SMBH masses for $M_\star \gtrsim 10^{10.5} M_\odot$. In this regime the simulations exhibit scaling relations with similar slope and, as remarked in Section 2.4, the normalization of the relation largely reflects differences in the calibration/choice of the subgrid AGN efficiency parameter, ϵ_f . There is poorer consensus at lower galaxy masses, a regime in which the

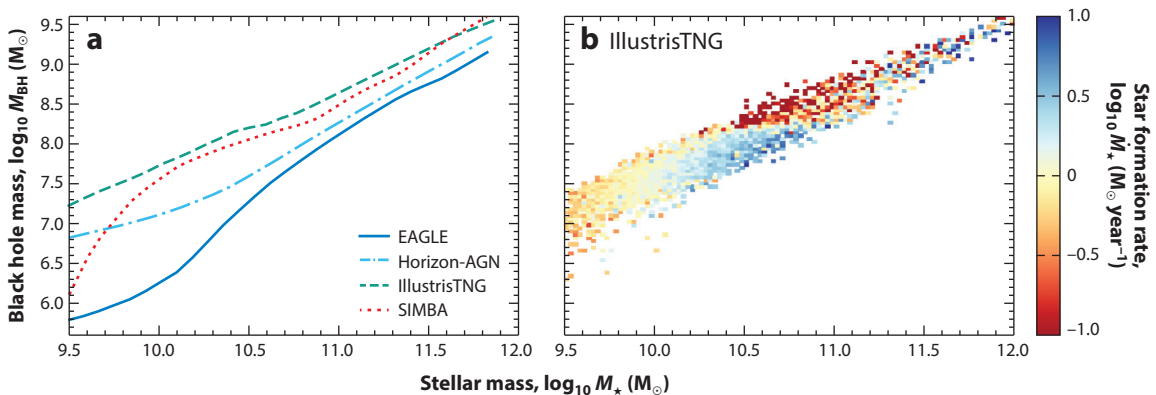


Figure 7

(a) The characteristic present-day SMBH mass (M_{BH}) as a function of stellar mass (M_\star) for EAGLE, Horizon-AGN, IllustrisTNG, and SIMBA (adapted from Habouzit et al. 2021). (b) As in panel a but showing the distribution of present-day central galaxies in IllustrisTNG, colored by their average SFR. For galaxies with $10^{10} \lesssim M_\star \lesssim 10^{11} M_\odot$, the SFR is lower in galaxies with more massive black holes, because they have experienced stronger AGN-driven outflows. Abbreviations: AGN, active galactic nucleus; SFR, star-formation rate; SMBH, supermassive black hole.

scatter in observed masses is also large (e.g., Kormendy & Ho 2013). In this regime, the growth of SMBHs is likely regulated by stellar feedback (e.g., Dubois et al. 2015, Bower et al. 2017), and the growth of simulated SMBHs is sensitive to degenerate details such as the seed mass, the subgrid accretion model, and subgrid treatments of unresolved dynamics (e.g., Bahé et al. 2022).

Figure 7b shows the present-day SMBH scaling relation in more detail for *IllustrisTNG*, with symbols colored by the SFR of the host galaxy. In massive galaxies ($M_{\star} \gtrsim 10^{10} M_{\odot}$), there is a strong negative correlation, at fixed stellar mass, between SMBH mass and SFR, which is indicative of the regulation of galaxy SFRs by AGN-driven outflows. Importantly, there is a qualitative consensus in state-of-the-art simulations for this particular correlation. However, as shown by Habouzit et al. (2021, see their figure 4), the simulations differ significantly in terms of the quantitative influence of AGN feedback on star-formation regulation.

4.2. The Star-Forming Main Sequence

Star-forming galaxies exhibit a tight correlation between their SFRs and stellar mass; this relation has become popularly known as the star-forming main sequence (SFMS; e.g., Noeske et al. 2007), whose normalization increases with increasing redshift. The ubiquity of the relation has led to the adoption of a more refined definition of what constitutes a passive galaxy, shifting from the canonical threshold at $\dot{M}_{\star} > 10^{-11} \text{ year}^{-1}$ to some number of decades from the median of the SFMS, thus naturally accounting for the mass and redshift dependence of the characteristic SFR of star-forming galaxies.

State-of-the-art simulations broadly reproduce the present-day SFMS, elucidating its origin and that of the scatter about it. Analyzing the SFMS in *IllustrisTNG*, Donnari et al. (2019) found that its normalization is mildly sensitive to both the timescale over and aperture within which the SFR is measured and that its intrinsic scatter (roughly at the factor of two level) is largely insensitive to stellar mass. However, they caution that the redshift evolution of the SFMS is not as strong as inferred from observations, which is a shortcoming in common with prior generations of simulations (e.g., Davé 2008). Matthee & Schaye (2019) examined the SFMS in *EAGLE*, finding a mild mass dependence. They showed that scatter in the present-day SFMS is due to both fluctuations on short timescales ($\lesssim 2 \text{ Gyr}$), associated with the physics of self-regulated gas flows, and long timescale ($\sim 10 \text{ Gyr}$) variations, due to differences in halo-formation histories. Related consequences of the short duration fluctuations are correlations at fixed stellar mass between the SFR and the outflow velocity of winds (Nelson et al. 2019) and between the SFR and the mass of cold gas that is (or will soon become) available for star formation (e.g., Lagos et al. 2016, Appleby et al. 2020), the latter being readily corroborated by observations (Saintonge & Catinella 2022).

4.3. The Tully–Fisher Relation

The Tully–Fisher (Tully & Fisher 1977) relation is a well-known scaling relation connecting the asymptotic rotation velocity of disk-dominated galaxies and their luminosity (or mass). The strongest correlation is recovered when considering the total baryonic mass of galaxies (McGaugh et al. 2000). A similar relation exists for elliptical galaxies between their central velocity dispersion and their mass (or luminosity; Faber & Jackson 1976). These relations are the traditional means by which the properties of galaxies were linked to those of their halos, because galaxy dynamics (in standard cold dark matter cosmogonies) are primarily governed by the structure of their halos.

The failure of prior generations of simulations to reproduce the observed rotation profiles of individual galaxies (summarized succinctly by Scannapieco et al. 2012), largely as a consequence of overcooling, translated into a failure of simulations of galaxy populations to reproduce the

high-mass end of the Tully–Fisher relation (e.g., McCarthy et al. 2012). The inclusion of AGN feedback and the calibration of feedback models in state-of-the-art simulations results in a more realistic stellar mass–halo mass relation and prevents artificial contraction of the halo due to the excessive condensation of stars (e.g., Schaller et al. 2015). As a result, these simulations broadly reproduce the observed Tully–Fisher relation (e.g., Vogelsberger et al. 2014, Ferrero et al. 2017, Sales et al. 2017) without having been explicitly calibrated to do so.

4.4. The Mass–Metallicity Relations

Feedback-driven outflows transport a fraction of the heavy elements synthesized in galaxies into the CGM (Tumlinson et al. 2011), and even beyond to the IGM (Aguirre et al. 2001). This establishes a relationship between the mass of a galaxy and the characteristic metallicities of its gas and stars (Larson 1974). This mass–metallicity relation was revealed in detail by the advent of highly multiplexed spectroscopic surveys such as SDSS, which enabled both the gas-phase metallicity (Tremonti et al. 2004) and that of the stars (Gallazzi et al. 2005) to be measured for tens to hundreds of thousands of galaxies. These studies thus reveal that a greater fraction of the metals synthesized by low-mass galaxies, relative to more massive counterparts, are transported into their gaseous environments. These trends are broadly reproduced by *EAGLE* (Schaye et al. 2015), *Horizon-AGN* (Dubois et al. 2014), *IllustrisTNG* (Torrey et al. 2019), and *SIMBA* (Davé et al. 2019) and are a natural consequence of mass-dependence of the macroscopic mass loading of outflows that emerges from their subgrid models for feedback (Beckmann et al. 2017, Nelson et al. 2019, Mitchell et al. 2020; see also Muratov et al. 2015, Christensen et al. 2016).

The observations reveal significant scatter in metallicity at fixed stellar mass, and other properties have been shown to correlate with this scatter, such as SFR (the fundamental metallicity relation); this behavior is reproduced by simulations (e.g., Lagos et al. 2016, Davé et al. 2019, Torrey et al. 2019). De Rossi et al. (2017) showed using *EAGLE* that one might expect a complex mass dependence on the correlation of metallicity with gas fraction and SFR. In low-mass galaxies, lower metallicities correspond to higher gas fractions, higher SFRs, and the presence of young stellar populations, whereas at higher-mass scales lower metallicities correspond to gas poor, quiescent galaxies. They interpreted this inversion as a consequence of the switch from stellar feedback–based regulation in low-mass galaxies to AGN-dominated regulation.

4.5. Cold Gas in Galaxies

Forging a holistic view of galaxy evolution requires an understanding of the evolution of their gaseous components, which can span many decades in density, temperature, and ionization fraction. Recovery of these properties is a prerequisite for the use of the simulations to elucidate the cycling of baryons into, and out of, galaxies (e.g., Tumlinson et al. 2017, Péroux & Howk 2020), which is driven by the interplay of large-scale cosmic structure evolution and the complex physics of galaxy formation.

Scaling relations relating the mass of cold, dense gas (primarily atomic and molecular hydrogen) can be readily constructed from observations of galaxies in the low-redshift cosmos (see, e.g., Saintonge & Catinella 2022). Such relations present an appealing benchmark, but because large simulations of the galaxy population do not in general employ an explicit radiation transport scheme, they cannot model the balance of molecular, atomic, and ionized hydrogen on the fly in a self-consistent fashion. The partitioning is therefore achieved as a postprocessing step based on empirical relationships and/or theoretical models.

Building on prior examinations of *EAGLE* (Lagos et al. 2015, Bahé et al. 2016, Crain et al. 2017) and *IllustrisTNG* (Diemer et al. 2018, Stevens et al. 2019), Davé et al. (2020) compared the cold

gas mass functions and scaling relations of these simulations with those of SIMBA. The comparison revealed a broad qualitative consensus but also a number of significant quantitative differences, such as in the characteristic break scale of the H I mass function and the direction in which this quantity shifts with increasing redshift. None of the simulations were found to universally reproduce the key cold gas observational diagnostics, with each exhibiting a different conflict with the data, despite all three simulations broadly reproducing the present-day GSMF and SFMS. This signals that the combined action of the subgrid treatments of interstellar gas, star formation, and feedback in these simulations is too simplistic to reproduce holistically the cold gas and stellar properties of present-day galaxies.

5. COSMIC GAS

The explicit modeling of cosmic gas flows and their interaction with galaxies is a major selling point of hydrodynamical simulations relative to simplified methods such as semianalytic modeling. The inclusion of energetic feedback processes in hydrodynamical simulations has indicated that a complex and intimate connection likely exists between the life cycles of galaxies and their gaseous environments, for which there is a growing body of corroborating observational evidence (e.g., Péroux & Howk 2020). Besides yielding fascinating insight in its own right, study of the gaseous cosmos is therefore now widely acknowledged as a complementary and profitable means of understanding galaxy formation and evolution, particularly in regard to constraining the physics of feedback and metal transport. In this section, we review key properties of cosmic gas, on the scale of the cosmic large-scale structure (the IGM) and within galaxy halos (the CGM/ICM), that emerge from state-of-the-art simulations of the galaxy population. The confrontation of these outcomes with observational measurements represent strong tests of the simulations, because such properties are less frequently considered than stellar properties when calibrating subgrid feedback models.

5.1. Absorption System Statistics

The low-redshift column density distribution functions of intermediate- and high-ionization state metal ions in EAGLE and IllustrisTNG have been presented in a number of studies (e.g., Schaye et al. 2015, Nelson et al. 2018a, Wijers et al. 2019). These demonstrate that the simulations are mostly compatible with observational measurements of the column density distribution function of absorption systems seen in quasar spectra, which exhibit significant scatter at fixed column density. Differences between the simulations are limited to the relatively high column regime ($N \gtrsim 10^{15} \text{ cm}^{-2}$), in which EAGLE yields too few OVI absorbers and IllustrisTNG yields too many. In both EAGLE and IllustrisTNG, the abundance of OVI absorbers increases slightly with increasing resolution. Wijers et al. (2019) demonstrated that the abundance of high column density OVII and OVIII absorbers is influenced by AGN feedback in a nontrivial fashion: AGN-driven outflows increase the metallicity of absorbers by transporting oxygen into halo outskirts but reduce the characteristic density of the absorbing gas. However, the close correspondence of the distribution functions of OVI, OVII, and OVIII in both simulations indicates that these diagnostics are, perhaps surprisingly, only mildly sensitive to the choice of hydrodynamics solver and details of how AGN and stellar feedback processes are implemented.

A related diagnostic is the covering fraction of absorption systems in the vicinity of galaxies. Prior generations of simulations failed to reproduce the high covering fraction of high column density ($N \gtrsim 10^{17} \text{ cm}^{-2}$) neutral hydrogen absorbers associated with star-forming galaxies at $z \simeq 2-3$ (e.g., Fumagalli et al. 2014). Using EAGLE, Rahmati et al. (2015) demonstrated that the H I covering fraction is largely determined by the supply of neutral hydrogen into halos, which

is governed primarily by gravitational infall and only weakly influenced by feedback processes. However, feedback crucially governs the relationship between galaxy stellar mass and halo mass, such that recovering realistic covering fractions for simulated galaxies of a fixed stellar mass requires that they form in realistic environments. The greater feedback efficiency of state-of-the-art simulations relative to earlier generations of simulations ensures that $z \simeq 2\text{--}3$ star-forming galaxies form in more massive halos, which exhibit radial H I covering fraction profiles consistent with those observed.

5.2. Physical Properties of the Circumgalactic Medium

The physical properties of the CGM have long been viewed as a powerful means of constraining the physics of galaxy formation, because the CGM interfaces galaxies with their incoming supply of fuel for star formation and their feedback-driven outflows of enriched gas. A salient example is the opportunity to elucidate how outflows operate: The stellar masses of galaxies (for example) are largely agnostic to whether cosmic gas inflows are regulated by violent, intermittent episodes of feedback or a gentler and more continuous process, but the physical conditions of the CGM are expected to differ strongly between these cases (e.g., van de Voort & Schaye 2012). Using EAGLE, Oppenheimer et al. (2018) examined the physical conditions traced by low ionization state absorption systems in the CGM of galaxies with mass comparable with that of the Milky Way to assist the interpretation of results from the COS (Cosmic Origins Spectrograph)-Halos survey (Werk et al. 2013, 2014). They recovered column densities for most ions within a factor of $\simeq 2$ of those reported by COS-Halos. They found little correlation of absorber column densities with the specific SFR of central galaxies, but a significant decline of column densities at larger impact factors, with both findings being consistent with the observations. Their analysis elucidated that these low ionization state absorption systems trace cool ($T \sim 10^4$ K) clumps of gas close to galaxy disks ($\lesssim 100$ kpc), distinct from the warmer ($T \simeq 10^{5.5}$ K) and more diffuse gas traced by O VI absorbers, usually found at $\gtrsim 150$ kpc. This suggests that low ionization state absorbers likely trace enriched gas that is reaccreting onto galaxies.

Perhaps unsurprisingly, comparison of similar CGM diagnostics from different simulations reveals areas of both consensus and discrepancy. Péroux et al. (2020) demonstrate that the gas mass fluxes and metallicities as a function of the azimuthal angle about star-forming galaxies are qualitatively similar in EAGLE and the high-resolution $L = 51.7$ cMpc TNG50 simulation of the IllustrisTNG suite. Both simulations exhibit a strong azimuthal dependence such that the major axis is dominated by relatively metal-poor inflows and the minor axis by enriched outflows. They also broadly reproduce the dichotomy of O VI column densities associated with star-forming and passive galaxies revealed by COS-Halos (Tumlinson et al. 2011). However, the origin of the relationship differs between the simulations. Based on analysis of EAGLE, Oppenheimer et al. (2016) argue that passive galaxies exhibit a deficit of O VI because they preferentially reside in more massive halos (in which oxygen tends to be more strongly ionized) than star-forming galaxies of comparable stellar mass. Nelson et al. (2018a) report a more direct connection between star formation and the transport of oxygen-enriched gas into the CGM by AGN feedback in IllustrisTNG. Such differences may be the tip of an iceberg, as the inclusion of physical processes not treated by either EAGLE or IllustrisTNG, for example, feedback by cosmic rays, may induce significant changes to the CGM (Ji et al. 2020, Butsky et al. 2022).

5.3. Gas Inflows and Outflows

The ISM and CGM are essentially gas reservoirs bound to galaxy halos, whose mass evolves in response to a shifting balance between, primarily, galaxy- and halo-scale inflows and outflows, with

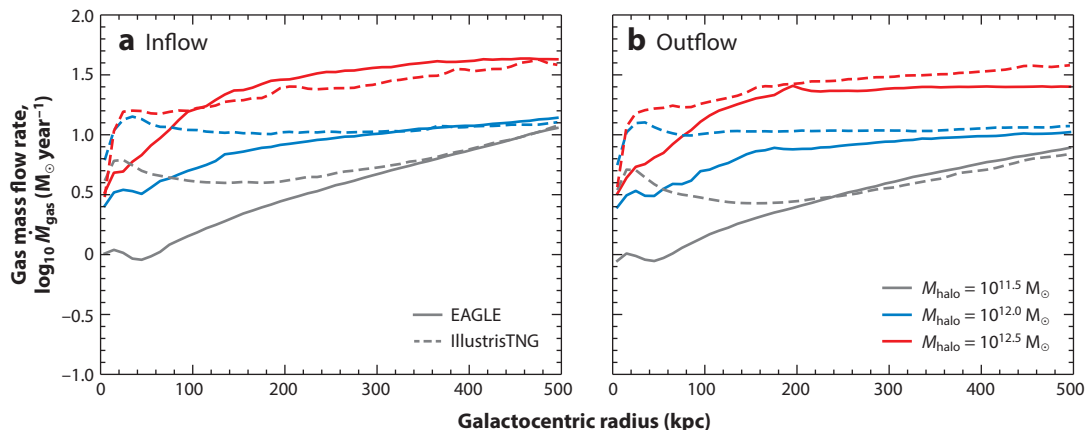


Figure 8

Inflow (a) and outflow (b) rates as a function of galactocentric radius for $M_{\text{halo}} = 10^{11.5} M_{\odot}$ (gray curves), $10^{12} M_{\odot}$ (blue curves), and $10^{12.5} M_{\odot}$ (red curves) in EAGLE (solid curves) and IllustrisTNG (dashed curves) at $z = 0$. Even though both simulations produce similarly realistic galaxies, the flow of gas into and out of galaxies and their halos can be substantially different. Additionally, the magnitude and sign of these differences depend on halo mass.

a generally minor contribution also played by star formation, stellar mass loss, and, in some cases, external environmental processes. In **Figure 8**, we show the halo mass–dependent present-day inflow and outflow rates as a function of galactocentric radius in EAGLE and IllustrisTNG. Comparison to the present-day SFRs of the central galaxies hosted by $10^{12} M_{\odot}$ halos shown in **Figure 6** highlights that the inflow and outflow rates are higher than the SFRs for both simulations.

Subgrid feedback models aim to approximate the impact of unresolved feedback processes on resolved scales. The flow rates presented in **Figure 8** illustrate that, despite the subgrid treatments in EAGLE and IllustrisTNG yielding broadly similar present-day galaxy populations, the radial transport of gas through their halos differs significantly (see also Mitchell et al. 2020). For example, inflow rates in IllustrisTNG halos are fairly constant as a function of galactocentric radius, whereas those in EAGLE decrease more strongly toward the halo center. Outflow rates in the two simulations are reasonably similar for the more massive halos ($10^{12.5} M_{\odot}$) and at large radii, whereas inflow rates are different by a factor of two in the low and high halo mass range. Within ≈ 50 kpc of the center of the less massive halos, the simulations differ by a factor of three or more.

Differing rates of inflow and outflow lead to differing halo baryon fractions, with high (net) inflow rates at the virial radius (as seen for low-mass halos in IllustrisTNG) being conducive to a high baryon fraction (see **Figure 9**) if sustained over a significant fraction of cosmic time. In the absence of feedback and large-scale outflows, the gas accretion rate at the virial radius is similar to the product of the cosmic baryon fraction, (Ω_b/Ω_m), and the DM accretion rate, but large-scale feedback-driven outflows can significantly suppress this accretion rate and so also act as a form of preventative feedback (e.g., Wright et al. 2020) that inhibits intergalactic gas from reaching the CGM and the ISM. The extent to which this happens is likely to depend strongly on the subgrid implementation of the feedback, and the differing flow rates shown in **Figure 8** indicate that state-of-the-art simulations do not present a consensus.

5.4. Halo Baryon Fractions

The baryon fraction of halos, i.e., the ratio of the baryonic and total masses within the virial radius, is a closely related diagnostic to halo gas flows. In the absence of radiative physics (cooling and

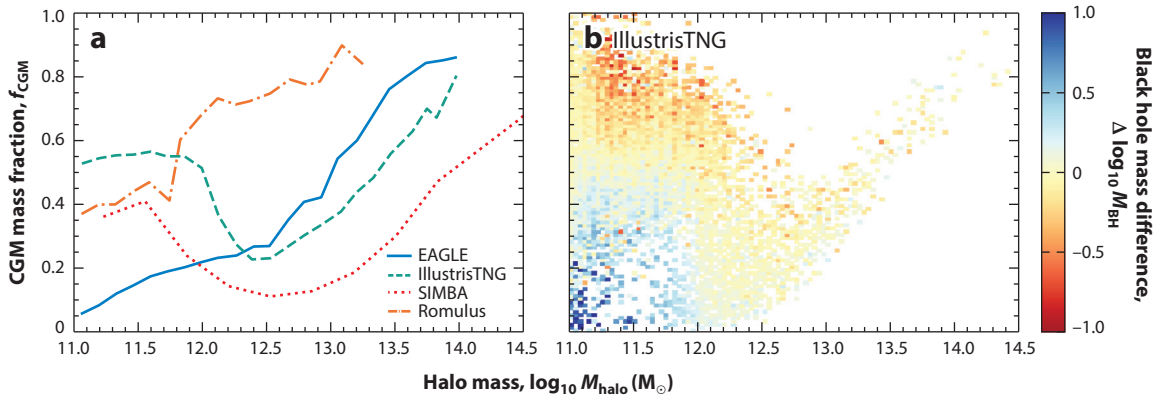


Figure 9

The gas fraction in the CGM (the CGM mass divided by the mass of the halo, normalized by the cosmic average baryon fraction), as a function of halo mass. Panel *a* shows the median for EAGLE, IllustrisTNG, SIMBA (data provided by R. Davé) and Romulus (data provided by M. Tremmel). The full distribution for IllustrisTNG is shown in panel *b*, color coded by the average logarithmic difference in SMBH mass relative to the median SMBH mass at the same halo mass. Halos with more-massive black holes have lower gas fractions for $M_{\text{halo}} \lesssim 10^{12.5} M_{\odot}$, because these halos have experienced more feedback from AGNs. Abbreviations: AGNs, active galactic nuclei; CGM, circumgalactic medium; SMBH, supermassive black hole.

feedback), simulations indicate that baryon fractions should be close to the cosmic average baryon fraction of $\Omega_b/\Omega_m \approx 0.16$ (Crain et al. 2007). Feedback-driven outflows, however, can markedly reduce the baryon fraction. **Figure 9** shows the median halo CGM mass fractions as a function of halo mass for EAGLE, IllustrisTNG, SIMBA, and Romulus (Tremmel et al. 2017). We include the latter here as it features a markedly different scheme for accretion onto SMBHs.

The baryon fractions of the massive halos that host rich galaxy groups and clusters are reasonably well constrained (Vikhlinin et al. 2005, Sun et al. 2009), and it is now well established that, owing to differences in the subgrid implementations of feedback processes, accurate reproduction of the GSMF offers no guarantee of accurately recovering realistic baryon fractions on these scales (e.g., Schaye et al. 2015, Haider et al. 2016). However, adjustment of the feedback implementations, particularly for AGNs, enables the correspondence with observed baryon fractions to be improved (Schaye et al. 2015, McCarthy et al. 2017, Weinberger et al. 2017). In less massive halos, gas fractions are largely unconstrained by observations and, as shown in **Figure 9**, the median relations that emerge from state-of-the-art simulations are strikingly dissimilar. As highlighted by Davies et al. (2020), the median present-day CGM mass fraction at $M_{\text{halo}} = 10^{12} M_{\odot}$ is $\simeq 0.2$ in EAGLE, but > 0.5 in IllustrisTNG.

Despite the dissimilar trends between CGM mass fraction and halo mass in EAGLE and IllustrisTNG, Davies et al. (2020) showed that in both simulations the scatter in CGM mass at fixed halo mass correlates strongly with the formation redshift of the halo (negatively), the mass of the central galaxy’s SMBH (negatively), and the central galaxy’s SFR (positively). The negative correlation with SMBH mass is shown for IllustrisTNG in **Figure 9b**, colored by the average deviation in logarithmic SMBH mass relative to the median value. Davies et al. (2020) interpreted these correlations as an indication that early growth enables the SMBH to begin delivering efficient AGN feedback sooner, expelling more gas from the halo and extending the radiative cooling timescale of the remaining CGM. The reduced cooling efficiency inhibits replenishment of the ISM as it is consumed by star formation (or ejected by the associated stellar feedback), reducing the SFR or even quenching the galaxy. Similar trends have also recently been found in SIMBA (Sorini

et al. 2022), and this qualitative consensus highlights that an intimate coevolution of galaxies and their CGM is a major prediction of state-of-the-art simulations.

6. THE INFLUENCE OF ENVIRONMENT

Environment, in the context of galaxy evolution, can refer to a variety of important physical processes and effects. The cosmological environment governs the accretion rate onto halos, which can be spatially inhomogeneous because of the filamentary structure of the cosmic web, as well as the rate of mergers. This profoundly impacts the growth of galaxies in combination with smaller-scale processes, such as galactic winds. The large-scale structure and its inhomogeneity can lead to differences between cosmological and idealized simulations, though we do not further discuss this here.

The environment of a galaxy can also refer to whether or not it lives at the minimum of its host halo's potential well, where the gas densities and cooling rates are highest. Satellite galaxies used to reside at the centers of their own halos but have since fallen into more massive halos and are thus offset from the center of the potential. Back-splash galaxies or flyby galaxies are those that were satellites in the past but have left the more massive halo they were previously in. They may therefore seem to not live in extreme environments but can still have been strongly affected in the past. Note that there is no well-defined halo boundary, and environmental effects driven by the halo can extend beyond the virial radius. When galaxies experience such effects before entering the more massive halo, this is also sometimes referred to as preprocessing.

6.1. Satellite Galaxies: Stripping and Starvation

Satellite galaxies make up a large proportion of the galaxy population and therefore must be included if we are to understand galaxy formation fully. Satellites make up more than 40% of the population below a stellar mass of $3 \times 10^{10} M_{\odot}$, though this decreases to less than 20% above $M_{*} = 10^{11} M_{\odot}$. Although we often think about processes involved in the formation of central galaxies and satellite galaxies separately, each satellite spent its early life (possibly most of its life) as a central and is therefore also shaped by the physical processes relevant for central galaxy formation. However, after becoming satellites, these galaxies evolve differently from central galaxies because they experience various environmental effects, such as ram pressure and tidal stripping, as they move through the halo of their more massive companion. Note that a fraction of the galaxy population considered central galaxies are back-splash galaxies and have thus previously experienced stronger environmental effects.

Ram pressure is exerted on a body of gas when it moves through a medium with a different velocity. It scales with the density of the surrounding medium and the square of the velocity difference. Ram pressure only directly affects gas, whereas tidal effects are purely gravitational and, thus, affect all components of a satellite, i.e., DM, gas, and stars. When the ram pressure is high enough to remove the ISM of the satellite galaxy, this is typically referred to as ram pressure stripping (e.g., Gunn & Gott 1972, Simpson et al. 2018). With the ISM removed, the satellite can no longer form stars, so this results in relatively fast quenching of star formation. When only the gas in the halo of the satellite is removed, the galaxy can continue forming stars from its ISM until it is depleted. This is often referred to as starvation or strangulation, and the quenching of star formation is much slower (e.g., van de Voort et al. 2017, Wright et al. 2022). In both cases, the satellite will not be able to replenish its ISM as it no longer has its own CGM acting as a reservoir from which it can accrete fresh gas.

Modern cosmological simulations have sufficient resolution to capture these stripping processes, though resolution effects can make a quantitative difference because galaxies that are less well resolved are more easily disrupted (e.g., Yun et al. 2019). If simulations were able to fully model the multiphase ISM, it is likely that the cold, dense gas would be less easily stripped than in current simulation suites. Wright et al. (2022) find that starvation and ram pressure stripping of the ISM contribute a similar amount to the reduction of the SFR in satellite galaxies. Some evidence has been found that ram pressure stripping is stronger along the central galaxy’s major axis than along its minor axis, possibly because the density in the polar direction is reduced due to feedback-driven outflows (Martín-Navarro et al. 2021).

Quenching, or the suppression of star formation, is a complicated and unsolved problem, as discussed before in Section 3.5. Quenching in satellites depends on internal feedback processes as well as external stripping. Simulations have struggled to quantitatively reproduce observed quenched fractions at all stellar masses. There is general agreement that the majority of satellites are quenched, with the quenched fraction increasing toward lower masses (Bahé et al. 2017, Donnari et al. 2021). However, observations seem to show the opposite trend and very low quenched fractions for $10^9 < M_* < 10^{10} M_\odot$ (Davies et al. 2019). For an unbiased comparison, it will be necessary to process the simulations in the same way as is done for the observational measurements.

6.2. Environmental Effects: Beyond Halo Mass

The environment of a galaxy is dominated by properties of the DM-dominated halo it lives in (e.g., Crain et al. 2009). Although straightforward to determine in simulations, halo mass is difficult to measure observationally. Instead, there are a variety of observational methods that can be used to quantify the environments of galaxies, such as distance to the N th nearest neighbor. These often-used environmental indicators correlate strongly with the halo mass (Haas et al. 2012, Marasco et al. 2016). Many of the differences identified based on environmental indicators exist because they are sensitive to host halo mass, and satellite galaxies behave differently from central galaxies, as discussed in Section 6.1.

In addition to a dependence of galaxy formation on halo mass, the larger-scale environment the halo lives in can also influence the evolution of embedded galaxies to a certain extent. To study environmental effects that go beyond the halo mass of the system, it is important to remove the halo mass dependence and characterize the large-scale environment. This usually requires a large-scale structure finder, in both observations and simulations, or visual classification. Rosas-Guevara et al. (2022) build a void catalog and compare galaxies living in those voids to those residing in the cosmic web. Overall the differences are fairly minor; e.g., average stellar masses can change by up to about 30%. They find that this large-scale environment affects low-mass galaxies differently from high-mass galaxies, which likely indicates that there are a variety of processes at play. The galaxy mass in low-mass halos is lower in voids, potentially due to a combination of lower accretion rates and lower merger rates, whereas it is higher in massive halos, potentially because of lower black hole growth and a lower amount of AGN feedback in void galaxies.

Ram pressure stripping, as discussed in Section 6.1, is not limited to only taking place inside galaxy halos. The filaments of the cosmic web can also harbor strong accretion shocks and provide ram pressure, which can result in stripping of the gas in and around dwarf galaxies that pass through (Benítez-Llambay et al. 2013, Herzog et al. 2022, Pasha et al. 2022). The ambient temperature of the gas in these filaments can be higher than the virial temperatures of dwarf galaxies, which means that they are not able to accrete fresh gas and, thus, remain gas deficient. These galaxies have lower SFRs and lower ISM masses than galaxies not affected by stripping in the cosmic web, yet they can appear quite isolated.

7. FUTURE OUTLOOK

The level of realism achieved by state-of-the-art cosmological hydrodynamical simulations of the galaxy population has advanced dramatically in the past decade. Despite this success, the outcomes are particularly sensitive to subgrid implementations of feedback processes. It is thus reasonable to argue that this success suffices only to establish that the basic sketch of galaxy-formation theory within the Λ CDM cosmogony is plausible. The development of simulations that can be used to stress test a truly comprehensive theory of galaxy formation and evolution will require confrontation of the devil in the details. Simulations of the galaxy population do not accurately reproduce the internal and vertical structure of galaxies and their reliance on subgrid methods to approximate the influence of unresolved physics leaves lingering degeneracies that diminish predictive power. Furthermore, the simulations neglect physical processes known to be significant in certain regimes. From our suite of examples, only *IllustrisTNG* models magnetic fields, which could impact cosmic gas flows and thus affect the evolution of galaxies (Pillepich et al. 2018a, van de Voort et al. 2021). The omission of magnetic fields also precludes realistic modeling of the influence of cosmic rays, which, owing to their ISM energy density being comparable with those of the thermal and magnetic pressures, likely influence galaxy-wide outflows (Uhlig et al. 2012). Similarly, thermal conduction influences the structure of the ICM and cooling onto massive galaxies (Carilli & Taylor 2002) but is generally neglected. Modeling the UVB as a spatially uniform radiation field that “switches on” at a fixed redshift is a particularly crude approximation for the evolution of galaxies during the EoR, when local sources of ionizing radiation can dominate the regulation of galaxy growth (Wise & Cen 2009, Trebitsch et al. 2017, Katz et al. 2020). Locally varying radiation fields, such as those produced by flickering AGNs, can also overionize gas and lead to cooling rates being overestimated by orders of magnitude if equilibrium conditions are assumed (Vasiliev 2011, Richings et al. 2014).

It is natural to consider how the field will build on the foundation provided by the current state-of-the-art generation of hydrodynamic simulations of the galaxy population. There are several routes by which we envisage that progress will be made, beyond the usual pursuit of superior resolution (which is needed to achieve converged results in the CGM; e.g., van de Voort et al. 2019). Despite remarkable recent demonstrations of simulation codes running on hundreds of thousands of computer cores (e.g., Schaller et al. 2016, Pakmor et al. 2022), subgrid models will remain a necessary component of simulations of the galaxy population for some years to come, precluding holistic, quantitative predictive power. There is, however, much scope to develop more detailed implementations and to interface them with the numerical calculation on shorter spatial scales, enabling more detailed confrontations with observations and increasing the diversity of the lines of inquiry for which the simulations can offer authoritative predictions. Furthermore, the ill-constrained parameters of the subgrid models used by leading simulations have been calibrated manually, by performing a small set of parameter-spanning simulations, analyzing their outputs, and choosing updated parameters for the subsequent set based on the practitioner’s intuition for the response of the galaxy population to the adjustment. Statistical methods exist to formalize this process and ensure that the plausible parameter space is efficiently explored. There is also much to be gained from improving not only the simulations themselves but also the techniques used to analyze their outcomes.

The development and testing of new or improved treatments of physical processes, whether implemented numerically or subgrid, is usually pioneered using simulations that adopt idealized or zoomed cosmological initial conditions. At fixed resolution, such simulations are markedly less expensive than simulations of representative volumes or, for a fixed number of computational core hours, they allow galaxies to be evolved at much higher resolution. Such simulations have been

used to pursue more detailed treatments of the multiphase ISM that require fewer assumptions, in particular relating to the star-formation efficiency of dense gas (e.g., Hopkins et al. 2014, Semenov et al. 2016, Kim & Ostriker 2017). Feldmann et al. (2023; see also **Figure 4**) showed recently that the use of these more detailed models in a periodic cosmological volume does not guarantee the emergence of realistic GSMF, even in the low stellar mass regime for which the included physical processes are expected to dominate galaxy regulation. A number of suites of zoom simulations focusing on dwarf galaxies have achieved sufficiently high resolution to model individual SN explosions within a multiphase ISM (e.g., Wheeler et al. 2019, Agertz et al. 2020, Gutcke et al. 2022). Detailed study of these processes in idealized or zoom simulations paves the way toward development of coarse-grained descriptions that, although still approximate, can be more realistic than the simple (and diverse) treatments of stellar feedback used by today’s galaxy population simulations. An extreme interpretation of this methodology is to treat outflows exclusively with a phenomenological approach, even on numerically resolved scales (e.g., Huang et al. 2020). Similarly, the multiphase structure of the CGM is not resolved in cosmological simulations, specifically the cool gas and the interfaces between the cool and hot CGM (e.g., McCourt et al. 2018, Fielding et al. 2020; see also Faucher-Giguère & Oh 2023, in this volume). Accurately treating the cool and intermediate temperature gas may require developing subgrid models within the CGM.

The recent discovery with JWST of multiple galaxies with spectroscopically confirmed redshifts of $z > 10$ (Curtis-Lake et al. 2023) highlights the urgent need to model the EoR galaxy population with explicit radiation hydrodynamics (RHD). Zoom simulations have proven useful for examination of the internal structure of individual galaxies in this regime (e.g., Pallottini et al. 2017) but suffer from potential selection biases and preclude examination of the influence of galaxies on the IGM as the latter undergoes a global phase transition. High-resolution RHD simulations of representative volumes remain extremely challenging, but Borrow et al. (2022a) were able to show, using simulations of small periodic volumes, that the use of a spatially uniform UVB fosters significant inaccuracies in the space density and internal properties of low-mass simulated galaxies in the EoR and results in the density and temperature structure of the IGM being too uniform.

In recent years, machine learning has steadily grown in influence in nearly every aspect of astrophysics (see the review by Smith & Geach 2023). Machine learning techniques have emerged as an effective means of elucidating formally the complex relationships between the properties of simulated galaxies and those of their host DM halos (Icaza-Lizaola et al. 2021, Piotrowska et al. 2022). The CAMELS (Cosmology and Astrophysics with Machine Learning Simulations) project (Villaescusa-Navarro et al. 2021) uses many simulations evolved with variations of the *IllustrisTNG* and *SIMBA* models, varying both cosmological and subgrid parameters, as a basis for the application of machine learning techniques. Jo et al. (2022) use the suite as a testbed for calibrating model parameters using neural networks trained as emulators. Interestingly, they found that the emulators could identify parameter sets that accurately reproduce a GSMF drawn from the input simulations chosen as a “target observable” but none that reproduce a real observationally inferred GSMF. It is unclear whether this signals a problem with the method, inconsistencies in the observationally inferred GSMF, or a genuine limitation of the physics implemented within the input simulations. If the latter, this methodology may prove a particularly effective means of guiding the development of more sophisticated subgrid models.

8. SUMMARY

Cosmological, hydrodynamical simulations have catalyzed significant strides in our understanding of the physics that governs the formation and evolution of galaxies. The much improved realism of the current generation of state-of-the-art simulations, relative to their predecessors, has

engendered greater confidence in the insight obtained from the confrontation of simulations with observational data and has diversified the lines of inquiry for which such comparisons are useful.

A crucial lesson learned from some of the first simulations of galaxies is that their evolution is strongly regulated by energetic feedback processes. Modern, realistic simulations model (at minimum) feedback associated with the formation and evolution of stellar populations (which dominates in low-mass galaxies) and that associated with the accretion of gas onto SMBHs (which dominates in massive galaxies). However, it remains the case (and will do so for some years to come) that this modeling is achieved in an approximate fashion using subgrid models, whose governing parameters are ill constrained. The greater realism of the current generation of simulations has therefore followed from the pragmatic approach of calibrating these parameters.

Simulations calibrated to reproduce the stellar masses and sizes of galaxies also reproduce many well-known observed scaling relations, e.g., the Tully–Fisher relation and the SFMS. This is often because they concern related properties or because reproduction of the relation relies primarily on ensuring that galaxies of a fixed stellar mass are associated with DM halos of the correct total mass and, hence, exhibit (for example) realistic space densities and cosmic matter inflow rates. Other scaling relations, particularly those related to the baryon cycle in halos such as the mass–metallicity relation, the baryon fraction–halo mass relation, and galaxy–absorber statistics, are more sensitive to the details of the adopted subgrid models. Quantitative agreement between the simulations and observations for these relations tends to require that the simulations be calibrated specifically to achieve it. The quantities in these relations also tend to be those for which numerical convergence is most challenging to achieve. Perhaps unsurprisingly then, there is a conspicuous absence of intersimulation consensus in this regime.

Different subgrid prescriptions therefore appear able to reproduce many galaxy properties equally well while yielding predictions for other properties (often those only weakly constrained by observations) that are markedly different. As we highlight in the summary boxes below, the current generation of state-of-the-art simulations of the galaxy population has therefore fostered important and enduring successes, but many outstanding challenges remain. The simulations provide a sound foundation from which to pursue the more sophisticated models needed to tackle these challenges and further our understanding of the galaxy population. This pursuit will, of course, also be driven by future observational discoveries, and more detailed characterization of the CGM is likely to prove among the most fruitful avenues for constraining, in particular, the macroscopic effects of energetic feedback.

Finally, we remark that the public data release (with detailed accompanying documentation) of major simulation campaigns, while being a major undertaking for the simulation teams, has proven a tremendously successful exercise. It has enabled their use by many researchers who were not involved with the development of the simulations, and many astronomers without prior specialism in simulations, allowing more thorough exploitation and analysis of the simulations and more diverse comparisons with observations. We energetically encourage the developers of future simulation campaigns to follow suit.

SUMMARY POINTS

1. By regulating star formation with plausible quantities of energetic feedback associated with the formation of stars and the growth of black holes, the current generation of state-of-the-art cosmological hydrodynamical simulations form a galaxy population with broadly realistic stellar masses and sizes.

2. Although the more realistic of these simulations were calibrated against present-day galaxy masses and sizes, the evolution of these quantities was not, and the observed evolutionary trends are also broadly reproduced.
3. Simulated galaxy populations exhibit the diversity of present-day morphologies exemplified by the Hubble sequence, as a natural outcome of the diversity of galaxy assembly histories and the intrinsic properties of their host halos.
4. The observed clustering of galaxies, as a function of stellar mass, galaxy color, and atomic gas content, is reproduced by simulations on the scales they are able to reliably sample and adequately resolve.
5. The simulations reproduce many key galaxy scaling relations, elucidating their origin. A key example is the stellar mass–halo mass relation, which has been shown to emerge primarily in response to gas expulsion in the low-mass regime and throttling of cooling from the circumgalactic medium (CGM) onto the ISM in more massive halos, due to AGN feedback.
6. The simulations also illuminate the origin of scatter about scaling relations, with key examples being the scatter in star-formation rate, and in metallicity, at fixed stellar mass. Both can be explained in terms of the balance of gas flows into and out of galaxies, star formation, and black hole growth.
7. Once armed with a realistic model, examination of partner simulations in which model components are adjusted or toggled has proven an effective approach to illuminating the sensitivity of galaxy properties and observables to physical processes.

FUTURE ISSUES

1. Simulations of the galaxy population still rely on simplified subgrid models to treat unresolved physical processes, including feedback associated with the formation of stars and the growth of black holes. Simulations thus remain distantly removed from *ab initio* theory.
2. Physical processes treated with subgrid models dominate the systematic uncertainty on the properties of simulated galaxies. Because different subgrid models can produce broadly realistic simulated galaxy populations, significant degeneracies remain between state-of-the-art simulation suites, and predictive power is limited.
3. Even in simulations whose subgrid models were designed to minimize their sensitivity to resolution, the convergence behavior of some physical properties (particularly those related to dense gas) can be poor, eroding confidence in outcomes.
4. The detail of simulations of the galaxy population remains relatively poor: The internal and vertical structure of galaxy disks is unrealistic owing to the simplistic subgrid modeling of interstellar gas.
5. There is relatively little quantitative intersimulation consensus concerning the properties of gas flows and the CGM, which are particularly sensitive to the implementation of feedback processes.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We extend our gratitude to the many colleagues and collaborators with whom we have enjoyed illuminating discussions that have helped to shape this review. We thank Romeel Davé, Jon Davies, Alex Hill, Rüdiger Pakmor, and Michael Tremmel for providing data used in the figures, and the EAGLE and IllustrisTNG teams for making their data publicly accessible. We are particularly grateful for Debora Sijacki's contribution to the early phase of assembling this review, and Tim Davis, Matthieu Schaller, and Renske Smit for valuable discussions. We thank Roselyn Lowe-Webb and Luis Ho, respectively our production and scientific editors at the *Annual Review of Astronomy and Astrophysics*, for their patience and encouragement, and Glenda Mahoney for assistance with figures, in particular for creating **Figure 2**. The authors are supported by Royal Society University Research Fellowships.

LITERATURE CITED

- Abadi MG, Navarro JF, Steinmetz M, Eke VR. 2003a. *Ap. J.* 597:21–34
Abadi MG, Navarro JF, Steinmetz M, Eke VR. 2003b. *Ap. J.* 591(2):499–514
Abel T, Bryan GL, Norman ML. 2002. *Science* 295(5552):93–98
Agertz O, Moore B, Stadel J, et al. 2007. *MNRAS* 380(3):963–78
Agertz O, Pontzen A, Read JI, et al. 2020. *MNRAS* 491(2):1656–72
Aguirre A, Hernquist L, Schaye J, et al. 2001. *Ap. J.* 560(2):599–605
Anglés-Alcázar D, Davé R, Faucher-Giguère CA, Özel F, Hopkins PF. 2017. *MNRAS* 464(3):2840–53
Appleby S, Davé R, Kraljic K, Anglés-Alcázar D, Narayanan D. 2020. *MNRAS* 494(4):6053–71
Artale MC, Pedrosa SE, Trayford JW, et al. 2017. *MNRAS* 470(2):1771–87
Bagla JS, Prasad J. 2006. *MNRAS* 370(2):993–1002
Bahé YM, Barnes DJ, Dalla Vecchia C, et al. 2017. *MNRAS* 470(4):4186–208
Bahé YM, Crain RA, Kauffmann G, et al. 2016. *MNRAS* 456:1115–36
Bahé YM, Schaye J, Schaller M, et al. 2022. *MNRAS* 516:167–84
Barnes J, Hut P. 1986. *Nature* 324(6096):446–49
Bate MR, Burkert A. 1997. *MNRAS* 288(4):1060–72
Beckmann RS, Devriendt J, Slyz A, et al. 2017. *MNRAS* 472:949–65
Behroozi P, Wechsler RH, Hearin AP, Conroy C. 2019. *MNRAS* 488(3):3143–94
Benítez-Llambay A, Navarro JF, Abadi MG, et al. 2013. *Ap. J. Lett.* 763(2):L41
Best PN, Heckman TM. 2012. *MNRAS* 421(2):1569–82
Black JH. 1987. *Interstellar Process.* 134:731–44
Boehringer H, Hensler G. 1989. *Astron. Astrophys.* 215:147–49
Bondi H, Hoyle F. 1944. *MNRAS* 104:273
Booth CM, Schaye J. 2009. *MNRAS* 398:53–74
Booth CM, Schaye J. 2010. *MNRAS* 405:L1–5
Borrow J, Kannan R, Garaldi E, et al. 2022a. arXiv:2212.03255
Borrow J, Schaller M, Bahe YM, et al. 2022b. arXiv:2211.08442
Borrow J, Schaller M, Bower RG, Schaye J. 2022c. *MNRAS* 511(2):2367–89
Bower RG, Schaye J, Frenk CS, et al. 2017. *MNRAS* 465:32–44
Bustamante S, Springel V. 2019. *MNRAS* 490(3):4133–53
Butsky IS, Werk JK, Tchernyshyov K, et al. 2022. *Ap. J.* 935(2):69
Carilli CL, Taylor GB. 2002. *Annu. Rev. Astron. Astrophys.* 40:319–48
Carraro G, Lia C, Chiosi C. 1998. *MNRAS* 297(4):1021–40

- Carrier J, Greengard L, Rokhlin V. 1988. *SIAM J. Sci. Stat. Comput.* 9(4):669–86
- Chabrier G. 2003. *Publ. Astron. Soc. Pac.* 115(809):763–95
- Chaves-Montero J, Angulo RE, Schaye J, et al. 2016. *MNRAS* 460(3):3100–18
- Chisholm J, Tremonti CA, Leitherer C, Chen Y. 2016a. *MNRAS* 463:541–56
- Chisholm J, Tremonti CA, Leitherer C, Chen Y, Wofford A. 2016b. *MNRAS* 457(3):3133–61
- Christensen CR, Davé R, Governato F, et al. 2016. *Ap. J.* 824:57
- Churazov E, Sazonov S, Sunyaev R, et al. 2005. *MNRAS* 363:L91–95
- Cicone C, Maiolino R, Sturm E, et al. 2014. *Astron. Astrophys.* 562:A21
- Conroy C, Gunn JE, White M. 2009. *Ap. J.* 699:486–506
- Conroy C, Wechsler RH, Kravtsov AV. 2006. *Ap. J.* 647:201–14
- Correa CA, Schaye J, Clauwens B, et al. 2017. *MNRAS* 472:L45–49
- Crain RA, Bahé YM, Lagos CdP, et al. 2017. *MNRAS* 464(4):4204–26
- Crain RA, Eke VR, Frenk CS, et al. 2007. *MNRAS* 377:41–49
- Crain RA, Schaye J, Bower RG, et al. 2015. *MNRAS* 450(2):1937–61
- Crain RA, Theuns T, Dalla Vecchia C, et al. 2009. *MNRAS* 399(4):1773–94
- Creasey P, Theuns T, Bower RG, Lacey CG. 2011. *MNRAS* 415(4):3706–20
- Cui W, Knebe A, Yepes G, et al. 2018. *MNRAS* 480(3):2898–915
- Curtis M, Sijacki D. 2015. *MNRAS* 454(4):3445–63
- Curtis-Lake E, Carniani S, Cameron A, et al. 2023. *Nat. Astron.* 7:622–32
- Dalgarno A, McCray RA. 1972. *Annu. Rev. Astron. Astrophys.* 10:375
- Dalla Vecchia C, Schaye J. 2008. *MNRAS* 387(4):1431–44
- Dalla Vecchia C, Schaye J. 2012. *MNRAS* 426:140–58
- Davé R. 2008. *MNRAS* 385:147–60
- Davé R, Anglés-Alcázar D, Narayanan D, et al. 2019. *MNRAS* 486(2):2827–49
- Davé R, Crain RA, Stevens ARH, et al. 2020. *MNRAS* 497:146–66
- Davé R, Hernquist L, Katz N, Weinberg DH. 1999. *Ap. J.* 511(2):521–45
- Davé R, Spergel DN, Steinhardt PJ, Wandelt BD. 2001. *Ap. J.* 547(2):574–89
- Davies JJ, Crain RA, Oppenheimer BD, Schaye J. 2020. *MNRAS* 491(3):4462–80
- Davies JJ, Crain RA, Pontzen A. 2021. *MNRAS* 501:236–53
- Davies JJ, Pontzen A, Crain RA. 2022. *MNRAS* 515:1430–43
- Davies LJM, Robotham ASG, Lagos CdP, et al. 2019. *MNRAS* 483(4):5444–58
- De Rossi ME, Bower RG, Font AS, Schaye J, Theuns T. 2017. *MNRAS* 472(3):3354–77
- Di Matteo T, Colberg J, Springel V, Hernquist L, Sijacki D. 2008. *Ap. J.* 676:33–53
- Diemer B, Stevens ARH, Forbes JC, et al. 2018. *Ap. J. Suppl.* 238(2):33
- Donnari M, Pillepich A, Nelson D, et al. 2019. *MNRAS* 485(4):4817–40
- Donnari M, Pillepich A, Nelson D, et al. 2021. *MNRAS* 506(4):4760–80
- Dubois Y, Devriendt J, Slyz A, Teyssier R. 2012. *MNRAS* 420(3):2662–83
- Dubois Y, Peirani S, Pichon C, et al. 2016. *MNRAS* 463(4):3948–64
- Dubois Y, Pichon C, Welker C, et al. 2014. *MNRAS* 444(2):1453–68
- Dubois Y, Teyssier R. 2008. *Astron. Astrophys.* 477:79–94
- Dubois Y, Volonteri M, Silk J, et al. 2015. *MNRAS* 452(2):1502–18
- Efstathiou G. 1992. *MNRAS* 256(2):43P–47P
- Efstathiou G, Davis M, White SDM, Frenk CS. 1985. *Ap. J. Suppl.* 57:241–60
- Faber SM, Jackson RE. 1976. *Ap. J.* 204:668–83
- Faucher-Giguère CA, Lidz A, Zaldarriaga M, Hernquist L. 2009. *Ap. J.* 703(2):1416–43
- Faucher-Giguère CA, Oh P. 2023. *Annu. Rev. Astron. Astrophys.* 61:131–95
- Feldmann R, Quataert E, Faucher-Giguère CA, et al. 2023. *MNRAS* 522(3):3831–60
- Feng Y, Di-Matteo T, Croft RA, et al. 2016. *MNRAS* 455(3):2778–91
- Ferland GJ, Chatzikos M, Guzmán F, et al. 2017. *Rev. Mex. Astron. Astrofis.* 53:385–438
- Ferrero I, Navarro JF, Abadi MG, et al. 2017. *MNRAS* 464:4736–46
- Fielding DB, Ostriker EC, Bryan GL, Jermyn AS. 2020. *Ap. J. Lett.* 894(2):L24
- François P, Matteucci F, Cayrel R, et al. 2004. *Astron. Astrophys.* 421:613–21

- Frenk CS, White SDM, Bode P, et al. 1999. *Ap. J.* 525(2):554–82
- Fumagalli M, Hennawi JF, Prochaska JX, et al. 2014. *Ap. J.* 780:74
- Furlanetto SR, Schaye J, Springel V, Hernquist L. 2005. *Ap. J.* 622:7–27
- Furlong M, Bower RG, Crain RA, et al. 2017. *MNRAS* 465:722–38
- Furlong M, Bower RG, Theuns T, et al. 2015. *MNRAS* 450(4):4486–504
- Gallazzi A, Charlot S, Brinchmann J, White SDM, Tremonti CA. 2005. *MNRAS* 362:41–58
- Garnier E, Adams N, Sagaut P. 2009. *Large Eddy Simulation for Compressible Flows*. Berlin: Springer
- Genel S, Bryan GL, Springel V, et al. 2019. *Ap. J.* 871:21
- Genel S, Nelson D, Pillepich A, et al. 2018. *MNRAS* 474(3):3976–96
- Genel S, Vogelsberger M, Springel V, et al. 2014. *MNRAS* 445:175–200
- Gerritsen JPE, Icke V. 1997. *Astron. Astrophys.* 325:972–86
- Gingold RA, Monaghan JJ. 1977. *MNRAS* 181:375–89
- Gnedin NY, Hollon N. 2012. *Ap. J. Suppl.* 202(2):13
- Gnedin NY, Tassis K, Kravtsov AV. 2009. *Ap. J.* 697:55–67
- Governato F, Mayer L, Wadsley J, et al. 2004. *Ap. J.* 607(2):688–96
- Governato F, Willman B, Mayer L, et al. 2007. *MNRAS* 374(4):1479–94
- Greggio L. 2005. *Astron. Astrophys.* 441(3):1055–78
- Greggio L, Renzini A. 1983. *Astron. Astrophys.* 118(2):217–22
- Greif TH, Glover SCO, Bromm V, Klessen RS. 2009. *MNRAS* 392(4):1381–87
- Guedes J, Callegari S, Madau P, Mayer L. 2011. *Ap. J.* 742(2):76
- Gunn JE, Gott J. Richard I. 1972. *Ap. J.* 176:1
- Guo Q, Gonzalez-Perez V, Guo Q, et al. 2016. *MNRAS* 461(4):3457–82
- Gutcke TA, Pakmor R, Naab T, Springel V. 2022. *MNRAS* 513:1372–85
- Haardt F, Madau P. 1996. *Ap. J.* 461:20
- Haardt F, Madau P. 2001. Modelling the UV/X-ray cosmic background with CUBA. In *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*, ed. DM Neumann, JTV Tran.
- Haardt F, Madau P. 2012. *Ap. J.* 746(2):125
- Haas MR, Schaye J, Jeesson-Daniel A. 2012. *MNRAS* 419(3):2133–46
- Habouzit M, Li Y, Somerville RS, et al. 2021. *MNRAS* 503(2):1940–75
- Hahn O, Abel T. 2011. *MNRAS* 415(3):2101–21
- Hahn O, Rampf C, Uhlemann C. 2021. *MNRAS* 503:426–45
- Haider M, Steinhäuser D, Vogelsberger M, et al. 2016. *MNRAS* 457(3):3024–35
- Henden NA, Puchwein E, Shen S, Sijacki D. 2018. *MNRAS* 479(4):5385–412
- Herzog G, Benítez-Llambay A, Fumagalli M. 2022. *MNRAS* 518:6305–17
- Hirschmann M, Dolag K, Saro A, et al. 2014. *MNRAS* 442(3):2304–24
- Hockney RW, Eastwood JW. 1981. *Computer Simulation Using Particles*. Boca Raton, FL: CRC Press
- Hopkins PF. 2013. *MNRAS* 428(4):2840–56
- Hopkins PF. 2015. *MNRAS* 450:53–110
- Hopkins PF, Kereš D, Oñorbe J, et al. 2014. *MNRAS* 445:581–603
- Hopkins PF, Wetzel A, Kereš D, et al. 2018. *MNRAS* 480:800–63
- Huang S, Katz N, Scannapieco E, et al. 2020. *MNRAS* 497(3):2586–604
- Huertas-Company M, Rodríguez-Gómez V, Nelson D, et al. 2019. *MNRAS* 489(2):1859–79
- Hummels CB, Smith BD, Hopkins PF, et al. 2019. *Ap. J.* 882(2):156
- Icaza-Lizaola M, Bower RG, Norberg P, et al. 2021. *MNRAS* 507(3):4584–602
- Jahnke K, Macciò AV. 2011. *Ap. J.* 734(2):92
- Jeans JH. 1928. *Astronomy and Cosmogony*. Cambridge, UK: Cambridge Univ. Press
- Jenkins A. 2010. *MNRAS* 403(4):1859–72
- Jenkins A. 2013. *MNRAS* 434(3):2094–120
- Ji S, Chan TK, Hummels CB, et al. 2020. *MNRAS* 496(4):4221–38
- Jo Y, Genel S, Wandelt B, et al. 2023. *Ap. J.* 944:67
- Kang X, Jing YP, Mo HJ, Börner G. 2005. *Ap. J.* 631:21–40
- Kannan R, Springel V, Hernquist L, et al. 2022. arXiv:2210.10066

- Katz H, Ramsoy M, Rosdahl J, et al. 2020. *MNRAS* 494(2):2200–20
- Katz N, Gunn JE. 1991. *Ap. J.* 377:365
- Katz N, Weinberg DH, Hernquist L. 1996. *Ap. J. Suppl.* 105:19
- Kaviraj S, Laigle C, Kimm T, et al. 2017. *MNRAS* 467(4):4739–52
- Kay ST, Thomas PA, Theuns T. 2003. *MNRAS* 343(2):608–18
- Keller BW, Wadsley J, Benincasa SM, Couchman HMP. 2014. *MNRAS* 442(4):3013–25
- Keller BW, Wadsley JW, Wang L, Kruijssen JMD. 2019. *MNRAS* 482(2):2244–61
- Kennicutt RC Jr. 1998. *Ap. J.* 498(2):541–52
- Khandai N, Di Matteo T, Croft R, et al. 2015. *MNRAS* 450(2):1349–74
- Kim CG, Ostriker EC. 2017. *Ap. J.* 846(2):133
- Kim J-h, Abel T, Agertz O, et al. 2014. *Ap. J. Suppl.* 210:14
- King A, Nealon R. 2021. *MNRAS* 502:L1–5
- Kormendy J, Ho LC. 2013. *Annu. Rev. Astron. Astrophys.* 51:511–653
- Kugel R, Borrow J. 2022. *J. Open Source Softw.* 7(72):4240
- Labbé I, van Dokkum P, Nelson E, et al. 2023. *Nature* 616:266–69
- Lackner CN, Cen R, Ostriker JP, Joung MR. 2012. *MNRAS* 425:641–56
- Lagos CdP, Crain RA, Schaye J, et al. 2015. *MNRAS* 452(4):3815–37
- Lagos CdP, Theuns T, Schaye J, et al. 2016. *MNRAS* 459(3):2632–50
- Larson RB. 1974. *MNRAS* 169:229–46
- Leitherer J, Schaerer D, Goldader JD, et al. 1999. *Ap. J. Suppl.* 123:3–40
- Leja J, Carnall AC, Johnson BD, Conroy C, Speagle JS. 2019. *Ap. J.* 876:3
- Lewis A, Challinor A. 2011. CAMB: Code for Anisotropies in the Microwave Background. Astrophysics Source Code Library, record ascl:1102.026
- Li C, White SDM. 2009. *MNRAS* 398(4):2177–87
- Lopez LA, Krumholz MR, Bolatto AD, Prochaska JX, Ramirez-Ruiz E. 2011. *Ap. J.* 731(2):91
- Lucy LB. 1977. *Astron. J.* 82:1013–24
- Ludlow AD, Schaye J, Schaller M, Richings J. 2019. *MNRAS* 488:L123–28
- Machacek ME, Bryan GL, Abel T. 2001. *Ap. J.* 548(2):509–21
- Madau P, Dickinson M. 2014. *Annu. Rev. Astron. Astrophys.* 52:415–86
- Maiolino R, Gallerani S, Neri R, et al. 2012. *MNRAS* 425:L66–70
- Mannucci F, Della Valle M, Panagia N. 2006. *MNRAS* 370(2):773–83
- Maoz D, Mannucci F, Nelemans G. 2014. *Annu. Rev. Astron. Astrophys.* 52:107–70
- Marasco A, Crain RA, Schaye J, et al. 2016. *MNRAS* 461(3):2630–49
- Martín-Navarro I, Pillepich A, Nelson D, et al. 2021. *Nature* 594(7862):187–90
- Matthee J, Schaye J. 2019. *MNRAS* 484:915–32
- McAlpine S, Helly JC, Schaller M, et al. 2016. *Astron. Comput.* 15:72–89
- McCarthy IG, Schaye J, Bird S, Le Brun AMC. 2017. *MNRAS* 465(3):2936–65
- McCarthy IG, Schaye J, Font AS, et al. 2012. *MNRAS* 427:379–92
- McCourt M, Oh SP, O’Leary R, Madigan AM. 2018. *MNRAS* 473(4):5407–31
- McGaugh SS, Schombert JM, Bothun GD, de Blok WJG. 2000. *Ap. J. Lett.* 533(2):L99–102
- McKee CF, Ostriker JP. 1977. *Ap. J.* 218:148–69
- Menon H, Wesolowski L, Zheng G, et al. 2015. *Comput. Astrophys. Cosmol.* 2:1
- Merloni A, Heinz S. 2008. *MNRAS* 388(3):1011–30
- Metha B, Trenti M. 2020. *MNRAS* 495:266–77
- Mihos JC, Hernquist L. 1994. *Ap. J.* 437:611
- Miralda-Escudé J. 2005. *Ap. J. Lett.* 620(2):L91–94
- Mitchell NL, McCarthy IG, Bower RG, Theuns T, Crain RA. 2009. *MNRAS* 395:180–96
- Mitchell PD, Schaye J, Bower RG, Crain RA. 2020. *MNRAS* 494(3):3971–97
- Monaghan JJ. 1985. *J. Comput. Phys.* 60(2):253–62
- Monaghan JJ. 1997. *J. Comput. Phys.* 136(2):298–307
- Muratov AL, Kereš D, Faucher-Giguère CA, et al. 2015. *MNRAS* 454(3):2691–713
- Navarro JF, Benz W. 1991. *Ap. J.* 380:320

- Navarro JF, Frenk CS, White SDM. 1995. *MNRAS* 275:56–66
- Navarro JF, Steinmetz M. 1997. *Ap. J.* 478:13–28
- Navarro JF, Steinmetz M. 2000. *Ap. J.* 538(2):477–88
- Navarro JF, White SDM. 1993. *MNRAS* 265:271
- Nelson D, Kauffmann G, Pillepich A, et al. 2018a. *MNRAS* 477:450–79
- Nelson D, Pillepich A, Genel S, et al. 2015. *Astron. Comput.* 13:12–37
- Nelson D, Pillepich A, Springel V, et al. 2018b. *MNRAS* 475:624–47
- Nelson D, Springel V, Pillepich A, et al. 2019. *Comput. Astrophys. Cosmol.* 6:2
- Noeske KG, Weiner BJ, Faber SM, et al. 2007. *Ap. J. Lett.* 660:L43–46
- Okamoto T, Eke VR, Frenk CS, Jenkins A. 2005. *MNRAS* 363(4):1299–314
- Omukai K, Tsuribe T, Schneider R, Ferrara A. 2005. *Ap. J.* 626(2):627–43
- Oppenheimer BD, Crain RA, Schaye J, et al. 2016. *MNRAS* 460(2):2157–79
- Oppenheimer BD, Davé R. 2008. *MNRAS* 387(2):577–600
- Oppenheimer BD, Davé R, Kereš D, et al. 2010. *MNRAS* 406(4):2325–38
- Oppenheimer BD, Schaye J, Crain RA, Werk JK, Richings AJ. 2018. *MNRAS* 481:835–59
- Orlando S, Peres G, Reale F, et al. 2005. *Astron. Astrophys.* 444(2):505–19
- O’Shea BW, Nagamine K, Springel V, Hernquist L, Norman ML. 2005. *Ap. J. Suppl.* 160:1–27
- Pakmor R, Springel V, Coles JP, et al. 2022. arXiv:2210.10060
- Pallottini A, Ferrara A, Gallerani S, et al. 2017. *MNRAS* 465(3):2540–58
- Pasha I, Mandelker N, van den Bosch FC, Springel V, van de Voort F. 2023. *MNRAS* 520(2):2692–708
- Peeples MS, Corlies L, Tumlinson J, et al. 2019. *Ap. J.* 873(2):129
- Péroux C, Howk JC. 2020. *Annu. Rev. Astron. Astrophys.* 58:363–406
- Péroux C, Nelson D, van de Voort F, et al. 2020. *MNRAS* 499(2):2462–73
- Pfeffer J, Kruijssen JMD, Crain RA, Bastian N. 2018. *MNRAS* 475(4):4309–46
- Pillepich A, Nelson D, Hernquist L, et al. 2018a. *MNRAS* 475:648–75
- Pillepich A, Springel V, Nelson D, et al. 2018b. *MNRAS* 473(3):4077–106
- Piotrowska JM, Bluck AFL, Maiolino R, Peng Y. 2022. *MNRAS* 512:1052–90
- Ploeckinger S, Schaye J. 2020. *MNRAS* 497(4):4857–83
- Pontzen A, Governato F. 2012. *MNRAS* 421(4):3464–71
- Portinari L, Moretti A, Chiosi C, Sommer-Larsen J. 2004. *Ap. J.* 604(2):579–95
- Power C, Navarro JF, Jenkins A, et al. 2003. *MNRAS* 338:14–34
- Price DJ. 2008. *J. Comput. Phys.* 227(24):10040–57
- Price DJ. 2012. *J. Comput. Phys.* 231(3):759–94
- Rahmati A, Pawlik AH, Raičević M, Schaye J. 2013. *MNRAS* 430(3):2427–45
- Rahmati A, Schaye J, Bower RG, et al. 2015. *MNRAS* 452(2):2034–56
- Reed DS, Bower R, Frenk CS, Jenkins A, Theuns T. 2007. *MNRAS* 374:2–15
- Richings AJ, Schaye J, Oppenheimer BD. 2014. *MNRAS* 440(4):3349–69
- Ritchie BW, Thomas PA. 2001. *MNRAS* 323(3):743–56
- Robertson A, Massey R, Eke V. 2017. *MNRAS* 467(4):4719–30
- Robertson BE, Kravtsov AV. 2008. *Ap. J.* 680(2):1083–111
- Rodriguez-Gomez V, Snyder GF, Lotz JM, et al. 2019. *MNRAS* 483(3):4140–59
- Romano D, Chiappini C, Matteucci F, Tosi M. 2005. *Astron. Astrophys.* 430:491–505
- Rosas-Guevara Y, Tissera P, Lagos CdP, Paillas E, Padilla N. 2022. *MNRAS* 517:712–31
- Rosen AL, Lopez LA, Krumholz MR, Ramirez-Ruiz E. 2014. *MNRAS* 442(3):2701–16
- Rosswog S. 2020. *MNRAS* 498(3):4230–55
- Saintonge A, Catinella B. 2022. *Annu. Rev. Astron. Astrophys.* 60:319–61
- Sales LV, Navarro JF, Oman K, et al. 2017. *MNRAS* 464(2):2419–28
- Scannapieco C, Wadepuhl M, Parry OH, et al. 2012. *MNRAS* 423:1726–49
- Schaller M, Frenk CS, Bower RG, et al. 2015. *MNRAS* 451(2):1247–67
- Schaller M, Gonnet P, Chalk ABG, Draper PW. 2016. In *PASC ’16: Proceedings of the Platform for Advanced Scientific Computing Conference, Lausanne, Switz., June 8–10*. Art. no. 2. <https://doi.org/10.1145/2929908.2929916>

- Schaye J, Crain RA, Bower RG, et al. 2015. *MNRAS* 446:521–54
- Schaye J, Dalla Vecchia C. 2008. *MNRAS* 383(3):1210–22
- Schaye J, Dalla Vecchia C, Booth CM, et al. 2010. *MNRAS* 402(3):1536–60
- Schirber M, Miralda-Escudé J, McDonald P. 2004. *Ap. J.* 610:105–16
- Schmidt M. 1959. *Ap. J.* 129:243
- Sembolini F, Yepes G, Pearce FR, et al. 2016. *MNRAS* 457:4063–80
- Semenov VA, Kravtsov AV, Gnedin NY. 2016. *Ap. J.* 826(2):200
- Shakura NI, Sunyaev RA. 1973. *Astron. Astrophys.* 24:337–55
- Shlosman I, Begelman MC, Frank J. 1990. *Nature* 345(6277):679–86
- Sijacki D, Springel V, Di Matteo T, Hernquist L. 2007. *MNRAS* 380(3):877–900
- Simpson CM, Grand RJJ, Gómez FA, et al. 2018. *MNRAS* 478:548–67
- Smith B, Sigurdsson S, Abel T. 2008. *MNRAS* 385(3):1443–54
- Smith BD, Bryan GL, Glover SCO, et al. 2017. *MNRAS* 466(2):2217–34
- Smith MJ, Geach JE. 2023. *R. Soc. Open Sci.* 10(5):221454
- Snyder GF, Torrey P, Lotz JM, et al. 2015. *MNRAS* 454(2):1886–908
- Sommer-Larsen J, Gelato S, Vedel H. 1999. *Ap. J.* 519(2):501–12
- Sorini D, Davé R, Cui W, Appleby S. 2022. *MNRAS* 516:883–906
- Sparre M, Hayward CC, Feldmann R, et al. 2017. *MNRAS* 466:88–104
- Springel V. 2010a. *Annu. Rev. Astron. Astrophys.* 48:391–430
- Springel V. 2010b. *MNRAS* 401(2):791–851
- Springel V, Di Matteo T, Hernquist L. 2005. *MNRAS* 361(3):776–94
- Springel V, Hernquist L. 2003. *MNRAS* 339(2):289–311
- Springel V, Pakmor R, Pillepich A, et al. 2018. *MNRAS* 475:676–98
- Steinmetz M, Mueller E. 1994. *Astron. Astrophys.* 281(3):L97–100
- Stevens ARH, Diemer B, Lagos CdP, et al. 2019. *MNRAS* 483(4):5334–54
- Stinson G, Seth A, Katz N, et al. 2006. *MNRAS* 373(3):1074–90
- Sun M, Voit GM, Donahue M, et al. 2009. *Ap. J.* 693(2):1142–72
- Sutherland RS, Dopita MA. 1993. *Ap. J. Suppl.* 88:253
- Tasker EJ, Brunino R, Mitchell NL, et al. 2008. *MNRAS* 390(3):1267–81
- Teyssier R. 2015. *Annu. Rev. Astron. Astrophys.* 53:325–64
- Theis C, Burkert A, Hensler G. 1992. *Astron. Astrophys.* 265(2):465–77
- Theuns T, Leonard A, Efstathiou G, Pearce FR, Thomas PA. 1998. *MNRAS* 301(2):478–502
- Thob ACR, Crain RA, McCarthy IG, et al. 2019. *MNRAS* 485:972–87
- Tornatore L, Borgani S, Dolag K, Matteucci F. 2007. *MNRAS* 382(3):1050–72
- Torrey P, Vogelsberger M, Marinacci F, et al. 2019. *MNRAS* 484(4):5587–607
- Trayford JW, Camps P, Theuns T, et al. 2017. *MNRAS* 470:771–99
- Trayford JW, Theuns T, Bower RG, et al. 2015. *MNRAS* 452(3):2879–96
- Trebitsch M, Blaizot J, Rosdahl J, Devriendt J, Slyz A. 2017. *MNRAS* 470:224–39
- Tremmel M, Governato F, Volonteri M, Quinn TR. 2015. *MNRAS* 451(2):1868–74
- Tremmel M, Karcher M, Governato F, et al. 2017. *MNRAS* 470:1121–39
- Tremmel M, Quinn TR, Ricarte A, et al. 2019. *MNRAS* 483(3):3336–62
- Tremonti CA, Heckman TM, Kauffmann G, et al. 2004. *Ap. J.* 613(2):898–913
- Truelove JK, Klein RL, McKee CF, et al. 1997. *Ap. J. Lett.* 489(2):L179–83
- Tully RB, Fisher JR. 1977. *Astron. Astrophys.* 54:661–73
- Tumlinson J, Peebles MS, Werk JK. 2017. *Annu. Rev. Astron. Astrophys.* 55:389–432
- Tumlinson J, Thom C, Werk JK, et al. 2011. *Science* 334(6058):948
- Uhlig M, Pfrommer C, Sharma M, et al. 2012. *MNRAS* 423(3):2374–96
- van de Voort F, Bahé YM, Bower RG, et al. 2017. *MNRAS* 466(3):3460–71
- van de Voort F, Bieri R, Pakmor R, et al. 2021. *MNRAS* 501(4):4888–902
- van de Voort F, Pakmor R, Grand RJJ, et al. 2020. *MNRAS* 494(4):4867–83
- van de Voort F, Schaye J. 2012. *MNRAS* 423(4):2991–3010
- van de Voort F, Springel V, Mandelker N, van den Bosch FC, Pakmor R. 2019. *MNRAS* 482:L85–89

van der Wel A, Franx M, van Dokkum PG, et al. 2014. *Ap. J.* 788:28

Vandenbroucke B, De Rijcke S. 2016. *Astron. Comput.* 16:109–30

Vasiliev EO. 2011. *MNRAS* 414(4):3145–57

Vikhlinin A, Markevitch M, Murray SS, et al. 2005. *Ap. J.* 628(2):655–72

Villaescusa-Navarro F, Anglés-Alcázar D, Genel S, et al. 2021. *Ap. J.* 915:71

Villaescusa-Navarro F, Genel S, Anglés-Alcázar D, et al. 2023. *Ap. J. Suppl.* 265:54

Vogelsberger M, Genel S, Sijacki D, et al. 2013. *MNRAS* 436(4):3031–67

Vogelsberger M, Genel S, Springel V, et al. 2014. *MNRAS* 444(2):1518–47

Volonteri M, Habouzit M, Colpi M. 2021. *Nat. Rev. Phys.* 3(11):732–43

Wadsley JW, Keller BW, Quinn T. 2017. *MNRAS* 471:2357–69

Wadsley JW, Veeravalli G, Couchman HMP. 2008. *MNRAS* 387:427–38

Weil ML, Eke VR, Efstathiou G. 1998. *MNRAS* 300(3):773–89

Weinberger R, Springel V, Hernquist L, et al. 2017. *MNRAS* 465(3):3291–308

Weinberger R, Springel V, Pakmor R. 2020. *Ap. J. Suppl.* 248(2):32

Wendland H. 1995. *Adv. Comput. Math.* 4:389–96

Werk JK, Prochaska JX, Thom C, et al. 2013. *Ap. J. Suppl.* 204(2):17

Werk JK, Prochaska JX, Tumlinson J, et al. 2014. *Ap. J.* 792:8

Wheeler C, Hopkins PF, Pace AB, et al. 2019. *MNRAS* 490(3):4447–63

White SDM. 1994. arXiv:astro-ph/9410043

Wiersma RPC, Schaye J, Smith BD. 2009a. *MNRAS* 393:99–107

Wiersma RPC, Schaye J, Theuns T, Dalla Vecchia C, Tornatore L. 2009b. *MNRAS* 399(2):574–600

Wijers NA, Schaye J, Oppenheimer BD, Crain RA, Nicastro F. 2019. *MNRAS* 488(2):2947–69

Wise JH, Cen R. 2009. *Ap. J.* 693:984–99

Wright AH, Robotham ASG, Driver SP, et al. 2017. *MNRAS* 470:283–302

Wright RJ, Lagos CdP, Power C, Mitchell PD. 2020. *MNRAS* 498(2):1668–92

Wright RJ, Lagos CdP, Power C, et al. 2022. *MNRAS* 516(2):2891–912

Yang H, Gao L, Frenk CS, et al. 2023. *MNRAS* 518(4):5253–59

Yepes G, Kates R, Khokhlov A, Klypin A. 1997. *MNRAS* 284:235–56

Yun K, Pillepich A, Zinger E, et al. 2019. *MNRAS* 483:1042–66

Zavala J, Frenk CS, Bower R, et al. 2016. *MNRAS* 460(4):4466–82

Zuo L. 1992. *MNRAS* 258:36–44