Dissertation

submitted to the

Combined Faculty of Mathematics, Engineering and Natural Sciences

of Heidelberg University, Germany

for the degree of

Doctor of Natural Sciences

Put forward by

M.Sc. Rahul Ramesh

born in: Bengaluru (India)

Oral examination: 24-04-2025

Deciphering galaxy evolution through the baryon cycle and circumgalactic medium in cosmological simulations

Referees: Dr. Dylan Nelson

Prof. Dr. Bjoern Malte Schaefer

Abstract

Observations and theoretical studies suggest that galaxies are surrounded by a halo of diffuse gas that extends far beyond the extent of the central stellar component. This region, known as the circumgalactic medium (CGM), serves as the area through which gas from larger scales accrete due to the gravitationally-driven growth of cosmic structures. Additionally, it contains gas that has been expelled from the galaxy due to feedback processes, as well as gas that circulates in the halo through recycling or fountain flows. The CGM is thus believed to be critically linked to the evolution of galaxies.

The non-trivial interactions of these various physical processes result in a complex structure within the CGM: while most of the volume is dominated by a warm-hot phase, there are also clouds of cooler gas that coexist. Despite significant progress over the recent past, many open questions remain regarding the origin of the multi-phase nature of this gaseous reservoir, its properties, and the role of different processes in shaping its existence. In this thesis, we explore various such puzzles using cosmological magnetohydrodynamical simulations run with AREPO and the IllustrisTNG galaxy formation model.

We begin by analyzing the publicly available TNG50 simulation, with a particular focus on Milky Way-like (MW-like) galaxies in most cases. Our findings suggest that, among other processes, feedback driven by galactic processes may significantly impact the CGM, including its overall temperature and velocity structure, the number of cold clouds, and the angular structure of magnetic fields. In the later parts of this thesis, we introduce and analyze the new GIBLE suite of simulations, also run with the same IllustrisTNG galaxy physics model, but exclusively simulating MW-like galaxies at ultra-high CGM gas mass resolutions. As a first scientific exploration with GIBLE, we study the draping of magnetic field lines around cold clouds – a phenomenon simply absent in simulations run at lower resolutions, highlighting the power of these new numerical experiments.

Zusammenfassung

Beobachtungen und theoretische Studien deuten darauf hin, dass Galaxien von einem Halo aus diffusem Gas umgeben sind, der sich weit über die Ausdehnung der zentralen stellaren Komponente der Galaxien hinaus erstreckt. Diese Region, die als zirkumgalaktisches Medium (CGM) bezeichnet wird, ist das Gebiet, durch welches das Gas gravitationsbedingt aus größeren Skalen von kosmischen Strukturen akkretiert wird. Weiterhin enthält das CGM Gas, welches aufgrund von Rückkopplungsprozessen aus der Galaxie ausgestoßen wurde, sowie Gas, das im Halo durch Recycling oder Fontänenströme zirkuliert. Man geht daher davon aus, dass das CGM eine entscheidende Rolle bei der Entwicklung von Galaxien spielt.

Die nicht trivialen Wechselwirkungen dieser verschiedenen physikalischen Prozesse führen zu einer komplexen Struktur innerhalb des CGM: Während der größte Teil des Volumens von einer warmen/heißen Gasphase dominiert wird, gibt es auch Wolken mit kühlerem Gas, die räumlich koexistieren. Trotz der bedeutenden Fortschritte in der jüngsten Vergangenheit bleiben viele Fragen offen, die den Ursprung der mehrphasigen Natur dieses Gasreservoirs, seine Eigenschaften und die Rolle der verschiedenen formenden Prozesse betreffen. In dieser Arbeit untersuchen wir verschiedene dieser Rätsel anhand von kosmologischen magnetohydrodynamischen Simulationen, die mit AREPO und dem IllustrisTNG Galaxienbildungsmodell durchgeführt wurden.

Wir beginnen mit der Analyse der öffentlich zugänglichen TNG50 Simulation, wobei wir uns in den meisten Fällen auf milchstraßenähnliche (MW-ähnliche) Galaxien konzentrieren. Unsere Ergebnisse deuten darauf hin, dass neben anderen Prozessen auch Rückkopplungen, die von galaktischen Prozessen angetrieben werden, einen erheblichen Einfluss auf das CGM haben können, einschließlich seiner allgemeinen Temperatur- und Geschwindigkeitsstruktur, der Anzahl kalter Wolken und der Winkelstruktur von Magnetfeldern. In den weiteren Teilen dieser Arbeit wird die neue GIBLE-Simulationsreihe vorgestellt und analysiert, die ebenfalls mit demselben IllustrisTNG Galaxien-

physikmodell durchgeführt wird, aber ausschließlich MW-ähnliche Galaxien mit sehr hoher Auflösung des CGM simuliert. Als erste wissenschaftliche Untersuchung mit GIBLE untersuchen wir die Topologie der Magnetfeldlinien um kalte Wolken - ein Phänomen, das in Simulationen mit niedrigeren Auflösungen nicht vorkommt und die Notwendigkeit und Aussagekraft dieser neuen numerischen Experimente unterstreicht.

To my parents and grandparents,
for their unconditional love and constant support.

Contents

Abstract **v**

Zusammenfassung **vii**

Contents **xi**

List of Figures **xiii**

1 Introduction **1**
1.1 Galaxies and their surrounding halos: A brief history 1
1.2 The Circumgalactic Medium and Its Physics 3
1.2.1 Radiative Cooling 4
1.2.2 Photo-Ionization 6
1.2.3 Galactic Feedback Processes – I: Supernovae 6
1.2.4 Galactic Feedback Processes – II: Supermassive Black Holes 8
1.2.5 Magnetic Fields 10
1.2.6 Cosmic Rays 10
1.2.7 Smooth IGM Accretion 12
1.2.8 Clumpy Satellite Accretion 13
1.2.9 Fluid Instabilities 14
1.3 Numerical Techniques 14
1.3.1 Gravity and Collisionless Components 15
1.3.2 Hydrodynamics 17
1.3.3 The IllustrisTNG Galaxy Formation Model 20
1.3.4 Initial Conditions 22
1.4 Structure of this thesis 22

2 The Circumgalactic Medium of Milky Way-like Galaxies in the TNG50 Simulation – I: Halo Gas Properties and the Role of SMBH Feedback **25**
2.1 Statement of contribution 25

3 The Circumgalactic Medium of Milky Way-like Galaxies in the TNG50 Simulation – II: Cold, Dense Gas Clouds and High-Velocity Cloud Analogs **51**

3.1 Statement of contribution 51

4 Azimuthal Anisotropy of Magnetic Fields in the Circumgalactic Medium Driven by Galactic Feedback Processes **73**

4.1 Statement of contribution 73

5 Zooming in on the Circumgalactic Medium with GIBLE: Resolving small-scale gas structure in cosmological simulations **85**

5.1 Statement of contribution 85

6 Zooming in on the Circumgalactic Medium with GIBLE: the Topology and Draping of Magnetic Fields around Cold Clouds **107**

6.1 Statement of contribution 107

7 Summary and Outlook **115**

List of Figures

1.1 Edwin Hubble's 1923 image of the *Andromeda Nebula*. 2
1.2 A visualization of gas flows around a simulated Milky Way-like galaxy. 4
1.3 A couple of illustrative examples of how cosmic rays may impact the evolution of galaxies and their CGMs. 11
1.4 A visualisation of refinement and de-refinement operations in AREPO. 20

Chapter 1

Introduction

1.1 Galaxies and their surrounding halos: A brief history

About a century ago, in October 1923, when Edwin Hubble put forward his findings suggesting that the *Andromeda Nebula* is about a million light-years away, it brought about a radical change in our understanding of the universe and the distribution of matter therein: galaxies other than our Milky Way exist. A few years later, in 1929, another seminal discovery of his showed that the outward radial velocities of distant galaxies are roughly proportional to their distance[1], in broad agreement with theories of an expanding universe. Through various observational, theoretical, and technical advancements over the years, it is now understood that our universe not only consists of many hundreds of billions of galactic structures assembled from stars, gas, and black holes, but also contains two additional mysterious components: (i) dark energy, which drives the accelerated (late-time) expansion of the universe, and (ii) dark matter, a peculiar substance that does not interact with light and is hence not visible, but outweighs *normal* baryonic matter by a factor of $\sim$5.6 in the mass budget of the universe (e.g. [1]). Altogether, these different matter and energy components form the basis of the Λ-CDM[2] model of cosmology.

While many observational results have provided independent evidence

[1] Although a similar result was also independently published by Georges Lemaître a couple of years earlier, in 1927, it did not gain widespread recognition at the time.

[2] Λ, the cosmological constant introduced by Einstein in his equations of general relativity, is a commonly accepted candidate for dark energy, although it was first proposed by him to obtain a solution for a static universe – in his own words, his 'biggest blunder'! CDM, or Cold Dark Matter, the widely accepted form of dark matter, is collisionless and hence imparts no pressure.

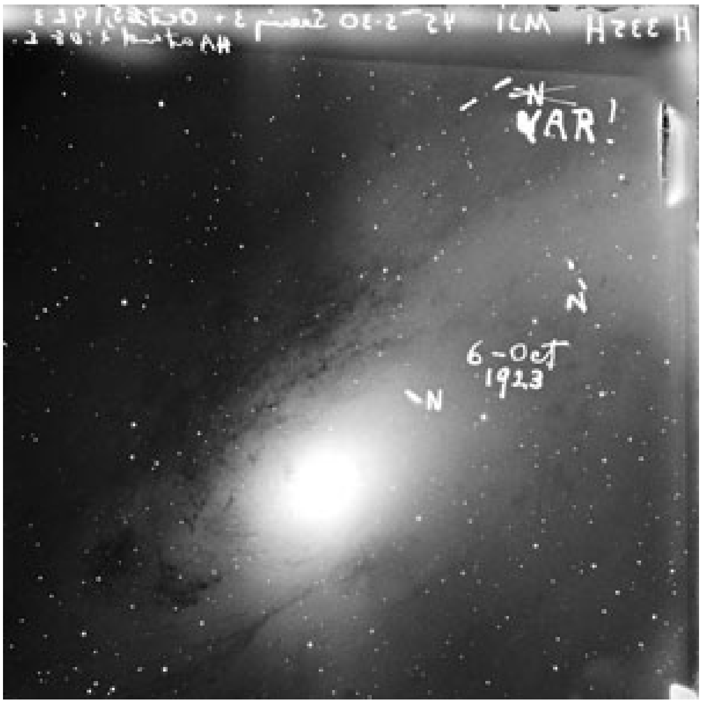

Figure 1.1. Edwin Hubble's 1923 image of the *Andromeda Nebula*. He estimated a distance of a few million light years, putting forward a seminal idea that galaxies other than our Milky Way exist. (Credits: Hubble Heritage Team)

for the existence of dark matter, the most elegant is arguably the fact that large-scale structure exists in the first place: since baryonic matter interacts with and is consequently heated by electromagnetic radiation, which is understood to dominate the mass-energy budget during the earliest epochs of the universe, associated primordial small-scale density perturbations cannot condense into the observed cosmic structure in the absence of pre-assembled *halos* of dark matter within the current lifetime of the universe [2, 3]. Even though the mass distribution of such extended dark matter halos cannot be directly imaged, yet it has been mapped at various scales using galactic rotation curves [4, 5], measurements of velocity dispersions of elliptical galaxies [6, 7], and through the phenomena of gravitational lensing [8, 9].

Although halos surrounding galaxies are dominated by dark matter in terms of their mass content, they also contain gas. The first evidence for this was put forth in the late-1950s when absorption line studies to background extraplanar stars in the Milky Way halo suggested the presence of Na I and Ca II [10]. The exploration of gas around external galaxies soon followed suit with Schmidt's 1963 discovery of quasars [11], revealing that these gaseous reservoirs are likely complex in terms of their density, temperature, and metallicity structure [12, 13, 14]. Fast forward to the 2000s, large-volume surveys began inferring the fraction of baryonic mass in galaxies to be smaller than expected within the realm of the Λ-CDM model [15, 16], and these surrounding halos were proposed to serve as hosts for at least a fraction of the remainder of the *missing baryons*. Soon, the concept of the circumgalactic medium (CGM) emerged, which is now recognized as a critical component of the interface between galaxies and their surrounding environments, and is considered to play a key role in their evolution.

1.2 The Circumgalactic Medium and Its Physics

Though loosely defined, the CGM is understood to be the volume sandwiched between the galaxy and the intergalactic medium (IGM). It is thus the region through which gas accreting from larger scales must flow, while simultaneously hosting gas ejected from the galaxy in the past via feedback processes, along with circulating fountain flows in the halo [17, 18]. As such, the CGM not only encodes information about the galaxy's past activity, but also potential future in-fluxes down towards the center of the halo.

Figure 1.2 provides an exemplary visualization of such gas flows around a Milky Way-like galaxy from a state-of-the-art cosmological magnetohydrodynamical simulation[3]. The circles are drawn at [0.15, 1.00] $\mathrm{R_{200c}}$, demarcating the two boundaries of the CGM[4]. The colors in the background image correspond to the radial velocity of gas, as shown by the colorbar at the bottom. Outflows triggered by galactic feedback processes (yellow; [19, 20]) intersect large-scale accretion driven by gravity (reddish-brown), both in a smooth fashion from the IGM [21, 22] and through clumpy satellite contributions [23, 24], in addition to the condensation of hot gas out of the halo and subsequent precipitation [25, 26]. These various flows, each with distinct kinematics, interact to produce evolving dynamics within the CGM.

As a consequence, the CGM is not a perfectly stratified atmosphere in hydrostatic equilibrium, but instead exhibits a highly complex *multi-phase, multi-scale* structure. Both theoretical and observational evidence suggest the existence of dense blobs of $\mathrm{T} \sim 10^4$ K cool gas traversing at velocities inconsistent

[3]This is from the TNG50 simulation, which we will discuss in detail in subsequent parts of this thesis.

[4]$\mathrm{R_{200c}}$, the radius within which the mean density is 200 times the critical density of the Universe, is a commonly used scale for defining the virial radius of the halo.

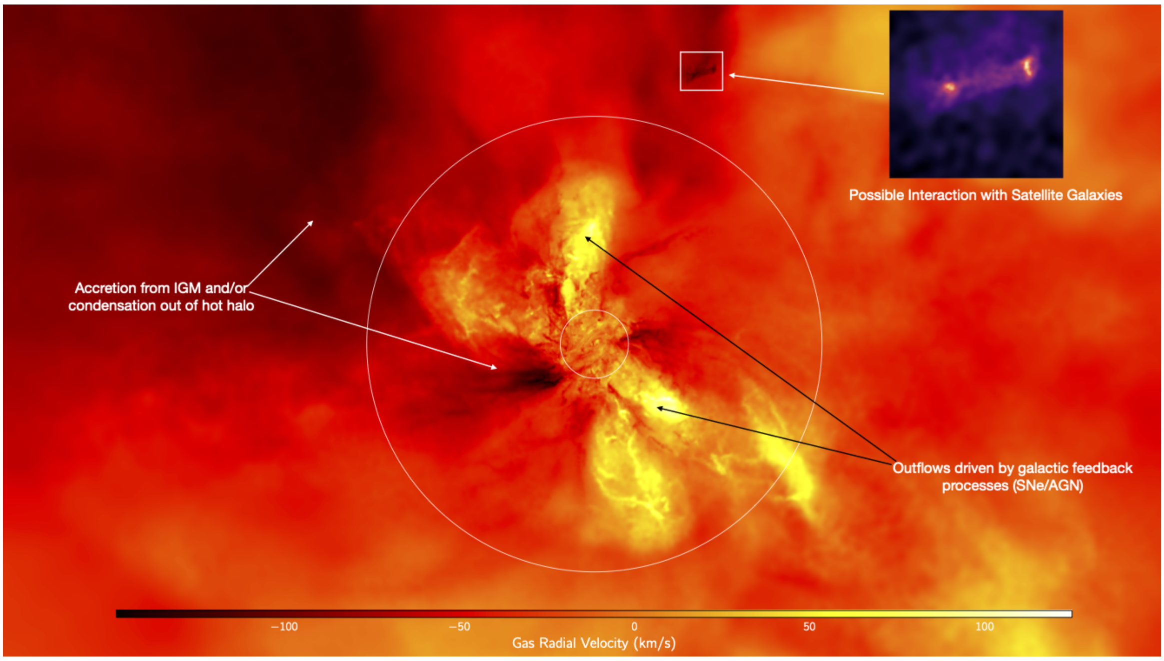

Figure 1.2. A visualization of gas flows around a simulated Milky Way-like galaxy. The circles are drawn at [0.15, 1.00] R_{200c}, demarcating the two boundaries of the CGM. Colors in the background image correspond to the radial velocity of gas, as shown by the colorbar at the bottom. Motion of gas in the CGM is complex, with outflows triggered by galactic feedback coexisting, both spatially and temporally, with large-scale accretion driven by gravity, as well as condensation of gas out of the hot halo.

with galactic (co-)rotation [27, 28], immersed in a volume-filling hot phase close to the virial temperature of the halo [29, 30], possibly with additional super-virial components such as the eROSITA bubbles observed recently on either side of the Milky Way disk [31]. This non-trivial structure of the CGM arises from a combination of *multi-physics* processes, each playing a role in its overall behavior. In the remainder of this section, we will explore the key underlying physics of the CGM, focusing on the interactions between thermal, dynamical, and feedback mechanisms that influence its evolution.

1.2.1 Radiative Cooling

A key physical process behind the multi-phase structure of the CGM is radiative cooling, which predominantly proceeds via a combination of the following two-body interactions [32]:

1. Collisional ionization, where a collision between a free electron and an

atom ionizes a bound electron, thereby reducing the kinetic energy of the free electron.

2. Collisional excitation, where a collision between a free electron and an atom excites a bound electron, and the subsequent decay radiates away a photon.

3. Recombination, where a free electron recombines with an ion, radiating energy as it decays to a lower energy state. A special additional case, for ions that contain at least one bound electron, is dielectronic recombination, where the interaction between the free electron and the ion initially excites a bound electron, and two photons are released as the two electrons successively decay.

4. Bremsstrahlung, where a free electron is accelerated by a nearby ion, subsequently releasing energy as accelerated charged particles radiate.

Note that while molecular cooling and fine-structure transitions could provide additional components, CGM gas densities and temperatures are typically not in the regime for them to be relevant, and hence they are not discussed herein. Similarly, the discussion of three-body cooling processes, which are very unlikely at these densities, is omitted.

For the simple case of an optically thin medium wherein all photons produced via the above processes escape, and assuming the absence of any external background radiation, a balance between the creation and destruction of various ionization species is set by recombination and collisional ionization. In most scenarios relevant to the CGM, the timescale of these processes is small compared to hydrodynamical timescales, and thus this yields a stable equilibrium for the system, commonly referred to as collisional ionization equilibrium (CIE).

Under CIE, the relative abundances of the various ions depend only on temperature. For a gas of primordial composition, cooling rates are dominated at relatively low temperatures by the collisional line excitations of H ($T \sim 10^{4.2}$ K) and He^+ ($T \sim 10^{5.0}$ K), and by bremsstrahlung at $T \gg 10^{6.0}$ K [33]. The cooling rate drops rapidly to zero at $T \lesssim 10^{4.0}$ K as collisions are no longer sufficiently energetic to ionize and/or excite atoms.

The presence of metals adds new line cooling channels, primarily between $\sim 10^{4.0}$ K and $\sim 10^{7.0}$ K. The increase in the cooling rate, and its temperature dependence, depends strongly on the total gas-phase metallicity (Z), as well as the relative abundance of the various metal species. For instance, assuming solar abundance, the total cooling rate at $T \sim 10^{6.0}$ K may increase by as large a factor as ~ 100 for $Z \sim Z_\odot$, but only by ~ 10 for $Z \sim Z_\odot/10$ [34]. While metals also provide possible cooling lines at $T \lesssim 10^{4.0}$ K, the associated rates are typically negligible[5], albeit non-zero [35].

1.2.2 Photo-Ionization

While we assumed the absence of external radiation sources in our discussion above (§ 1.2.1), this condition is rarely satisfied in realistic scenarios. From local radiation fields near quasars and young stellar populations [36, 37], to the metagalactic UV background (UVB) at $z \lesssim 6 - 20$ [38, 39], the universe is believed to be permeated by ionizing radiation.

In addition to directly heating the gas, photo-ionization may also affect the cooling rates discussed above. This occurs because a photo-ionized parcel of gas is over-ionized for its temperature as compared to the CIE case, which reduces the fraction of ions that would otherwise be collisionally excited, ultimately lowering the associated cooling rates [34].

Unlike the case of CIE where the relative abundance of ions is purely a function of temperature, the density of gas plays an important role once photo-ionization is considered. At high densities, where the rate of recombination overpowers that of photo-ionization, cooling is well described by and is similar to CIE. In contrast, at low densities, where photo-ionization dominates and ionizes a significant portion of the gas, cooling rates drop considerably. Note however that the cooling of $T \gg 10^{6.0}$ K gas is only mildly affected, as bremsstrahlung remains the dominant mechanism in this regime.

1.2.3 Galactic Feedback Processes – I: Supernovae

As mentioned earlier, gas in the CGM is greatly influenced by activity occurring within the galaxy. The first of many processes we will discuss is that of supernovae (SNe), which either proceed via the violent explosion of massive stars as they undergo core collapse at the end of their lifetimes, or through runaway nuclear fusion of white dwarfs.

In the former case, the collapse begins when the rate of nuclear fusion

[5] Note again that we ignore fine-structure transitions in this discussion, which may otherwise contribute significantly.

in the stellar core is no longer able to sustain the outward pressure gradient needed to balance the inward pull of gravity. What follows is a rapid infall of matter toward the center of the star, leading to an increase in density and temperature, and ultimately to the formation of an inner core of neutrons, sometimes referred to as a *proto-neutron star*. The still-infalling outer layers of the progenitor star *bounce* off this inner core, producing an outward-propagating shock wave[6], expelling the surrounding gas at rapid velocities [40, 41].

The collapse of white dwarfs is instead triggered by accretion from a companion star or, possibly, by a merger between two white dwarfs, causing the mass of the progenitor to exceed the Chandrasekhar limit[7]. This initiates rapid nuclear fusion[8], leading to the release of $\sim 10^{51}$ erg of energy within a few seconds, thereby ejecting gas in the vicinity at velocities of $\sim O(1000$ - $10000)$ km / s [45, 46]. While the kinetic energy dump in this case is similar to that of core-collapse, the integrated neutrino flux is a factor of ~ 100 - 1000 times larger in the latter, leading to a much larger total energy release [47, 48].

Depending on the initial mass function (IMF) and the assumed SNe progenitor mass, a kinetic energy input of $\sim O(10^{49})$ erg occurs for every unit of solar mass of stars formed [49, 50]. For star-formation rates typical of low-mass, $10^9 \lesssim \mathrm{M}_\star / \mathrm{M}_\odot \lesssim 10^{11}$, galaxies ($\sim 0.1$ - $10\,\mathrm{M}_\odot$ / yr; [51, 52]), this provides a steady input of energy that not only plays a key role in keeping the star formation rate in check, but also in distributing significant amounts of mass, metals, and energy into the CGM [53, 54]. Note that the circulation of matter through the CGM typically extends to larger distances, when normalized by the halo virial radius, for lower-mass galaxies, as they have shallower gravitational potential wells compared to their more massive counterparts [55].

[6]Most theories suggest that this initial shock is likely stalled as it propagates outwards, and re-invigoration by an alternate process is required for its revival. Outflowing neutrinos are the widely postulated candidate, although this remains a topic of ongoing research.

[7]This is the maximum possible mass beyond which electron degeneracy can no longer counterbalance gravity, and is thus the upper limit for a white dwarf. For a non-rotating object, it is $\sim 1.44\,\mathrm{M}_\odot$ [42, 43].

[8]It is, however, believed by some authors that, as a result of increased core density and temperature, as well as perturbations therein, fusion may even begin just before the Chandrasekhar limit is attained (within ~ 1 %), and the collapse takes place soon after [44].

1.2.4 Galactic Feedback Processes – II: Supermassive Black Holes

The evolution of galaxies that populate the higher-mass regime of the spectrum, and their surrounding gaseous halos, is relatively less impacted by SNe-driven feedback since (a) they are typically quenched or have very low star-formation rates [56, 57], (b) their stronger gravity limits the propagation of such outflows to large distances [55], and (c) possibly most importantly, they are dominated by feedback driven by another source – supermassive black holes (SMBHs) [58, 59].

While the formation of stellar-mass black holes (BHs) is believed to proceed via the core-collapse of stars that are sufficiently massive, i.e., those that produce a remnant above the Tolman–Oppenheimer–Volkoff limit ($\sim 3\,\mathrm{M}_{\odot}$; [60]) following the initial core collapse, the origin of SMBHs remains widely debated. The simplest formation pathway involves them growing through a series of BH-BH mergers, starting from a population of stellar-mass BH *seeds*, in addition to accreting material[9] between merger events [63].

Although elegant, recent observations, including those with JWST, suggest the presence of $\gtrsim 10^7\,\mathrm{M}_{\odot}$ black holes at $z \gtrsim 6$ [64, 65], posing a problem, as this may not provide sufficient time for SMBHs to form via the channel discussed above. Alternate pathways include starting from seeds produced either by the collapse of Pop-III stars ($\mathrm{M}_{\mathrm{seed}} \sim 10^2$ - $10^3\,\mathrm{M}_{\odot}$; [66, 67]) or through the collapse of massive primordial gas clouds ($\mathrm{M}_{\mathrm{seed}} \sim 10^5\,\mathrm{M}_{\odot}$; [68, 69]), or from primordial black holes produced soon after the Big Bang [70]. Reality likely involves a combination of all of the above, but irrespective of their origin, most, if not all, luminous galaxies are believed to host an SMBH at their centers [71, 72].

Similar to many other scenarios in astrophysics, SMBHs are typically surrounded by an accretion disk, primarily composed of gas and plasma. At any given radius, a net transfer of angular momentum takes place in the outward direction as a result of viscous stresses[10], leading to an infall of gas and an increase in temperature as matter descends deeper into the potential well. If radiative cooling is efficient and the dominant cooling mechanism, the re-

[9] The accretion is typically assumed to proceed at the Bondi-Hoyle-Lyttleton rate [61, 62], limited by the Eddington rate – the point at which the inward force of gravity balances the outward force from radiation pressure exerted by the infalling matter.

[10] In addition to viscosity produced by the rubbing of particles off each other, turbulent eddies [73] and/or magnetic fields, through the magnetorotational instability [74], are believed to play a key role.

sulting disk is geometrically thin and optically thick, and the spectrum is well described by a sum of black bodies [75, 76]. If disks are radiatively inefficient and are instead cooled primarily by advection, depending on the accretion rates onto the SMBH, they fall under the categories of ADAFs (Advection Dominated Accretion Flows; [77]), Polish doughnuts [78], or slim disks [79]. All three categories show deviations from purely thermal spectra, although slim disks are rather close to that of a black body. Note that the spin of the SMBH may additionally have an impact on the radiative efficiency of the disk, and thus possibly also on setting the optical/geometrical thickness [76, 80].

In addition to the above *thermal*-mode feedback, which not only heats but also ionizes the surrounding gas, accretion towards the SMBH may lead to the formation of powerful (bi-polar) jets. Although this remains an ongoing area of research, the widely preferred description to date is that these are produced through the Blandford–Znajek process[11] [83], wherein poloidal magnetic fields produced by the hot plasma in the inner regions of the disk are twisted as the SMBH spins, ultimately accelerating gas to large velocities along the axis of rotation at the expense of the SMBH's rotational energy.

This mechanical *kinetic*-mode feedback, sometimes referred to as *radio*-mode feedback as such jets are prominently observed in radio wavelengths extending to large distances from the central galaxy, is increasingly prominent and dominates over thermal radiation for massive SMBHs, as they typically have higher accretion rates and spins, leading to more energetic jets [84]. Given the tight correlation between SMBH and galaxy masses [72, 85], relatively larger galaxies are generally dominated by the kinetic-mode feedback to a greater extent [86, 87]. It is believed that these powerful injections of kinetic energy play a role in the quenching of massive central galaxies (M_{200c}[12] $\gtrsim 10^{12}\,M_\odot$; [88, 89]), and in their transformation from (star-forming) spirals to (quenched) ellipticals [90, 91].

As a final note in this sub-section, and as already hinted at above, we mention that the two modes of feedback are often theoretically associated with two distinct types of flows onto SMBHs: the thermal-mode with comparatively higher accretion rates in a standard disk scenario [75, 73], and the kinetic-mode with lower rates through a relatively quasi-spherical hot flow [92, 93], a distinction that is often used in cosmological simulations to (sub-grid) model these processes [89, 94].

[11] Closely related are the Penrose [81] and Blandford-Payne processes [82].
[12] M_{200c}, the virial mass, is the total mass of the halo within R_{200c}.

1.2.5 Magnetic Fields

The universe is believed to be permeated by magnetic fields on all scales, from black hole accretion disks and stars to galaxies and the intergalactic medium (IGM). Although still debated, magnetic fields at galactic, halo, and IGM scales likely originate from weak primordial remnants of the Inflation phase [95, 96], and/or from weak seed fields generated by the Biermann battery [97] or the Weibel instability [98] during the Epoch of Reionization ($z \sim 5-30$). These are thereafter understood to be amplified through the process of structure formation, small-scale dynamos, differential rotation within disk galaxies, and feedback processes [99], resulting in $O(1-10\,\mu\text{G})$ fields by $z \sim 0$ at the centers of galaxies [100], and weaker by a factor of $\sim 10-100$ in their surrounding halos [101].

In the CGM, magnetic fields are thought to play a crucial role in the evolution and survival of small-scale cold gas clouds. They may provide additional pressure support [102, 103], suppress fluid instabilities via magnetic tension [104, 105], affect their shape [106] and kinematics [107, 108], or enhance the drag force on these clouds when surrounded by a draped layer, thus reducing the time required for them to co-move with their surroundings [109]. On larger scales, these fields may impact the properties of halo gas [101, 110], outflows [111], and the hot intracluster medium [112], ultimately influencing the evolution of the galaxy embedded within.

1.2.6 Cosmic Rays

Another non-thermal physical component of significant interest for galaxy evolution is cosmic rays (CRs), a collection of charged particles traveling at nearly the speed of light and primarily composed of protons and alpha particles. Interestingly, the elemental abundance of CR metals is similar to that of the solar system, with the exceptions of lithium, beryllium, and boron, as a fraction of these elements are produced via the spallation of carbon nuclei [113].

CRs are produced by a variety of processes, including supernovae, black hole jets, and cosmic shocks [114]. They span a very large range of energies, from $\sim 10^6$ eV to $\sim 10^{18}$ eV, with their energy density peaking at $\sim 10^9$ eV. Overall, the CR energy reservoir in galaxies is non-negligible and in rough equipartition with magnetic, thermal, and turbulent energies [115, 116].

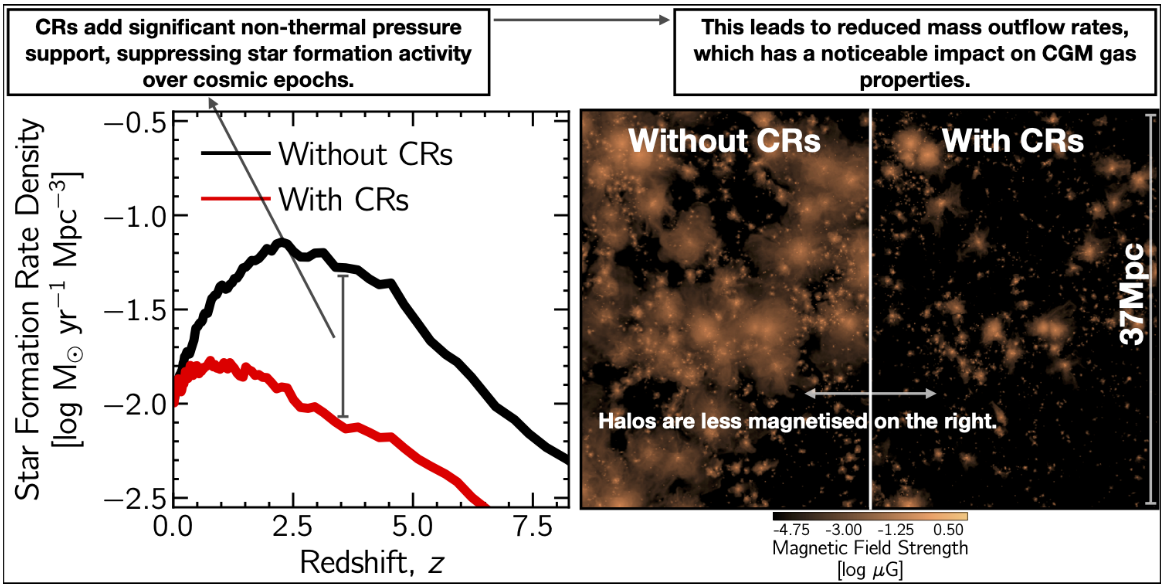

Figure 1.3. A couple of illustrative examples of how CRs may impact the evolution of galaxies and their CGMs. On the left, we contrast the evolution of the cosmic star formation rate density in a case with no CRs (black) versus one with CRs (red), finding that star formation activity is suppressed when this non-thermal component is included. The images on the right compare the large-scale distribution of magnetic fields between the two cases, with halos being less magnetized when CRs are included.

Depending on their momenta, CRs can influence gas in various ways. Cosmic ray particles in the $\lesssim 100\,\mathrm{MeV}\,c^{-1}$ range of the spectrum can ionize their surrounding gas [117, 118], while CRs with $\sim$ GeV c^{-1} can drive powerful galactic scale outflows [119, 120] through gradients in their pressure structure [121, 122]. Much like magnetic fields, their added non-thermal pressure support may also provide additional stability to small-scale cold gas structures [123, 124].

More energetic CRs ($\gtrsim$ TeV c^{-1}) are not dynamically relevant, as they are sub-dominant in terms of their contribution to the total energy budget. However, pion production via hadronic interactions leads to a subsequent decay into γ-rays, providing an important and observable non-thermal radiative signature [125]. While these high energy CRs are less impacted by streaming losses [126], their lower energy counterparts excite Alfvén waves[13] through the CR streaming instability when they propagate faster than the Alfvén velocity [127], effectively heating up gas as these waves damp on short time

[13] A class of plasma waves wherein ions oscillate in response to magnetic tension providing the restoring force.

scales [128].

CRs may thus play an important role in the evolution of galaxies and their gaseous halos. They may have an impact on the accretion of gas onto galaxies by altering the flow structure of CGM gas, thereby modifying structural properties such as galactic disk sizes [129], and also time-integrated quantities like the star formation rates within the ISM [130]. Outflows driven by CRs can be smoother, cooler and denser than those by supernovae [131], which may in turn affect the escape of Lyα photons from galaxies [132]. Their added non-thermal pressure may also result in halos that are significantly cooler [133].

We conclude this sub-section with Figure 1.3, which provides a couple of illustrative examples of how CRs may impact the evolution of galaxies and their CGMs. These are from a suite of simulations[14] run with the same base galaxy physics model as Figure 1.2, but also with an additional module for CR physics in one of the two cases. On the left half of Figure 1.3, we contrast the evolution of the cosmic star formation rate density in a case with no CRs (black) versus one with CRs (red), finding that star formation activity is suppressed when this non-thermal component is included. The images on the right compare the large-scale distribution of magnetic fields between the two cases, showing how halos are less magnetized once CRs are included.

1.2.7 Smooth IGM Accretion

In the Λ-CDM paradigm, the gravitational collapse of gas initially closely follows that of dark matter. As structures begin to form, the situation becomes more complex: according to the first theories developed in the late 1970s, gas accreting from the IGM into massive galaxies is shock-heated to the virial temperature at roughly the virial radius by the virial shock, settles into a pressure-supported layer, and descends further into the halo only after it cools [135, 136]. This *hot-mode* accretion, however, is absent in halos less massive than a given threshold (e.g. $\mathrm{M_{200c}} \lesssim 10^{11.5}\,\mathrm{M}_\odot$; [137]), since these cannot support a stable hot atmosphere, and gas instead accretes through a *cold-mode* channel without being shock-heated [138].

While the prediction of these early analytic works was largely accurate in that a transition between two modes takes place, modern cosmological simulations suggest that even halos above the mass threshold, i.e., those that sup-

[14] These results are from [134], a paper written towards the end of my PhD, but not included in the main body of this cumulative thesis.

port a stable virial shock, may accrete cold gas without being shock-heated, particularly at $z \gtrsim 1.5$. This proceeds through accretion via the filamentary structure of the cosmic web, which is able to penetrate deep into the halo without being heated, as a result of relatively higher densities [139, 140]. Note that such *cold streams* cease at lower z, or once halos cross a second mass threshold at higher z, since the survivability of these filamentary flows scales inversely with the overdensity of the stream to the halo [137].

Over time, accretion from the IGM, both via hot and cold modes, plays an indispensable role in assembling the baryonic structure of both galaxies and their surrounding CGMs, with certain simulations suggesting that ~ 80 % of the gaseous supply of $z = 0$ MW-like CGMs may have been accreted from the IGM over the last ~ 10.5 Gyr [141].

1.2.8 Clumpy Satellite Accretion

The other mode of cosmological accretion through which (central) galaxies acquire mass is via satellites. These are objects that collapsed into their dark matter halos in the past, after which they drifted along cosmic filaments into the potential well of a more massive structure.

As satellites pass through their host halos, their orbits gradually decay due to dynamical friction [142], causing them to spiral inward toward the potential minimum and eventually merge with the central galaxy[15]. Among other effects, such mergers may contribute a substantial amount of cold gas to the central [145], trigger bursty periods of star formation [146], perturb the disks of spiral galaxies [147], and, in some cases, play a role in transforming spirals into ellipticals [148].

While these effects may have an indirect impact on the gaseous reservoir of the CGM, the phase during which satellites spiral through the halo is of direct importance. In addition to tidal stripping due to differential gravity [149], their relative motion with respect to their ambient background drives an effective force in the opposite direction due to ram pressure [150], causing them to deposit gas into the surrounding medium. While these contributions may directly add to the cold mass budget of the host CGM, particularly when the ISM of the satellite is stripped [151], they may also trigger the condensation of hot halo gas via density perturbations [152]. In certain contrasting

[15]Note although that a fraction of galaxies may be completely disrupted by gravitational and hydrodynamical forces before they can merge [143, 144].

cases, the motion of satellites may instead induce turbulence and heat the surrounding gas [153]. Similar to IGM accretion, the acquisition of gas via satellites is believed to play an important role, with certain simulations suggesting that $\sim 20\,\%$ of the gaseous supply of $z=0$ MW-like CGMs may have been contributed by satellites over the last ~ 10.5 Gyr [141].

1.2.9 Fluid Instabilities

Hydrodynamic fluid instabilities play a key role in the origin, evolution, and survival of cold CGM gas [154]. The most important of these is unarguably the Kelvin–Helmholtz instability (KHI; [155, 156]), which occurs due to a velocity shear between two fluids[16]. In the limit of weak surface tension, small perturbations along the interface are rapidly amplified, driving the mixing of the two fluids. Also relevant in certain contexts is the Rayleigh–Taylor instability (RTI; [157]), which instead takes place in the presence of a density contrast between two fluids, again leading to their mixing. The exact impact that these instabilities have on cloud growth depends on multiple factors, particularly the efficiency of radiative cooling [158, 159], and also the orientation and strength of magnetic fields to a smaller extent [109, 105], which we discuss in more detail in subsequent parts of this thesis[17].

1.3 Numerical Techniques

As motivated above, modeling and simulating a realistic CGM is a challenging endeavor due to the plethora of physical processes and the magnitude of scales involved. Over the years, various types of simulation setups have been experimented with to explore puzzling and unanswered questions regarding the CGM. Of notable mention are turbulent-box [159], wind-tunnel [158], tall-box [160], and isolated galaxy [126] setups. The first two aim to simulate $\sim O(1\,\mathrm{kpc})$ scale regions – i.e. patches of the CGM – while the latter two model slightly larger volumes, typically with a galaxy at the center, surrounded by the disk-halo interface and, in some cases, the inner CGM even farther away.

While these numerical experiments have contributed immensely to advancing the field, particularly because they are relatively computationally

[16]Note that a velocity shear within a single continuous fluid also gives rise to the same instability, although this is less important for gas in the CGM.

[17]A detailed analysis was also presented in [141], a paper written toward the end of my PhD, but not included in the main body of this cumulative thesis.

inexpensive to run and thus make it possible to explore a large range of parameter space, they have a key limitation: they are all *idealized*, i.e., they typically make a number of simplifying assumptions. Relatedly, and perhaps most importantly, they lack a full Λ-CDM treatment. *Cosmological magnetohydrodynamical* simulations are a natural pathway forward in this regard and offer an avenue to self-consistently reproduce the complexity and diversity of the CGM. Throughout the rest of this introduction and thesis, we will almost exclusively focus on such setups, although note that a number of numerical techniques discussed herein – particularly that of hydrodynamics – are commonly applied in simpler idealized simulations as well.

1.3.1 Gravity and Collisionless Components

Gravity, being the primary driver of cosmological structure formation in the universe, is one of the most important physical process to accurately account for. Collisionless components – i.e. dark matter, stars, and black holes – are modeled under the continuum limit of the collisionless Boltzmann equation

$$\frac{\mathrm{d}f}{\mathrm{d}t} \equiv \frac{\partial f}{\partial t} + \vec{v}\frac{\partial f}{\partial \vec{x}} - \frac{\partial \Phi}{\partial \vec{x}}\frac{\partial f}{\partial \vec{v}} = 0, \tag{1.1}$$

where $\vec{x}$ and $\vec{v}$ correspond to position and velocity, $f(\vec{x}, \vec{v}, t)$ is the mass density in six-dimensional phase-space, and the gravitational potential Φ satisfies the Poisson equation

$$\nabla^2\Phi(\vec{x}, t) = 4\pi G \int f(\vec{x}, \vec{v}, t)\mathrm{d}\vec{v}, \tag{1.2}$$

with G being the universal gravitational constant. While one could in principle directly solve these equations using finite volume methods (see below), the standard approach involves discretizing the underlying mass distribution into a finite number ($= N$) of particles.

In this *N-body* approach, the net $\Phi(\vec{x})$ at any given position $\vec{x}$ is simply the sum of Newtonian potentials contributed by each particle

$$\Phi(\vec{x}) = -G\sum_{i=1}^{N} \frac{m_i}{(|\vec{x} - \vec{x}_i|^2 + \epsilon^2)}, \tag{1.3}$$

where m_i and x_i are the mass[18] and position of the i^{th} particle, and ϵ, the

[18] In practice, when the distance to a particle is smaller than a length scale that is a proxy for its radius, only a fraction of the mass is used in the calculation. This requires the assumption of a mass distribution within each particle, i.e. a kernel function, which is similar to the SPH

gravitational softening length, is included to suppress spurious scattering events at small separations. Note that this naturally sets a resolution cut-off below which gravitational motions are no longer captured. The acceleration of the $i^{\rm th}$ particle is thereafter set by $\ddot{\vec{x}}_i = -\vec{\nabla}\,\Phi(\vec{x}_i)$.

Of the various techniques to numerically compute $\Phi(\vec{x})$, the simplest involves brute-forcing the calculation as a pairwise summation for each particle with every other particle in the domain. While this yields the highest accuracy relative to other methods we will discuss, i.e. down to machine precision, it is an $O(N^2)$ operation, and understandably the slowest. Broadly speaking, all other methods fall into two categories, or a combination thereof.

The first class consists of so called *tree*-based algorithms [161, 162], wherein particles are hierarchically partitioned into smaller groups, i.e. nodes, using a tree structure. The computation of the potential at any given position thereafter proceeds by *walking* the tree from top to bottom: if the angular size of any given node is smaller than a critical limit θ_c, typically referred to as the tree opening criterion, an effective contribution is added by assuming that all the mass of the node is concentrated at its center of mass. If not, the node is opened, and the tree walk continues to the next set of daughter nodes. For a given θ_c, the expected number of nodes to be opened scales as $\log(N)$ [163], and thus the overall force computation through such a method is an $O(N\,\log N)$ operation.

The alternate approach, referred to as particle-mesh (PM), is to construct a mesh of grid cells onto which the mass of particles is deposited. The force calculation can then be performed either in Fourier space [164], or in real space using approximate iterative methods [165]. These methods are also $O(M\,\log M)$ operations, with M being the total number of grid points.

In practice, combinations of both methods are typically used, since PM calculations are relatively less accurate on small scales when particles are close together, and tree methods become increasingly more expensive as the number of particles grows larger. One of the most common approaches, which is also used in all simulations presented in this thesis, is the TreePM algorithm [166]. This employs a tree-based calculation for short-range forces and transitions to a Fourier PM method for long-range interactions, thus providing the perfect trade-off between speed and accuracy.

treatment for hydrodynamics that we will shortly discuss.

1.3.2 Hydrodynamics

While gas interacts gravitationally in the same way as the components discussed above, it is not collisionless, and thus requires further treatment to account for its fluid interactions. In what follows, we assume gas to be ideal and inviscid, in which case its dynamics is set by the following system of hyperbolic partial differential equations, commonly called the Euler equations, which signify the conservation of mass, momentum, and energy, respectively

$$\frac{\partial \rho}{\partial t} + \vec{\nabla}.(\rho\vec{v}) = 0, \tag{1.4}$$
$$\frac{\partial (\rho\vec{v})}{\partial t} + \vec{\nabla}.(\rho\vec{v}\vec{v}^T + P) = 0, \tag{1.5}$$
$$\frac{\partial (\rho e)}{\partial t} + \vec{\nabla}.((\rho e + P)\vec{v}) = 0, \tag{1.6}$$

where ρ is the gas density, $\vec{v}$ the velocity, $e = u + \frac{|\vec{v}|^2}{2}$ is the total energy per unit mass with u being the thermal energy per unit mass, and the thermal pressure P is set by the equation of state

$$P = (\gamma - 1)\rho u, \tag{1.7}$$

where γ, the adiabatic index, is commonly set to $\frac{5}{3}$ under the assumption of a monoatomic gas. In addition to the above, including more physics would naturally include new terms, and possibly also new conservation equations. Of particular relevance to this thesis is the inclusion of magnetic fields. Under the assumption of ideal magnetohydrodynamics (MHD), i.e. in the limit of zero resistivity and no magnetic diffusion, the above equations are modified to

$$\frac{\partial \rho}{\partial t} + \vec{\nabla}.(\rho\vec{v}) = 0, \tag{1.8}$$
$$\frac{\partial (\rho\vec{v})}{\partial t} + \vec{\nabla}.(\rho\vec{v}\vec{v}^T + P_{\text{tot}} - \vec{B}\vec{B}^T) = 0, \tag{1.9}$$
$$\frac{\partial (\rho e_{\text{tot}})}{\partial t} + \vec{\nabla}.((\rho e_{\text{tot}} + P_{\text{tot}})\vec{v} - \vec{B}(\vec{v}.\vec{B})) = 0, \tag{1.10}$$

where $\vec{B}$ is the magnetic field vector, and $P_{\text{tot}} = P + \frac{|\vec{B}|^2}{2}$ and $e_{\text{tot}} = e + \frac{|\vec{B}|^2}{2\rho}$ include respective contributions of pressure and energy from the magnetic field. A final conservation equation

$$\frac{\partial \vec{B}}{\partial t} + \vec{\nabla}.(\vec{B}\vec{v}^T - \vec{v}\vec{B}^T) = 0 \tag{1.11}$$

sets the evolution of the magnetic field with time. In practice, note that various source terms – i.e. quantities on the right-hand side of the equations – may additionally be included. For instance, if divergence cleaning methods are used to enforce $\vec{\nabla}\,.\,\vec{B}=0$ [167], if a gravitational field is present, if radiative cooling is included, or if the equations are extended to comoving coordinates, as is common in cosmological simulations.

As with collisionless particles, numerically solving these equations requires the discretization of the underlying mass distribution into a number of finite elements. The first commonly used approach in astrophysics, *smoothed particle hydrodynamics* (SPH; [168]), is similar to the N-body technique in that the gas density field is partitioned into particles. A given fluid property $p(\vec{x})$ at a sample point $\vec{x}$ is then approximated as a weighted sum over n neighboring gas particles

$$p(\vec{x}) = \sum_{i=1}^{n} V_i p_i W(|\vec{x} - \vec{x_i}|, h), \tag{1.12}$$

where V_i is the volume of the i^{th} particle, and W is a suitably chosen kernel function that splats the relevant quantity over a length scale h. These interpolated values are then used to evolve particles forward in time, solving a modified version of the above equations [169].

An alternate[19] approach involves discretizing in space rather than mass. In these *finite volume methods* (FVM; [170]), the simulation domain is divided into a mesh of grid points. This facilitates solving a conservation equation of the form

$$\frac{\partial \vec{u}}{\partial t} + \vec{\nabla}.\vec{f}(\vec{u}) = 0 \tag{1.13}$$

by evaluating fluxes at the surface of each mesh cell: taking a volume integral over a cell i, Equation 1.13 is now

$$\int_{v_i} \frac{\partial \vec{u}}{\partial t} \partial v + \int_{v_i} \vec{\nabla}.\vec{f}(\vec{u}) \partial v = 0, \tag{1.14}$$

where v_i is the volume of the cell. Using the divergence theorem, the integral on the right can be converted to a surface intergral:

$$v_i \frac{\partial \bar{\vec{u}}_i}{\partial t} + \int_{S_i} (\vec{f}(\vec{u}).\vec{n}) dS = 0, \tag{1.15}$$

[19]Note that while other methods such as that of finite element and discontinuous Galerkin also exist, we do not discuss them here.

where $\vec{n}$ is the normal vector to the surface, and the first integral is now replaced with the volume average of $\vec{u}$ over the cell. In effect, a partial differential equation is solved instead in an algebraic form, and the evolution of the cell over a timestep Δt is computed as

$$\vec{u_i}(n+1) = \vec{u_i}(n) - \frac{\Delta t}{v_i} \sum_{\text{faces}} f_{\text{face}} \Delta S_{\text{face}}, \tag{1.16}$$

where $\vec{u_i}(n)$ is the value in the n^{th} timestep, and f_{face} is the flux through a face of area ΔS_{face}. A number of methods exist to compute fluxes, including upwind, central difference, Godunov, and MUSCL schemes, among others, each of which are accurate to varying degrees and find application in different problems.

While both SPH and FVM have been extensively used for running various simulations in the past, each has its own set of limitations. For instance, FVM methods are not Galilean-invariant, which poses a problem in the case of highly supersonic flows, a common feature in astrophysics. SPH, on the other hand, has been shown to spuriously suppress the onset and/or growth of fluid instabilities, leading to inefficient mixing at the interface between two fluids. To overcome these limitations, various new iterations have been proposed. New classes of methods that combine ideas from both approaches have also emerged, such as the meshless finite volume (MFV) and meshless finite mass (MFM) schemes [171].

Throughout the rest of this thesis, we will use the AREPO code [172], which employs yet another hybrid of the SPH and FVM techniques: it uses a moving, unstructured mesh defined by the Voronoi tessellation of a set of points. The computational domain is divided into a number of cells, and the conservation equations are solved using a FVM approach, with fluxes computed by solving the Riemann problem across the faces of neighboring cells, using the second-order MUSCL-Hancock scheme, which utilizes not only cell averages but also linear gradients, and the approximate Riemann HLLD solver in the case of MHD. Analogous to SPH, the cells move along with the local flow. Moreover, the cells are maintained close to a *target mass resolution*, thereby naturally providing improved spatial resolution in regions where the gas density is higher. This is achieved through a series of refinement and derefinement operations, where a cell that is too massive for its target mass is split into two smaller ones, while a cell that is under-massive for its target mass is dissolved into the background mesh, as visualised in Figure 1.4. This

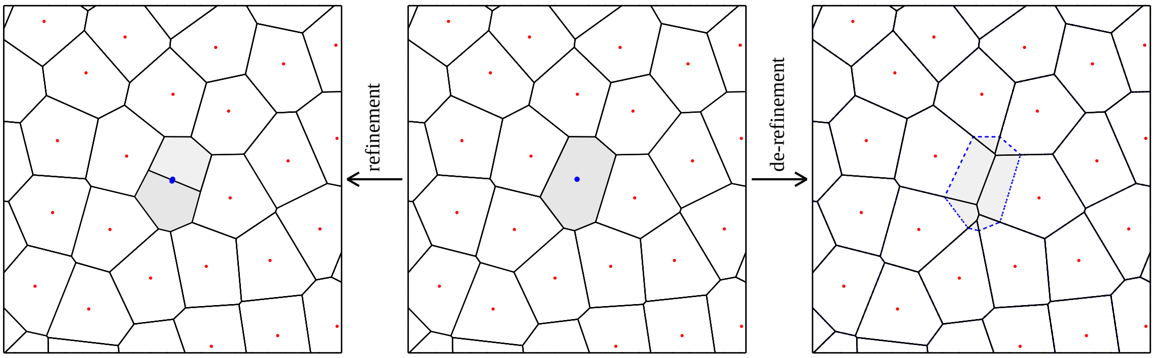

Figure 1.4. A visualisation of refinement and de-refinement operations in AREPO. In the central panel, an exemplary Voronoi cell is highlighted in grey. If this cell is too massive for its target mass, it is split into two smaller ones, as shown on the left. If instead it is under-massive relative to its target mass, it is dissolved into the background mesh, as shown on the right (image taken from [172]).

scheme also makes it possible to effortlessly adjust the resolution in a particular region by predefining a target mass criterion, which forms the basis for a new set of simulations described in Chapters 5 and 6.

1.3.3 The IllustrisTNG Galaxy Formation Model

With gravity and hydrodynamics set up, the next crucial step is to define a suitable model for the physics of galaxy formation and evolution. In this thesis, we adopt the IllustrisTNG model, which is described in detail in [94, 173]. We provide a summary of the key points here:

1. Star formation: Gas clouds that are sufficiently massive, i.e. those that exceed the Jeans mass, may collapse into stars. The corresponding Jeans length in typical galactic star-forming regions is ~ 0.1 pc, much smaller than the resolution limit of present-day cosmological simulations. The physics of star-formation is thus captured using the *sub-grid* model presented in [49], wherein star particles form stochastically from gas cells that exceed a density threshold of $n_{\mathrm{H}} \sim 0.1\,\mathrm{cm}^{-3}$ assuming a Chabrier initial mass function [174] and the Kennicutt-Schmidt relation [175].

2. Stellar evolution: As noted above, stars in the simulation are tracked in the form of a population grouped into a star particle. These populations are evolved to return mass and metals to their surrounding gas through

Type-Ia and Type-II SNe and stellar winds driven by asymptotic giant branch (AGB) stars according to pre-computed tables.

3. Radiative cooling and photo-ionization: Primordial cooling rates are computed by iteratively solving the equations described in [33], while the contribution of metals is added using pre-tabulated values interpolated as a function of metallicity, temperature, density and redshift [34]. Cooling is modulated by the meta-galactic UV-background prescription from [38], as well as contributions from AGN radiation fields [176], with corrections for self-sheilding in case of high gas densities [177].

4. Black hole formation and growth: SMBHs are seeded with a mass of $\sim 10^6\,\mathrm{M}_\odot$ at the potential minima of halos that cross a virial mass threshold of $\mathrm{M}_{200c} \sim 10^{10.8}\,\mathrm{M}_\odot$. Over time, they grow by accreting mass at the Eddington-limited Bondi-Hoyle-Lyttleton rate, or by merging with other SMBHs. The motion of SMBHs is set by re-positioning them to the local potential minima, and not captured gravitationally since dynamical friction at such scales is unresolved in these simulations.

5. Stellar Feedback: Feedback driven by SNe proceeds through a decoupled wind scheme [49], wherein star-forming gas cells stochastically launch so-called wind-particles with a given velocity, mass loading factor, and metal content. These recouple with the mesh after traveling a prescribed distance, thereby depositing mass, momentum, metal, and energy into the surrounding gas, effectively driving galactic scale outflows.

6. Black hole feedback: SMBH feedback proceeds in two modes: at high accretion rates relative to the Eddington limit, thermal energy is dumped into the surrounding gas. At low accretion rates, feedback instead proceeds via a kinetic kick injection. The transition between the two channels occurs at an accretion rate threshold that depends on SMBH mass, such that SMBHs with $\mathrm{M}_{\mathrm{SMBH}} \gtrsim 10^8\,\mathrm{M}_\odot$ are preferentially active in the kinetic state.

7. Magnetic fields: A uniform primordial field of 10^{-14} comoving Gauss is seeded at the start of the simulation, which is subsequently (self consistently) amplified as a combined result of structure formation, small-scale dynamos, and feedback processes [101]. The approximate Riemann HLLD solver is used to evaluate fluxes, and the [167] cleaning scheme is used to maintain $\vec{\nabla}\,.\,\vec{B} = 0$ over time.

The IllustrisTNG model has been calibrated, i.e. fine-tuned, to reproduce a number of statistical properties of galaxies in the known universe, particularly the cosmic star formation rate density, the stellar mass function, and the stellar-to-halo mass, stellar-to-SMBH mass, stellar size-to-mass, and halo gas fraction relations, all at $z = 0$, except the first [173]. A number of other galaxy properties are also in good agreement with observational constraints, such as the bimodal galaxy color distribution [178], quenched fractions of satellite galaxies [179], the galaxy size-mass relation at $z > 0$ [180], among others, highlighting the success and robustness of this galaxy physics model.

1.3.4 Initial Conditions

Last, but certainly not least, is defining the initial conditions (ICs) to start the simulation from. Current understanding suggests that all large-scale structure in the Universe likely arose from initially Gaussian fluctuations in the density field. Observations of the cosmic microwave background (CMB) radiation [181, 1], i.e. the relic radiation from the Big Bang, allow us to constrain not only the various cosmological parameters associated with the ΛCDM model but also the power spectrum $P(k)$ of primordial density fluctuations. This, in turn, enables the reconstruction of the initial density field. Using the first-order Zel'dovich approximation [182], or a more accurate treatment with second order Lagrangian perturbation theory (2LPT; [183]), the initial field can be evolved to the beginning of the nonlinear regime, thereby setting up the initial conditions for a cosmological simulation. Commonly used and publicly available codes for IC generation include N-GENIC [184] and MUSIC [185], which use the Zel'dovich approximation and 2LPT, respectively. For all simulations discussed in this thesis, ICs are generated using the former at $z = 127$.

1.4 Structure of this thesis

This thesis is a compilation of my published work on addressing and exploring key questions regarding the CGM, with a particular focus on Milky Way-like (MW-like) galaxies and, in most cases, the cold phase of gas.

In Chapters 2 - 4, we primarily analyze data from TNG50 [186, 19], the highest resolution box in the IllustrisTNG suite of simulations, which evolves a volume of $\sim (50\,\mathrm{Mpc})^3$ at a baryonic mass resolution of $\sim 10^5\,\mathrm{M}_\odot$. We focus on characterizing the properties of CGM gas of MW-likes galaxies in Chapter 2, the occurrence of cold clouds and their properties in Chapter 3, and the impact of galactic feedback processes on the the angular structure of magnetic fields in Chapter 4.

In Chapters 5 and 6, we introduce and analyze the new GIBLE simulations, a suite of cosmological magnetohydrodynamical zoom-in simulations[20] of eight MW-like galaxies with additional refinement in the CGM, also run with the same IllustrisTNG galaxy physics model. These new simulations make it possible to study clouds in cosmological simulations at an unprecedented level compared to earlier efforts. In particular, in the simulation analyzed in Chapter 6, we observe fascinating magnetic field structures at $\sim$kpc scales, which are simply absent at lower resolutions due to sparser sampling of the underlying gas distribution. We summarize in Chapter 7 and provide an outlook on future work and extensions.

[20]These are setups where only the primary object (in this case, the MW-like halo) is simulated at high resolution, with the rest of the box simulated at progressively coarser levels as one moves farther away.

Chapter 2

The Circumgalactic Medium of Milky Way-like Galaxies in the TNG50 Simulation – I: Halo Gas Properties and the Role of SMBH Feedback

2.1 Statement of contribution

- Scientific Analysis: My contribution was central. In addition, the analysis benefited greatly from feedback provided by collaborators.
- Figures: I independently produced the figures based on the planned scientific analysis. The process also greatly benefited from iterative input by collaborators.
- Writing: I primarily wrote most of the manuscript text. The drafts were further refined based on feedback from collaborators.
- Code and Simulation Development: This paper made use of the publicly available TNG50 simulation. The code to analyze the simulation data and generate plots was written by me, but it also benefited from input provided by colleagues.

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **518,** 5754–5777 (2023)
Advance Access publication 2022 December 2

https://doi.org/10.1093/mnras/stac3524

The circumgalactic medium of Milky Way-like galaxies in the TNG50 simulation – I: halo gas properties and the role of SMBH feedback

Rahul Ramesh [1]★ Dylan Nelson [1] and Annalisa Pillepich [2]

[1] *Universität Heidelberg, Zentrum für Astronomie, Institut für theoretische Astrophysik, Albert-Ueberle-Str 2, D-69120 Heidelberg, Germany*
[2] *Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

Accepted 2022 November 28. Received 2022 November 25; in original form 2022 November 1

ABSTRACT

We analyse the physical properties of gas in the circumgalactic medium (CGM) of 132 Milky Way (MW)-like galaxies at $z = 0$ from the cosmological magneto-hydrodynamical simulation TNG50, part of the IllustrisTNG project. The properties and abundance of CGM gas across the sample are diverse, and the fractional budgets of different phases (cold, warm, and hot), as well as neutral H I mass and metal mass, vary considerably. Over our stellar mass range of $10^{10.5} < M_\star/\mathrm{M}_\odot < 10^{10.9}$, radial profiles of gas physical properties from $0.15 < R/R_{200c} < 1.0$ reveal great CGM structural complexity, with significant variations both at fixed distance around individual galaxies, and across different galaxies. CGM gas is multiphase: the distributions of density, temperature, and entropy are all multimodal, while metallicity and thermal pressure distributions are unimodal; all are broad. We present predictions for magnetic fields in MW-like haloes: a median field strength of $|B| \sim 1\,\mu$G in the inner halo decreases rapidly at larger distance, while magnetic pressure dominates over thermal pressure only within $\sim 0.2 \times R_{200c}$. Virial temperature gas at $\sim 10^6$ K coexists with a subdominant cool, $< 10^5$ K, component in approximate pressure equilibrium. Finally, the physical properties of the CGM are tightly connected to the galactic star formation rate, in turn dependent on feedback from supermassive black holes (SMBHs). In TNG50, we find that energy from SMBH-driven kinetic winds generates high-velocity outflows ($\gtrsim$500–2000 km s^{-1}), heats gas to supervirial temperatures ($>10^{6.5-7}$ K), and regulates the net balance of inflows versus outflows in otherwise quasi-static gaseous haloes.

Key words: galaxies: haloes – galaxies: kinematics and dynamics.

1 INTRODUCTION

The circumgalactic medium (CGM) is believed to play a vital role in the evolution of galaxies. It is the region through which gas accreting from larger scales and from the intergalactic medium (IGM) must flow. It harbours gas previously ejected from the galaxy due to feedback processes, as well as 'recycling' or fountain flows that circulate in the halo. Each of these gaseous flows will have distinct kinematics, producing complex and evolving dynamics in the CGM. In addition, gaseous haloes are thought to have complex spatial and phase structure, being both multiscale and multiphase. In particular, gas co-exists in both cool $\sim 10^4$ K and volume-filling hot components, the latter at or near the halo virial temperature (see Tumlinson, Peeples & Werk 2017 for a recent review, and Fielding et al. 2020 for a recent cross-simulation comparison of CGM properties).

Our own Milky Way (MW) galaxy is embedded in a multiphase and multiscale CGM. Observations of the MW halo have detected cool gas via Hα emission at $T \lesssim 10^{4.5}$ K (Putman et al. 2003), as well as through the H I 21-cm line at $T \lesssim 10^4$ K (Peek et al. 2011). At the same time, observations of X-ray lines such as O VI, O VII and O VIII suggest the presence of a warm-hot phase at $T > 10^5$ K, with *XMM–Newton* (Henley et al. 2010; Miller & Bregman 2015; Miller, Hodges-Kluck & Bregman 2016); *Chandra* (Gupta et al. 2012; Fang, Bullock & Boylan-Kolchin 2013); and *HaloSat* (Kaaret et al. 2020). In addition, absorption features from HST/COS spectroscopy towards nearby stars at varying distances (∼4–15 kpc) reveals multiple cool to warm metal ionization species, including Ca II, Fe II, Si IV, C IV (Werk et al. 2019).

The CGM of the MW exhibits a prominent, large-scale feature. Namely, the Fermi (Su, Slatyer & Finkbeiner 2010) and and eROSITA (Predehl et al. 2020) bubbles: these are two bipolar cocoon-like structures emerging from the Galactic Centre, extending below and above the Galaxy's disc up to 10–14 kpc, and emitting in gamma-ray and X-ray, respectively. Recently, Ashley et al. (2022) observed the existence of multiphase gas clouds embedded within the two Fermi bubbles on either side of the galactic plane, adding support to the picture of a multiphase structure of the MW CGM.

The MW CGM is highly multiscale. It is believed to host small structures of the order of $\lesssim$1–10 kpc in the form of neutral hydrogen rich, high-velocity clouds (Wakker & van Woerden 1997). The resulting kinematics have also been observed to be complex. While a non-negligible amount of cold gas is both inflowing and outflowing (Clark, Bordoloi & Fox 2022), at least a fraction of hot gas is believed to rotate around the centre of the halo, with the amount of angular momentum comparable to that of the stellar disc (Hodges-Kluck, Miller & Bregman 2016). Furthermore, the gas within the

★ E-mail: oorc06@gmail.com

Fermi/eROSITA bubbles has been measured to have radial velocities ranging from 330 km s^{-1} (for H I clouds close to the disc: Di Teodoro et al. 2018) to 900–1300 km s^{-1} for UV clouds at higher altitudes (Fox et al. 2015; Bordoloi et al. 2017; Ashley et al. 2020).

This picture of a multiphase, multiscale CGM has also gained support from cosmological hydrodynamical simulations. For instance, using cosmological simulations, Kereš & Hernquist (2009), Fernández, Joung & Putman (2012), and Joung, Bryan & Putman (2012) showed the existence of cold gas clouds in the CGM of MW-like galaxies, and Sokołowska et al. (2016) pointed out the existence of a hot diffuse halo around similar galaxies. With the HESTIA simulations, Damle et al. (2022) noted the presence of both the cold (H I and Si III) and warm-hot phases (O VI, O VII, and O VIII) in the CGM of haloes resembling the MW-M31 system. With the FOGGIE simulations, a set of simulations that preferentially increase resolution in the CGM, Peeples et al. (2019) and Corlies et al. (2020) noted the existence of a multiphase structure in the CGM of MW like haloes, as did Hani et al. (2019) with the Auriga suite of simulations, a set of high-resolution zoom-in simulations of MW-like galaxies. Hafen et al. (2019) used a sample of FIRE-2 simulations to show that gas in the CGM has multiple origins, leading to a diversity of gas properties. Collectively, these studies reveal how the MW gaseous halo assembles to its redshift zero state.

While cosmological simulations are important tools to study properties of realistic CGM realizations, they are computationally expensive. As an alternative, various semi-analytic models have been developed. For instance, Stern et al. (2016) used measurements of ionic column densities to develop a model describing cool gas in the CGM, deriving a mean cool gas density radial density profile. This model was extended to include steady-state cooling flows as a single-parameter family of solutions (Stern et al. 2019). On the other end of the temperature spectrum, Faerman, Sternberg & McKee (2017) used observations of O VI, O VII, and O VIII lines to build a model for the warm-hot phases of the CGM, thereby suggesting that hot galactic coronae can contain significant amounts of gas, possibly accounting for the previously 'missing baryons' (see also Faerman, Sternberg & McKee 2020). Voit et al. (2017) proposed a global model to describe the condensation and precipitation of gas in the CGM, arguing that precipitation can proceed in two modes: either by gas 'uplift' in galactic outflows for $t_{\rm cool}/t_{\rm ff} \lesssim 10$, or due to the slope of the entropy gradient (see also Sharma et al. 2012). Using observations of Si IV, O VI, and C IV, Qu & Bregman (2019) built two-dimensional models for the distribution of warm gas ($T \sim 10^5$ K), suggesting that the disc components of Si IV and O VI have similar density profiles, while the distribution and kinematics of C IV is similar to Si IV (Qu et al. 2020; Qu, Lindley & Bregman 2022).

While early observations focused on the nearest CGM – the gaseous halo of our own MW – we now have detailed observations of other haloes, out to $z \sim 6$ (Leclercq et al. 2017). Extragalactic halo samples increase the scope of observational probes, studying features and physics not dominant in MW-like systems. Observations of neutral hydrogen (Cai et al. 2017; Prochaska et al. 2017; Chen et al. 2018), Mg II (Zhu et al. 2014), and Ly α (MUSE: Leclercq et al. 2017, 2020; see also Byrohl et al. 2021 for comparisons with the TNG50 simulation) have shown the presence of non-negligible amounts of cold gas in the CGM around galaxies across cosmic epochs. For example, Hennawi et al. (2015) reported that the most massive structures in the distant universe ($z \sim 2$; $M_{\rm halo} \sim 10^{13}\,{\rm M}_\odot$) can have as much as $10^{11}\,{\rm M}_\odot$ of cold gas. Focusing on ionized gas phases, observations of O VI with HST/COS (Stocke et al. 2013; Werk et al. 2016) have detected warm gas, while hot X-ray emitting gas in extended haloes around other galaxies has also been detected (with e.g. Chandra Bogdán et al. 2013b,a; Goulding et al. 2016; *XMM–Newton*: Li et al. 2016; CASBaH: Burchett et al. 2019; and eROSITA via stacking: Comparat et al. 2022; Chadayammuri et al. 2022). The CGM of the Andromeda galaxy has been studied in detail with multiple independent quasar sightlines, due to its proximity (Lehner et al. 2020). Finally, the warm-hot phase of extragalactic gaseous haloes is also accessible through the Sunyaev-Zel'dovich effect (tSZ; de Graaff et al. 2019; Lim et al. 2021).

In addition to the observable tracers of different gas phases mentioned above, non-thermal components including magnetic fields and cosmic rays may play an important role with respect to CGM gas. Using the Auriga suite of simulations, Pakmor et al. (2020) showed that the CGM of MW-like haloes are highly magnetized ($|B| \gtrsim 0.1\,\mu$G) before $z \sim 1$, mainly due to outflows that transport magnetized gas outwards away from the galaxy. van de Voort et al. (2021) showed, with one Auriga halo, that these magnetic fields can alter the physical properties of CGM gas, including inflow velocities, temperatures, and pressures. Magnetic pressure may be important in gaseous haloes, including its ability to stabilize dense, cold clouds surrounded by a hot ambient medium (Nelson et al. 2020). Pressure support due to cosmic rays may also stabilize cold gas clouds (Butsky & Quinn 2018; Ji et al. 2020; Huang, Jiang & Davis 2022), which can change characteristics of cold-phase gas (Butsky et al. 2020), and can make the CGM less thermally supported and thus cooler (Hopkins et al. 2020). The formation, and survival, of such clouds remains, however, an open theoretical question (e.g. McCourt et al. 2015; Gronke & Oh 2020; Li et al. 2020; Dutta, Sharma & Nelson 2022).

Finally, the CGM does not exist in isolation. It is influenced by external perturbations, such as galaxy mergers and gas accretion, as well as by feedback-driven outflows from the central galaxy. Recent studies based on current large-scale cosmological hydrodynamical galaxy simulations – EAGLE and IllustrisTNG – have quantified the importance of supermassive black hole (SMBH) feedback on properties of the CGM (Davies et al. 2020; Oppenheimer et al. 2020; Truong et al. 2020; Zinger et al. 2020; Truong, Pillepich & Werner 2021a). The injection of energy by SMBHs drives gas out of galaxies at high velocities, increases the average entropy of the halo gas, and lengthens typical cooling times, not only preventing future star formation but also directly impacting the thermodynamics of the gaseous halo. Feedback episodes can cause gas to deviate from hydrostatic equilibrium (Oppenheimer 2018), ultimately altering the morphology (Kauffmann et al. 2019) and kinematics (Nelson et al. 2015) of gas in the CGM.

Furthermore, SMBH feedback contributes to the enrichment of the metal content of the CGM (Sanchez et al. 2019), and to the abundance of various ionization species of these metals (Segers et al. 2017). For example, using the IllustrisTNG simulations, Nelson et al. (2018b) showed that galaxies with more massive SMBHs, which are thus preferentially quenched and red, have a lower abundance (i.e. column density) of O VI in their CGM. The IllustrisTNG simulations have also shown that low-accretion state energy input by the central SMBH can create X-ray eROSITA-like bubbles in disc galaxies (Pillepich et al. 2021) and anisotropic X-ray signatures (Truong et al. 2021b) in the CGM gas. For less massive galaxies, winds driven by stellar feedback including supernovae are responsible for driving gas out of the galaxy and into the CGM (Fielding et al. 2017; Li & Tonnesen 2020).

Despite mounting observational data and ever-improving models, a complete picture of the physical properties of the CGM of galaxies remains remote. For example, the multiphase nature of the CGM remains a nebulous concept. We do not know the CGM mass and volume fractions in different gas phases, as a function of galaxy mass and redshift, on average, and as a function of morphological and star-formation states, and from galaxy to galaxy. Crucially, it is not currently empirically known whether all disc-like galaxies similar to our MW are surrounded by a hot, X-ray emitting gaseous halo. We have no comprehensive view of the distribution, structure, or kinematics of the bulk of the CGM surrounding galaxies less massive than brightest cluster/group galaxies.

In this work, we provide insights on the properties of the CGM of MW-like galaxies by using the outcome of the TNG50 simulation of the IllustrisTNG project. We quantify the amount and distribution of the CGM gas, and its key physical properties including temperature, density, entropy, pressure, metallicity, and magnetic field strength, as well as observable signatures including H I, various metal ions, and X-ray emission. We investigate the connection between CGM properties and galaxy global properties, including star formation and recent feedback from central SMBHs. Importantly, we characterize and contrast the gaseous haloes of a large sample of TNG50 MW-like galaxies, numbering 132 in total. This allows us to quantify the galaxy population-wide median behaviour, as well as study the amount, and origin, of the scatter in trends as galaxies evolve in unique ways.

This paper is organized as follows: Section 2 describes the simulations, the galaxy sample, and our definitions and analysis/methodological choices. We present our main results in Section 3, and summarize our conclusions in Section 4.

2 METHODS

2.1 The TNG50 simulation

In this paper, we present results from the TNG50 simulation (Pillepich et al. 2019; Nelson et al. 2019b) of the IllustrisTNG suite (hereafter TNG; Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018a; Pillepich et al. 2018b; Springel et al. 2018), a set of cosmological magneto-hydrodynamical simulations. TNG is *The Next Generation* of the original Illustris simulation (Genel et al. 2014; Vogelsberger et al. 2014a,b; Sijacki et al. 2015), which was one of the first cosmological simulations to reproduce a reasonably realistic, diverse population of galaxies as in the observed universe. TNG was run using the AREPO moving-mesh code (Springel 2010), with a new underlying physics model including magneto-hydrodynamics (Pakmor, Marinacci & Springel 2014), and modified feedback processes (Weinberger et al. 2017; Pillepich et al. 2018a).

TNG50 employs a (periodic) box of side length ∼50 comoving Mpc (cMpc). Despite the relatively large box, the resolution is comparable to many zoom-in simulations, with an average baryonic (dark matter) mass resolution of $\sim 8 \times 10^4\,\mathrm{M}_\odot$ ($4 \times 10^5\,\mathrm{M}_\odot$). This unique combination of resolution and volume allows one to study highly resolved structures over a large sample size of galaxies.

The TNG simulations evolve from $z = 127$ to present day ($z = 0$), and their output is stored at 100 points in time between $z = 20$ and 0. As detailed in Pillepich et al. (2018a), TNG employs an extensive and well-validated galaxy formation model. We briefly mention here the key features of the model related to star formation. Stars form stochastically from gas cells above a threshold of neutral hydrogen density $n_\mathrm{H} \gtrsim 0.1\,\mathrm{cm}^{-3}$, according to the two-phase ISM model of Springel & Hernquist (2003). Gas colder than 10^4 K is not explicitly modelled in the TNG simulations. Metal enrichment due to stellar evolution and supernovae is also included, as well as the transfer of supernovae feedback energy in the form of a kinetic, galactic-scale decoupled wind (Springel & Hernquist 2003).

In the TNG model, SMBHs are 'seeded' when the mass of a friends-of-friends halo exceeds $\sim 7 \times 10^{10}\,\mathrm{M}_\odot$. Starting with a mass of $\sim 1 \times 10^6\,\mathrm{M}_\odot$, SMBHs grow over time by merging with others, and by accreting nearby gas with the accretion rate set as the minimum of the Bondi and Eddington accretion rates. As described in detail in Weinberger et al. (2017), the accretion rate then determines the mode in which active galactic nuclei (AGN) feedback energy is injected into the surroundings. At high-accretion rates, SMBHs continuously deposit thermal energy into neighbouring gas cells. At low-accretion rates, they instead impart kinetic energy in discrete events after a certain amount of energy is accumulated. In both modes, energy injection is isotropic at the injection scale. The former is the dominant state for less massive SMBHs, while the latter dominates for high-mass SMBHs, with the transition occurring roughly at $M_\mathrm{BH} \sim 10^8\mathrm{M}_\odot$ (Weinberger et al. 2017). Additional details on the effective functioning of the SMBH feedback in TNG and the ensuing flows can be found in Pillepich et al. (2021).

TNG adopts a cosmology consistent with the Planck 2015 analysis (Planck Collaboration et al. 2016), with: $\Omega_\Lambda = 0.6911$, $\Omega_\mathrm{m} = 0.3089$, $\Omega_\mathrm{b} = 0.0486$ and $h = 0.6764$ [$100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$].

2.2 The TNG50 MW-like galaxy sample

Among the thousands of galaxies realized within TNG50 at $z = 0$, in this paper, we focus on a sample of MW-like systems.

To this aim, we adopt the MW/M31 selection and sample as presented in Pillepich et al. (in prep), and first used in Engler et al. (2021) and Pillepich et al. (2021). At $z = 0$, this sample is based on galaxies (1) having a stellar mass, measured within a 3D aperture of 30 kpc in the range $10^{10.5}$ to $10^{11.2}\,\mathrm{M}_\odot$, (2) being discy, either based on a constraint on the minor-to-major axis ratio of the stellar mass distribution ($s < 0.45$) or through a visual inspection by-eye, of stellar light maps (see Pillepich et al. in prep for details), (3) not being in overly massive haloes ($M_{200,\mathrm{c}} < 10^{13}\,\mathrm{M}_\odot$), and (4) being reasonably isolated with no other galaxy having $M_\star > 10^{10.5}\,\mathrm{M}_\odot$ within a distance of 500 kpc.

In this work, to select MW-like galaxies only, we further restrict the mass range to $10^{10.5} < M_\star/\mathrm{M}_\odot < 10^{10.9}$. These criteria return a diverse sample of 132 galaxies, with total halo masses $M_{200,\,\mathrm{c}}$ between $10^{11.7}$–$10^{12.5}\,\mathrm{M}_\odot$, and virial radii $R_{200,\,\mathrm{c}}$ spanning ∼150–300 kpc. These are all the central galaxies of their underlying dark-matter and gaseous halo. Most of these galaxies have a stellar bar (Pillepich et al. in prep), exhibit eROSITA-like X-ray bubbles in their CGM (Pillepich et al. 2021), and are surrounded by a number of classical satellites consistent with the observations of the MW and similar nearby galaxies (Engler et al. 2021). Moreover, they are the result of a variety of merger and assembly histories, including several with gas-rich major mergers in the last 5 billion years (Sotillo-Ramos et al. 2022). As a reuslt, they exhibit a diversity of stellar bulge, disc, and halo properties, as well as star formation histories. Throughout the rest of the paper, we refer to this sample as 'TNG50 MW-like galaxies'.

2.3 Physical definitions

While the term CGM is widely understood as the region between the galaxy and the IGM, there is no consensus for the definition of the CGM. In this paper, unless otherwise stated, we define the CGM to be the region bounded by $[0.15, 1.0] \times R_{200\mathrm{c}}$ of the corresponding halo. Our analysis always considers all gas in the simulation volume, and is not restricted to gas cells which belong to the corresponding friends-of-friends halo. To focus on the CGM of the central galaxy itself, we always exclude all gas that is gravitationally bound to satellite galaxies according to the SUBFIND algorithm (Springel et al. 2001).

Unless otherwise stated, we adopt the following definitions:

(i) Temperature of star forming gas: the TNG ISM model employs a subgrid pressurization model for star forming gas (Springel & Hernquist 2003). Due to this, the temperature of such gas is 'effective' and not directly physical. We therefore always set the temperature of star forming gas to 10^3 K, its cold-phase value.

(ii) Stellar mass of the galaxy: total mass of all stellar particles within an aperture of 30 physical kpc.

(iii) Star formation rate (SFR): rather than summing up the instantaneous SFR of gas, we instead quote the SFR as the total stellar mass formed within an aperture of 30 physical kpc, averaged over the last 1 Gyr (see Donnari et al. 2019; Pillepich et al. 2019).

(iv) Specific star formation rate (sSFR): ratio of the star formation rate to the stellar mass of the galaxy, both measured as above.

(v) Neutral atomic hydrogen (H I): while the simulation directly tracks and outputs the neutral hydrogen content of gas, we use the H_2 model of Gnedin & Kravtsov (2011) to compute the fraction of neutral hydrogen that is molecular to remove this component in order to derive the atomic H I mass, following Popping et al. (2019).

(vi) O VI density: similarly, while the total oxygen mass is directly tracked and output by the simulation, the relative abundances of different ionization species are not calculated. Following Nelson et al. (2018b), we model these metal ionization states using CLOUDY (Ferland et al. (2013), v13.03): accounting for (a) photo-ionization in the presence of a high-energy (X-ray + UV) background (the 2011 update of Faucher-Giguère et al. 2009), (b) collisions, and (c) self-shielding by highly dense gas (Rahmati et al. 2013), we iterate to equilibrium using CLOUDY's single zone mode.

(vii) X-ray luminosity: we compute X-ray emission in the energy range 0.5–2.0 keV (soft X-rays), using the APEC emission model (Smith et al. 2001). This accounts for both emission lines by highly ionized species of metals, and the continuum emission produced through bremsstrahlung.

(viii) Mass flow rate: following Nelson et al. (2019b), the mass flow rate at a radius r_0 across a shell of thickness Δr is:

$$\dot{M} = \frac{1}{\Delta r} \sum_{\substack{i=0 \\ |r_i - r_0| \leq \Delta r/2}}^{n} m_i \times v_{\rm rad,i}, \tag{1}$$

where m_i and $v_{\rm rad,\,i}$ are the mass and radial velocity of the $i^{\rm th}$ gas cell. Unless otherwise stated, we consider gas to be outflowing if it has positive radial velocity (i.e. $v_{\rm rad,\,i} > 0$), and inflowing otherwise. Throughout the paper, we adopt a fixed value of $\Delta r = 5$ kpc.

3 RESULTS

3.1 Global properties of the CGM of TNG50 MW-like galaxies

The TNG50 simulation produces a diverse set of more than 100 MW-like galaxies (at $z = 0$), which in turn exhibit diverse CGM properties both across and within galaxies.

3.1.1 CGM maps

We begin by visualizing this diversity in Fig. 1, where we show gas column density projections for a random subset of 20 galaxies/haloes.[1] Note that we show random projections of galaxies, i.e. the central gaseous discs are occasionally seen edge-on, face-on, or at intermediate inclinations. The projections extend $\pm 1.5 R_{200,\rm c}$ from edge to edge, as well as along the line of sight direction. The two white circles mark $(0.15, 1.0) \times R_{\rm 200c}$ of the corresponding halo, which we adopt as our fiducial definition for the spatial region of the CGM (see Section 2.3).

We see that the sample includes a variety of (a) galaxy/halo sizes, as can be inferred from the scale bar in the top-left of each panel, (b) stellar masses, bottom-left of each panel, and (c) most importantly for the sake of this paper, CGM gas distributions. To set a scale, the virial radii of TNG50 MW-like galaxies vary in the range $R_{\rm 200c} \sim$ 150–300 kpc, with a median value of ∼210 kpc. The distribution of CGM gas is, in general, not azimuthally symmetric (Péroux et al. 2020; Pillepich et al. 2021; Truong et al. 2021b), occasionally due to CGM gas that is directly connected to the central galactic disc. At the same time, there are often tidal features including cold gas tails and debris in the CGM, due to ram-pressure stripping of satellite gas (Ayromlou et al. 2019; Yun et al. 2019).[2]

In Fig. 2, we take a closer look at a single gaseous halo, selected as a relaxed example without significant asymmetric features. We show four different physical properties in the different quadrants. In the top-left quadrant, we show the column gas density projection: gas is most dense close to the disc, reaching $\gtrsim 10^8$ M$_\odot$ kpc^{-2} ($\gtrsim 10^{21}$ cm^{-2}), and becomes more rarified at larger distances into the halo outskirts and the IGM. While gas outside of the virial radius (i.e. outside the upper limit of the CGM) shows minimal (visible) structure, i.e. mostly smooth, we observe noticeable 'clouds' of gas in the CGM. These gas density fluctuations are local overdensities that are typically cooler than their surroundings. The top-right-hand panel shows the gas temperature: the disc is dominated by cold gas, while more extended gas is much warmer. Halo gas has a complex temperature structure with hot gas at or near the halo virial temperature ($10^{5.8-6}$ K) co-existing with colder gas structures. In the bottom-left quadrant, we show the thermal pressure $P_{\rm th} = (\gamma - 1)\,\rho\, U_{\rm th}$ of the gas, where ρ and $U_{\rm th}$ are the density and internal energy of gas, respectively. For star forming gas, note that this is the pressure based on the effective equation of state. A clear radial pressure gradient exists, moving from the galaxy (dominated by high-thermal pressure) to the IGM (low pressure). Finally, in the bottom-right-hand panel, we show the gas entropy ($\sim P_{\rm th}/\rho^\gamma$). Gas in the disc is at relatively low entropy, while the IGM hosts gas at higher entropies. Similar to the other quantities discussed above, gas in the CGM shows a mix of the two: while most of the volume is filled by gas at high entropy, small overdensities at lower values of entropy are realized by the TNG50 model (see also Nelson et al. 2020).

Fig. 3 shows four additional views of the same halo, focusing on quantities that are either directly observable, or commonly inferred from observables of the gas. In the top-left quadrant, we show the H I column density: while $N_{\rm H\,I}$ is qualitatively similar to total gas column density, it decreases more rapidly as one transitions from the disc towards the IGM, because self-shielding from the metagalactic background radiation field becomes negligible at densities characteristic of the outer CGM (Rahmati et al. 2013). The top-right quadrant shows the gas metallicity: while the disc is dominated by

[1] See www.tng-project.org/ramesh22 for the CGM maps of all our TNG50 MW-like galaxies.

[2] Recall that satellites are intentionally excluded in our analysis (Section 2.3), and hence are missing in these gas projections, but several would in general be expected in $z = 0$ MW-like haloes (Engler et al. 2021).

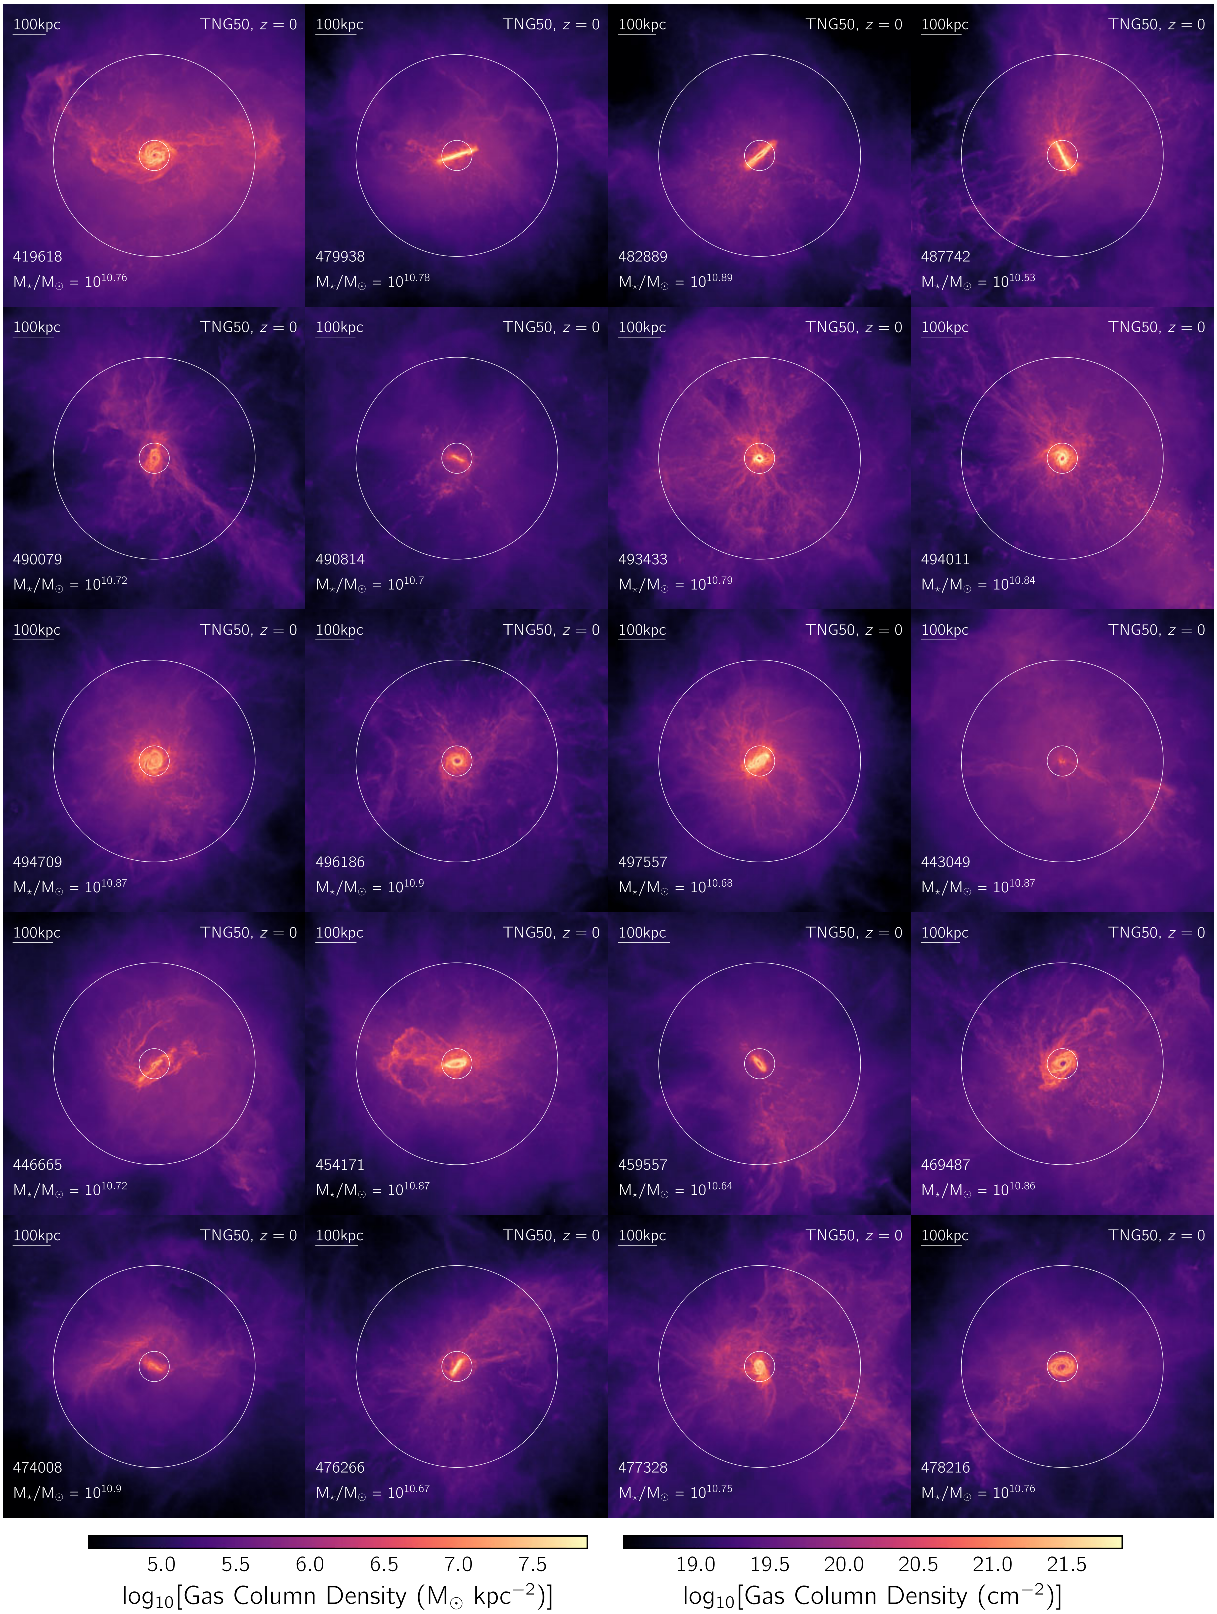

Figure 1. Gas density projections at $z = 0$ (along random orientations in a box that extends $\pm 1.5 R_{200c}$, along the perpendicular axis) of 20 random TNG50 MW-like galaxies, and their surrounding CGM excising satellites. In each panel, the two circles mark $(0.15, 1.0) \times R_{200c}$ of the parent halo – we refer to the region between the two circles as the CGM. On the bottom-left of each panel, the numbers indicate the subhalo ID, and stellar mass of the corresponding (central) galaxy. These halo-scale gas density projections reflect the wide variety we find across the 132 MW-like TNG50 galaxies in our sample.

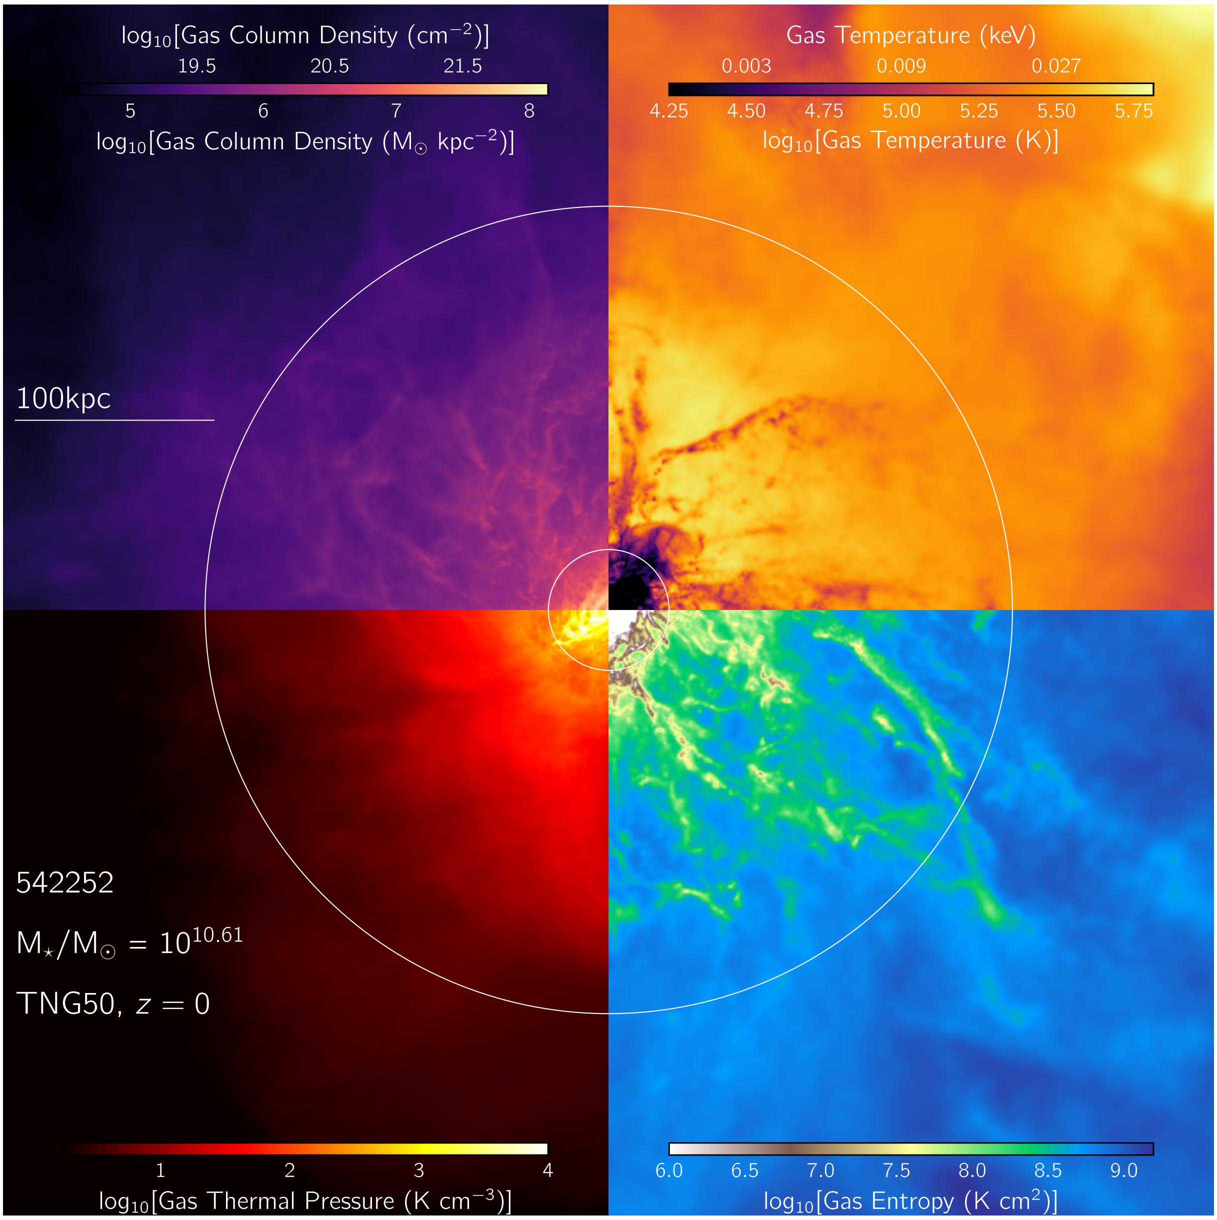

Figure 2. Visualization of four different physical quantities, each shown in one quadrant, for a representative gaseous halo around a MW-like galaxy (subhalo ID 542242) from the TNG50 simulation at $z = 0$. The orientation is random, the projections extend $\pm 1.5 R_{200,\rm c}$ perpendicular to the image plane, each quadrant extends $\pm 0.75 R_{200,c}$ in extent from side to side, and satellite gas is excised. As above, the two circles mark $(0.1, 1.0) \times R_{200\rm c}$ of the corresponding halo. The panels show: gas column density (top left), gas temperature (top right), gas pressure (bottom left), and gas entropy (bottom right). While the galaxy at the centre of the halo is dominated by dense cold gas with large pressure, the CGM has different physical properties – halo gas is, on average, hotter and less dense. There are strong inhomogeneities on top of the background state, which are overdense, cool, and low entropy, i.e. 'clouds' of cool, dense gas.

metal-enriched gas $\gtrsim Z_{\odot}$, the CGM host regions that are significantly less metal enriched at $0.1 Z_{\odot}$, and lower, in fact, columns of highly metal-enriched gas typically emerge above and below the galactic disc (Pillepich et al. 2021), as it can be seen at the edge of the quadrant. The bottom-left (bottom-right) quadrant shows O VI column density (X-ray luminosity), which is an indirect tracer for warm (hot) gas. $N_{\rm O\,VI}$ declines from $\sim 10^{16}$ to $\sim 10^{13}$ cm^{-2} from 0.15 to 1 $R_{200,\rm c}$, while $L_{\rm X}$ declines rapidly with distance primarily due to its n^2 dependence, dropping ten orders of magnitude, from $\sim 10^{40}$ to $\sim 10^{30}$ erg s^{-1}, over the same distance. Both these tracers of the warm and hot phases are more diffuse/volume-filling than the cold gas traced by neutral hydrogen. Although this example is typical, significant variation across the sample is evident.

3.1.2 *CGM integrated properties*

To explore this diversity in CGM gas properties among MW-like galaxies, Fig. 4 presents trends of integrated properties of CGM gas as a function of stellar mass. The markers show the values of individual galaxies, coloured by specific star formation rate (sSFR), while lines plot the median as a function of mass.

The top-left-hand panel of Fig. 4 shows the total CGM gas mass, where the circles/black median curve correspond to our fiducial definition of the CGM, i.e. the region bounded by $[0.15, 1.0] \times R_{200\rm c}$. For comparison, in crosses/grey median line, we also include an alternate definition of the CGM: the region bounded by [30,

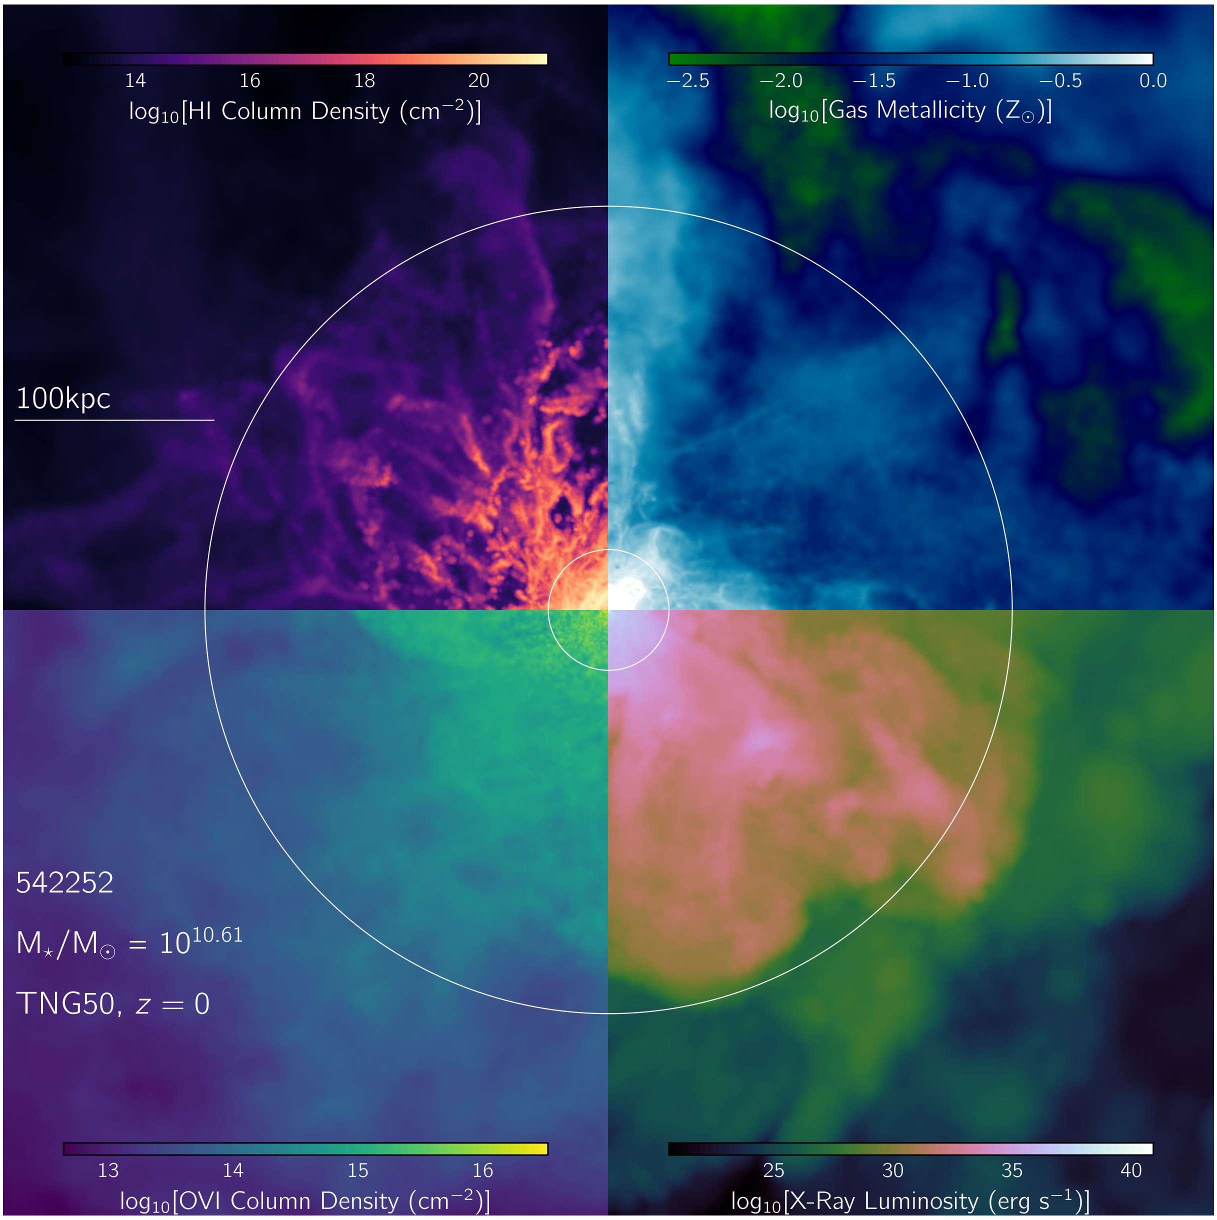

Figure 3. Visualizations of four additional physical properties for the same gaseous halo shown in Fig. 2. Here we focus on more directly observable quantities. We show: H I column density (top left), gas metallicity (top right), O VI column density (bottom left), and X-ray luminosity in the 0.5–2.0 keV band (bottom right). In general, the CGM of MW-like galaxies can host regions of much lower metallicity than the galaxy (also at less than one percent $Z_\odot$: green). Cold gas, as traced by neutral hydrogen, is particularly clumpy, while the warm-hot phase traced through O VI and X-rays is much more spatially smooth.

200] kpc. While differences between the two definitions are negligible at low-stellar masses, more massive galaxies, on average, are surrounded by more extended haloes, and defining the outer boundary of the halo at 200 kpc reduces the amount of gas that we can measure in the CGM. To best capture the CGM across our sample, which covers a range of dark-matter halo masses and thus sizes, we therefore adopt the R_{200c}-relative definition as our fiducial choice.

The median total CGM gas mass is almost constant as a function of stellar mass with a difference of $\sim$0.1 dex across the stellar mass range. Two more important trends emerge: (a) at fixed stellar mass, there is a significant diversity in amount of gas present in the CGM, as visible in the scatter, and (b) galaxies with a higher sSFR preferentially have more gas in their CGM, i.e. blue points are systematically above their green and red counterparts. Conversely, TNG50 MW-like galaxies below the star-forming main sequence are surrounded by less massive CGMs than their more star-forming counterparts. We have checked that this trend is not driven by halo mass, i.e. it is in place even though more massive haloes tend to host more massive CGMs and even though, in the TNG model, at fixed stellar mass, quiescent galaxies tend to reside in somewhat more massive haloes.

To make a comparison with other models, we show a similar measurement of the (total) CGM gas mass (at $z = 0$) from the Auriga simulations, a set of zoom-in simulations run with AREPO, albeit with a different underlying galaxy physics model. Shown in the orange diamond, the CGM of Auriga Halo 6 (van de Voort et al. 2021, hereafter, Au6) – a relatively isolated MW-like galaxy – is comparable to

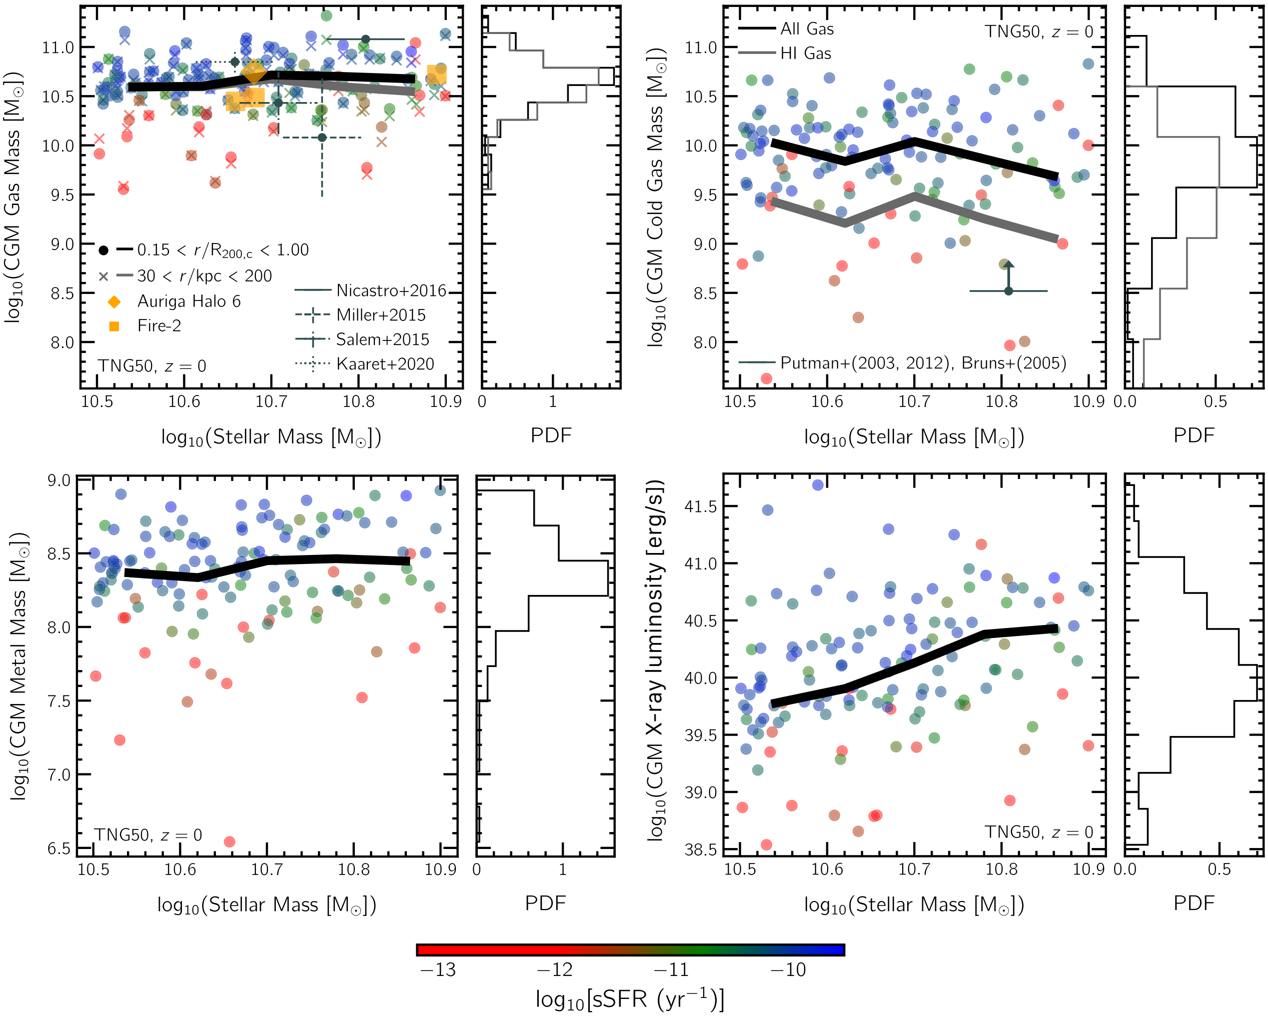

Figure 4. Selected integral properties of the CGM plotted as a function of galaxy stellar mass for all 132 MW-like galaxies from the TNG50 simulation at $z = 0$: total CGM gas mass (top-left), CGM cold gas (T $< 10^{4.5}$ K) mass (along with the median line of CGM H I mass; top-right-hand panel), CGM metal mass (bottom-left), and CGM X-ray luminosity (0.5–2.0 keV; bottom right-hand panel). The first three quantities depend only weakly on galaxy stellar mass, while X-ray luminosity rises sharply with increasing stellar mass of the galaxy. In each panel, the circular markers are colour coded by the sSFR of the galaxy – in all cases, we note that galaxies with a lower sSFR have smaller values of the corresponding quantity. In the top-left-hand panel, we consider two different definitions for the CGM: variable volume (extends between 0.15 and 1.0 $R_{200,\,\rm c}$) and constant volume (extends between 30 and 200 kpc) – the two definitions yield statistically similar results at the low mass end, while differences are apparent towards the high mass end. In all other panels, we adopt the variable volume definition, which is our fiducial choice.

the median behaviour of the TNG50 sample at that stellar mass.[3] We also show properties of three haloes from the FIRE-2 simulations, run with the code Gizmo (Hopkins 2015) and with rather different physical models compared to TNG – the (total) CGM gas mass of m12b, m12c, and m12w, all at $z = 0.25$, are shown in orange squares (Esmerian et al. 2021).[4] The measurements from FIRE-2 are roughly comparable with TNG50, although the two galaxies at $\sim 10^{10.65}$ M$_\odot$ are lower by ∼0.2 dex in their CGM gas mass content with respect to the TNG50 median, albeit well within the TNG50 scatter, which could be due to more powerful/ejective stellar feedback.

On the observational side, the total gas mass of the CGM of the MW is highly uncertain: while Nicastro et al. (2016) report that this value may be as high as 1.2×10^{11} M$_\odot$, Miller & Bregman (2015) estimate the total mass of CGM gas out to 250 kpc to be $1.2^{+0.9}_{-0.8} \times 10^{10}$ M$_\odot$. Through a study of the orbital motion of the Large Magellanic Cloud (LMC), Salem et al. (2015) infer the mass of CGM gas out to 300 kpc to be $2.7^{+1.4}_{-1.4} \times 10^{10}$ M$_\odot$. Using X-ray observations, Kaaret et al. (2020) estimate the mass out to 260 kpc to lie in the range $(5.5–8.6) \times 10^{10}$ M$_\odot$. We show these measurements with dark grey solid, dashed, dot–dashed, and dotted errorbars, respectively. In each case, for the stellar mass of the MW, we use the measurement reported by McMillan (2011): $M_{\star,\rm MW} \sim (6.43 \pm 0.63) \times 10^{10}$ M$_\odot$, and spread these four data points somewhat along the x-axis for visibility. All four of these observational inferences are well bracketed within the TNG50 scatter.

In the top-right-hand panel of Fig. 4, we show the cold gas mass (T $< 10^{4.5}$ K) in the CGM. In addition, the grey curve shows the median H I gas mass, which resembles a down-scaled version of the black curve. As with the total CGM gas mass, the trends of cold gas mass are more or less constant as a function of stellar

[3] Note that van de Voort et al. (2021) defines CGM gas mass as the total mass of all non-star forming gas within the virial radius, a minor difference.
[4] Esmerian et al. (2021) considers the outer boundary of the CGM to be 1.0 $R_{\rm vir}$, while the inner boundary is set to max $(1.2\,R_{\rm gal}, 0.1\,R_{\rm vir})$.

mass. There is again a noticeable gradient in the sSFR colours across the vertical scatter: galaxies with a higher sSFR have substantially more cold gas in their CGM. We have checked, and the trend with sSFR is present also when adopting a larger inner boundary of the CGM (0.3 instead of 0.15 × $R_{200,\,c}$), i.e. by ensuring that extended cold gaseous discs – which may be expected in highly star-forming galaxies – do not 'contaminate' our definition of CGM.[5] Moreover, although not shown here, such a gradient is not present when points representing MW-like galaxies are coloured by either the halo mass or SMBH mass, but it does appear to depend on the cumulative energy injected via the kinetic mode of SMBH feedback.

Observationally, Putman, Peek & Joung (2012) estimate the cold (H I) gas mass contained in high-velocity clouds (HVCs), using 21-cm emission to be $\sim 2.6 \times 10^7\,\mathrm{M}_\odot$. Putman et al. (2003) and Brüns et al. (2005) report that the Magellanic System (i.e the Magellanic Stream, Leading Arm and Magellanic Bridge) contributes at least $\sim 3 \times 10^8\,\mathrm{M}_\odot$ to the H I mass of the MW, resulting in a *lower limit* to the total H I mass of $\gtrsim 3.3 \times 10^8\,\mathrm{M}_\odot$ for the MW CGM. This measurement is shown with the grey solid errorbar. Once again, we use the measurement reported by McMillan (2011) for the stellar mass. Given this lower limit, the TNG50 median is ~0.5 dex above, although some outliers from our sample have consistent H I masses, as apparent with the marginalized distributions on the right.

The total amount of metals in CGM gas is shown in the bottom-left-hand panel of Fig. 4: while galaxies with a higher (lower) sSFR have more (fewer) metals in the CGM gas, the metallicity content in CGM gas shows no strong trend with respect to stellar mass. Since the stellar mass range investigated is rather small (0.4 dex), such a flat trend may not be altogether surprising even though there is a slight trend of higher CGM metallicity in higher mass haloes within this galaxy sample (not shown).

Lastly in the bottom-right-hand panel of Fig. 4, we show the X-ray luminosity of CGM gas. We see a clear increase in luminosity as a function of stellar mass, by a factor of ~10 across the stellar mass range explored. Additionally, at fixed stellar mass, higher sSFR galaxies have more luminous gaseous haloes in the 0.5–2.0 keV (soft) X-ray band. Similar to other quantities, there is a large scatter around the median value, i.e larger than the mass trend itself. While the computed X-ray luminosity includes continuum emission, plasmas at temperatures typical of MW-like systems are dominated by metal line emission (Anderson et al. 2015), making the metal content of the CGM, and the enrichment history of the galaxy critically important for extended X-ray haloes. The bottom-right-hand panel of Fig. 4 confirms, also for the case of discy TNG50 MW-like galaxies, the predictions previously put forward by Truong et al. (2020) and Oppenheimer et al. (2020) who showed that, according to both TNG and EAGLE, at the transitional mass scale of $\sim 10^{10.5-11}\,\mathrm{M}_\odot$ in stars, star-forming galaxies are surrounded by X-ray brighter haloes than similarly-massive more quiescent galaxies – a phenomenon dubbed the X-ray luminosity halo dichotomy.

Overall, Fig. 4 demonstrates that according to TNG50 and even within the relatively narrow selection of MW-like galaxies, the abundance and physical properties of gas in the CGM of MW-like haloes exhibit a large diversity. Although three of the four integral properties studied here have a weak dependence on stellar mass, across the stellar mass range of $10^{10.5}$–$10^{10.9}\,\mathrm{M}_\odot$, the star formation activity or status of the central galaxy appear closely linked to the state of its CGM.

Following previous works (see e.g. Davies et al. 2020; Terrazas et al. 2020; Truong et al. 2020, 2021b; Zinger et al. 2020), we have checked that, particularly in the intermediate stellar mass range of MW-like galaxies, a connection of CGM gas properties also exists with the (cumulative, in time) SMBH energy injection in the kinetic mode. We find that more energetic SMBH kinetic feedback implies lower CGM mass, CGM cold mass, CGM metal mass, and X-ray luminosity (not explicitly shown). In fact, within the TNG model, SMBH kinetic feedback – rather than SMBH thermal energy injection at high-accretion rates or stellar feedback – is the cause for the suppression of star-formation in massive galaxies (Weinberger et al. 2017; Nelson et al. 2018a; Terrazas et al. 2020). In fact, there would be no low-SFR galaxies in our TNG50 MW-like sample if SMBH kinetic feedback was switched off.

We therefore conclude that the modulations in total CGM mass, CGM metal mass, and X-ray luminosity of MW-like galaxies (see Truong et al. 2020) are primarily due to the ejection and heating of CGM gas driven by the SMBH kinetic feedback in TNG (Nelson et al. 2019a; Zinger et al. 2020). Star formation status is to zeroth order a proxy for this SMBH activity. On the other hand, we speculate that the gradient with sSFR in the *cold* gas mass in the CGM may arise due to a physical connection between the CGM and galactic star formation. This could be caused by a link with the flow of cold gas around galaxies, although the direction of causation is unclear. A greater amount of cold halo gas may, for example, imply larger gas accretion rates on to the galaxy, resulting in higher sSFRs.

3.1.3 The baryon budget of the CGM

Inspired by efforts to make an empirical consensus of the baryonic content of haloes (Peeples et al. 2014; Werk et al. 2016), we explore the budget of baryons across different components and gas phases. Fig. 5 shows the baryonic components of galaxies/haloes, split into seven different categories. We calculate each as a fraction with respect to the halo virial mass multiplied by the cosmic baryon fraction, $M_{\mathrm{comp}}/(f_{\mathrm{b}}\,\mathrm{M}_{200\mathrm{c}})$. We consider seven different components: CGM gas split into cold ($T < 10^{4.5}$ K), warm ($10^{4.5}\,\mathrm{K} < T < 10^{5.5}$ K), and hot (T $> 10^{5.5}$ K) phases, galaxy stellar mass, galaxy gas mass, CGM metal mass, and galaxy metal mass. Each MW-like galaxy in our sample is represented by a single bar with its components stacked vertically, and arranged in ascending stellar mass from left to right (see top x-axis). For better visibility, the latter two components have been multiplied by a factor of fifty. To avoid double-counting gas cells between the galaxy and the CGM in this figure (only), we use the term galaxy to refer to the part of the halo within the inner boundary of the CGM, i.e $< 0.15 \times R_{200\mathrm{c}}$.

Galaxies with similar stellar masses can vary significantly in terms of how their baryons are distributed across different components: while one galaxy may be surrounded by a CGM that is dominated by hot gas, a different galaxy with a similar stellar mass can instead be surrounded by a largely cold gas dominated CGM. The median fractional mass of gas in the CGM with respect to $f_{\mathrm{b}} M_{200\mathrm{c}}$ is ~29

[5]This is the case for all four CGM integrated properties of Fig. 4. While setting the inner boundary of the CGM to 0.15 × R_{vir}, sometimes leads to a part of the disc being included in the CGM (as can be seen in some panels of Fig. 1), we confirm here that this does not strongly impact the conclusions of Fig. 4. For instance, setting the inner boundary of the CGM instead to 0.3 × R_{vir} leads to a systematic drop of ~0.5 dex in values of cold gas mass, metal mass, and X-ray luminosity. However, the median trends and the dependencies on star formation rates remain almost unchanged, albeit somewhat weakened in the case of X-ray luminosity.

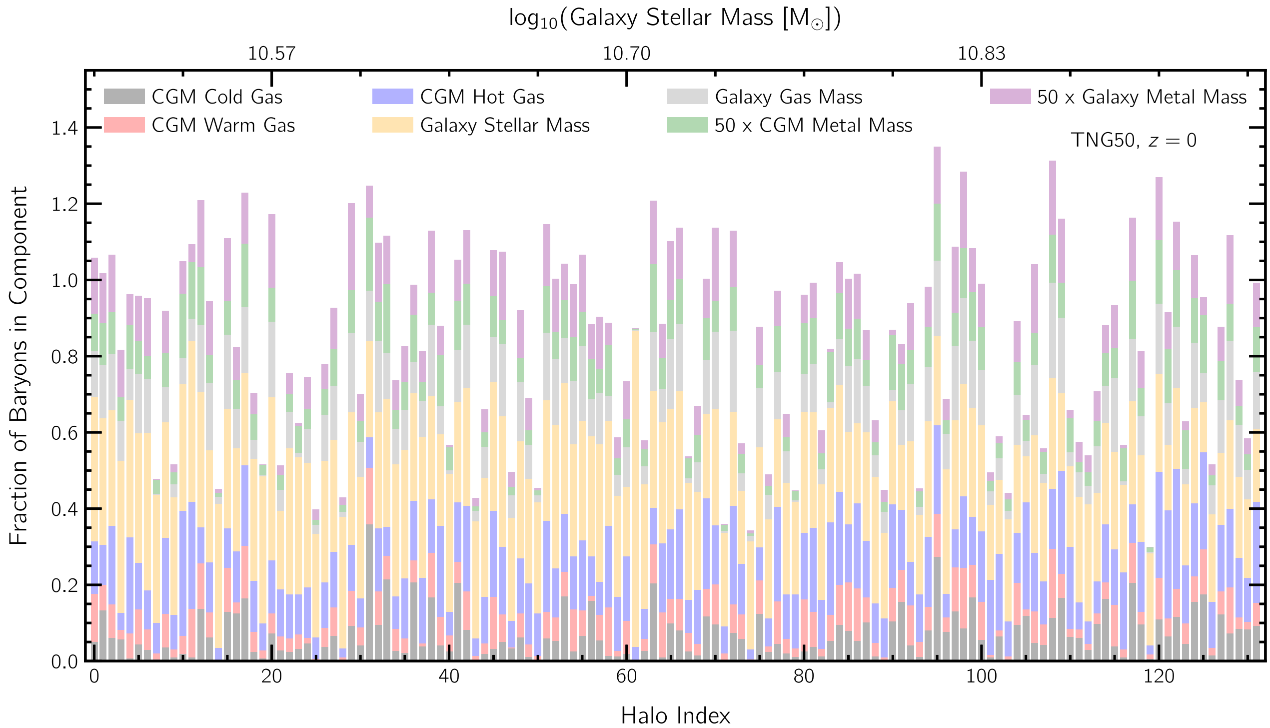

Figure 5. Fraction of baryonic matter in different components and phases, relative to the expected baryon budget of the halo, i.e. $M_{\rm comp}/(f_{\rm b}\,{\rm M}_{\rm 200c})$ at $z = 0$. Each of the 132 MW-like galaxies from TNG50 is represented as an individual vertical bar. The galaxies/haloes are arranged in *ascending order* of their galaxy stellar mass from left to right (see top x-axis). The different components shown are: CGM cold gas (grey), CGM warm gas (red), CGM hot gas (blue), galaxy stellar mass (orange), galaxy gas mass (grey), CGM metal mass (green), and galaxy metal mass (purple). Note that the latter two quantities have been multiplied by a factor of fifty for better visibility. We find a large diversity across the sample, with no clear trend in fractional abundances as a function of galaxy stellar mass. For this Figure, 'galaxy' refers to material within the inner CGM boundary, i.e. $<0.15R_{\rm 200c}$.

per cent, although this value varies considerably across the sample from ∼15 (16$^{\rm th}$ percentile) to ∼41 per cent (84$^{\rm th}$ percentile). This variability is also visible in the fractions of the different components: the median values and percentiles for the fractional mass of cold gas, warm gas, hot gas, and metals, all in the CGM are 5^{+7}_{-4}, 6^{+5}_{-4}, 14^{+8}_{-6}, and $0.17^{+0.08}_{-0.10}$ per cent, respectively. Because of this significant galaxy to galaxy variation, no strong trends as a function of stellar mass are evident. However, generally, the CGM mass of MW-like galaxies is largely dominated by hot, and thus X-ray emitting gas.

3.2 Spatial distribution of gas around TNG50 MW-like galaxies

Having explored various global properties of the CGM, we now consider how gas in the CGM is distributed, and how its properties vary across the CGM.

To begin, Fig. 6 shows spherically-averaged radial profiles as a function of galactocentric distance. We include density (upper left), H I density (upper right), temperature (centre left), metallicity (centre right), thermal pressure (lower left), and entropy (lower right). Each (thin) curve shows a single profile of an individual halo (mass-weighted median for temperature, metallicity, pressure, and entropy), whereas the black curve shows the median behaviour across all haloes in the sample. The individual curves are coloured according to the stellar mass of the corresponding galaxy, as shown by the colourbar at the bottom. While the (lower) x-axes of all panels show the galactocentric distance normalized by the virial radius of the corresponding halo, the upper x-axis on the top-left-hand panel shows the (median) galactocentric distance in physical kpc units, for reference.

The total gas density profile (top-left-hand panel of Fig. 6) peaks close to the disc, i.e. at the inner radius of the CGM, and decreases rapidly outwards into the halo, where the average CGM density can become as low as $10^{2-3}\,{\rm M}_\odot\,{\rm kpc}^{-3}$. The median density profile is well fit by a function $r^{-\alpha}$, where r is the galactocentric distance, and $\alpha \sim 0.36$. Perhaps unexpectedly (see previous Section), galaxies with larger stellar masses tend to have smaller gas densities in their CGM: a difference of ∼0.1 (outer halo) to ∼0.3 dex (inner halo) when the sample is split into two bins based on stellar mass (not explicitly shown). Observationally, based on ram pressure stripping of the LMC, Salem et al. (2015) constrain the CGM density of the MW at 48 ± 5 kpc to be $1.1^{+0.44}_{-.45} \times 10^{-4}$ cm^{-3}, while Kaaret et al. (2020) estimate the density of the MW CGM at the virial radius to be $(4.8 \pm 1.0) \times 10^{-5}\,{\rm cm}^{-3}$. These are shown in red solid and dashed errorbars, respectively. The inner halo measurement is below the median TNG50 profile, though well within the scatter, while the outer halo measurement is ∼0.3 dex higher than that expected average across the TNG50 MW-like sample.

The H I density (top-right-hand panel, Fig. 6) behaves qualitatively similarly to the total gas density, although the drop in density is much sharper as self-shielding becomes ineffective in the low-density environment of the outer halo, as discussed in Section 3.1. Galaxies with larger stellar masses generally have lower H I density profiles, and the trend is stronger in comparison to the total gas density: a difference from ∼0.2 (outer halo) to ∼1 dex (inner halo) when splitting the sample into two bins at the median.

The centre left-hand panel of Fig. 6 shows gas temperature: while gas close to the disc is predominantly cold ($<10^{4.5}$ K), typically (see the solid black curve representing the

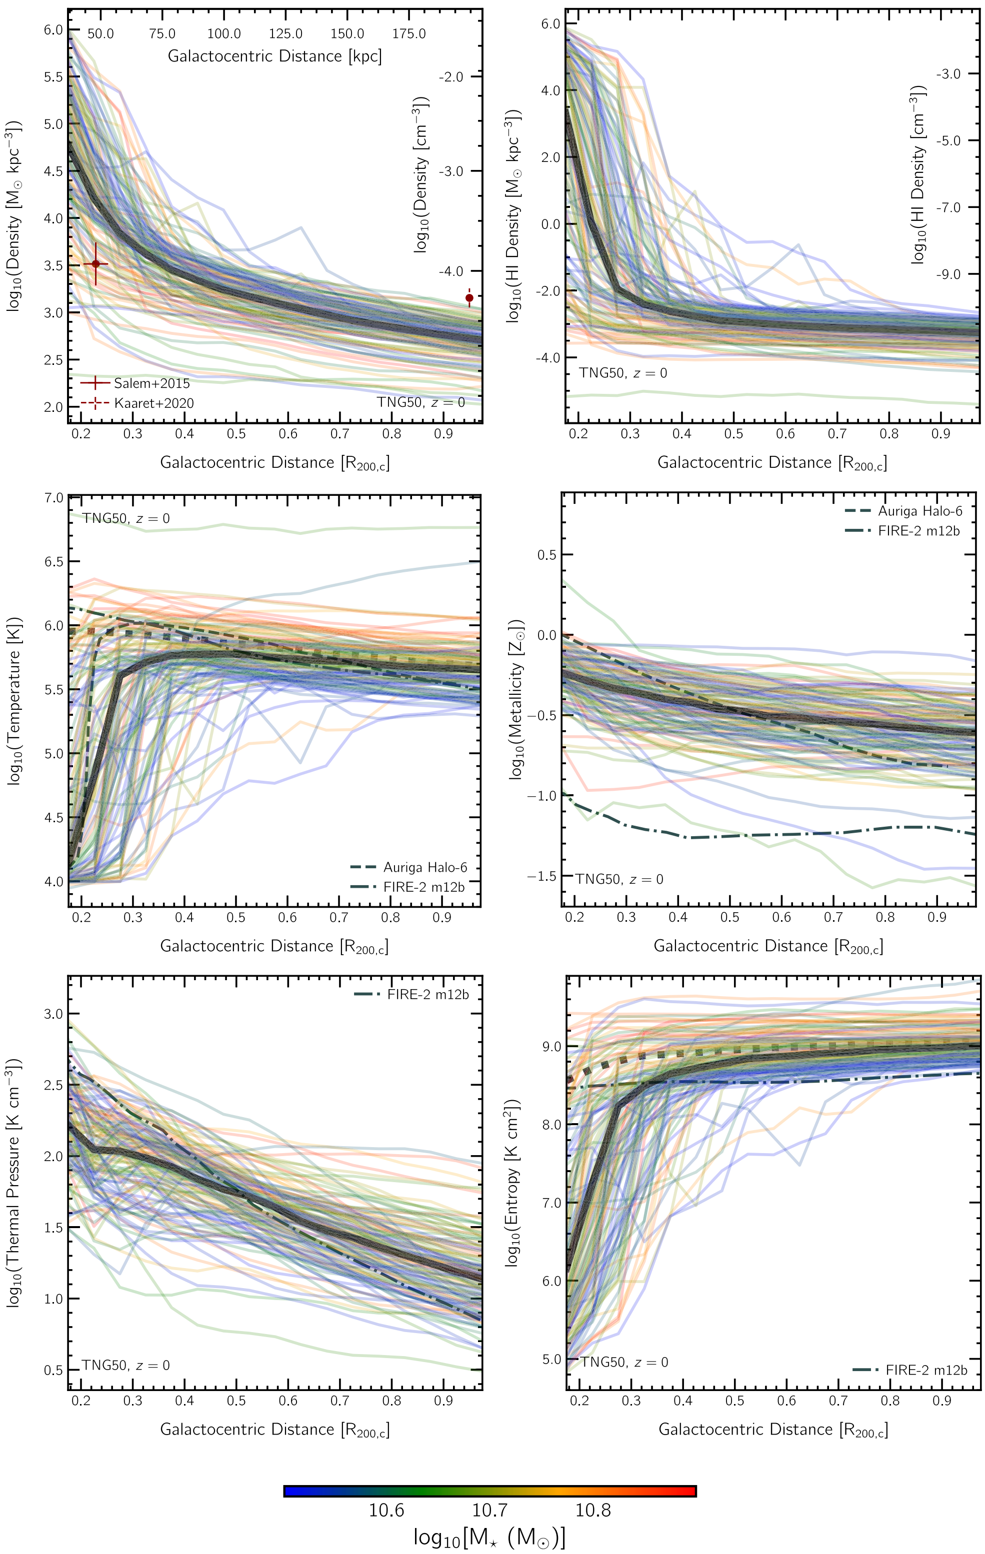

Figure 6. Spherically-averaged radial profiles of six physical gas quantities as a function of (3D) galactocentric distance around TNG50 MW-like galaxies. We consider density (upper left), H I density (upper right), temperature (centre left), metallicity (centre right), thermal pressure (lower left), and entropy (lower right). Each panel shows the median profile across the galaxy sample (black line), while the thinner curves correspond to profiles of individual haloes, coloured by galaxy stellar mass. Such profiles represent the spherically-averages of the depicted integrated or mass-weighted gas property at the given radius. Whereas the temperature and entropy of gas is generally lower in the inner halo as compared to the outskirts, due to the presence of a second cooler component becoming dominant, no such trend is clearly seen in the other four panels. The galaxy-to-galaxy scatter of individual curves about the median profile is large, and often correlates well with stellar mass.

population-wide average) the temperature rises sharply to the virial temperature (around 0.25–0.3 R_{200c}) before slowly declining outwards to larger galactocentric distances. This behaviour is a result of cold gas dominating the disc-halo interface. To demonstrate this, we also show the median trend including only gas hotter than 10^5 K (thick dotted black curve), a lower cut-off of roughly $0.1T_{\rm vir}$ which excludes the cold phase. In the inner halo, this curve is higher and flatter than the profile shown in the solid black curve with no central depression. A noticeable trend is also seen with respect to stellar mass: lower-mass galaxies have cooler temperature profiles, with a difference of ∼0.3 dex between the two bins when the sample is split at the median stellar mass. This trend is expected, since both the galaxy stellar mass and halo virial temperature scale with halo mass.

For comparison, in the temperature-profiles panel, we include results from the Auriga simulations (Au6 at $z = 0$; van de Voort et al. 2021) – these are shown in dashed dark grey curves. While Au6 reaches its peak temperature at roughly the same distance as the TNG50 median (∼$0.3R_{200c}$), its value is slightly higher (∼0.3 dex), although the virial radius of Au6 is close to the average virial radius of the TNG50 sample at ∼210 kpc. We also show a profile from one of the FIRE-2 simulated galaxies (m12b; Esmerian et al. 2021).[6] Shown in dark grey dot–dashed curves, the gas temperature profile of this galaxy is higher in the inner halo ($\lesssim 0.45R_{200c}$) in comparison to the TNG50 median, possibly due to stronger heating from the supernovae feedback model of FIRE.-hand

We quantify trends of gas metallicity in the centre-right-hand panel of Fig. 6: gas close to the disc is more enriched than farther away into the halo, although the radial change is moderate, declining by ∼0.3 dex from 0.15 to 1.0 R_{200c}. On average, haloes hosting galaxies with lower versus higher stellar mass have roughly the same metallicity out to a distance of ∼$0.4R_{200c}$, with more massive galaxies having greater gas-phase metallicity beyond this distance, up to a peak difference of ∼0.2 dex when splitting the sample into two bins at the median stellar mass. Although the comparison has caveats, the Au6 metallicity profile is qualitatively different than the median TNG50 behaviour: it is steeper-hand with higher values within the inner halo ($\lesssim 0.5R_{200c}$), but lower values in the outer halo ($\gtrsim 0.5R_{200c}$). The Au6 profile is, none the less, bracketed within the population diversity of the TNG50 sample, as indicated by the scatter among the coloured lines. The FIRE-2 galaxy (m12b) profile, on the other hand, exhibits a considerably metal-poorer CGM in comparison to the TNG50 median (a difference of ∼0.6–0.8 dex, depending on distance): this is possibly because only hot gas is included, and/or is related to the differing details of metal mixing in the gaseous haloes between the numerical techniques of FIRE-2 versus TNG.

Returning to TNG50, Fig. 6 includes thermal pressure radial profiles (lower left-hand panel). As expected given the rapid decline of density to large distance, pressure is also highest towards the halo centre. Gas in the CGM of galaxies with lower stellar masses has lower thermal pressure, with an offset of ∼0.1 to ∼0.2 dex when splitting the sample into two bins at the median. The FIRE-2 galaxy thermal pressure profile is qualitatively different than the median TNG50 behaviour: it is steeper, and so has higher pressures within the inner halo ($\lesssim 0.5R_{200c}$), but lower pressure in the outer halo ($\gtrsim 0.5R_{200c}$). None the less, we can still identify a few TNG50 MW-like galaxies with similar pressure profiles.

Finally, the bottom-right panel-hand of Fig. 6 shows the behaviour of gas entropy as a function of galactocentric distance. Moving away from the disc, entropy increases sharply out to a distance of around $0.3R_{200c}$, beyond which its value increases only gradually. As with temperature, the drop at small distance is due to the dominance of cold gas in the inner halo. In the dotted black curve, we show the median trend including only gas hotter than 10^5 K. In the inner halo, this curve is once again higher than the profile shown in the solid black curve, and asymptotically approaches the latter at large distances ($\gtrsim 0.5R_{200c}$). Lower-mass galaxies have relatively lower entropy profiles with differences of ∼0.3 to ∼0.4 dex across the sample. Except for the inner halo ($\lesssim 0.3R_{200c}$), the FIRE-2 galaxy profile is lower by ∼0.4 dex than the median TNG50 profile.

In each of the above panels, except for marginal cases, most haloes behave similar to the median except with an overall normalization offset, and albeit with a large scatter.

3.2.1 *Radial profiles of the gas beyond azimuthal averages*

The previous analysis of radial profiles has one major limitation: spherically-averaged quantities (at a given distance) do not carry information about non-radial structure, i.e. they remove all detail about the diverse nature of gas at various angular positions at a given (galactocentric) distance. To explore this, Fig. 7 shows two-dimensional distributions of the same six physical quantities of Fig. 6, albeit only for one TNG50 MW-like galaxy (the same halo visualized in Figs 2 and 3) – averaging over the entire sample would hide any localized features. The different colours in each panel show the (relative) amount of gas mass, such that the full distribution function of each physical property at each radius is visible; black curves show median trends, namely the spherically-averaged quantity as in the thin curves of Fig. 6.

Some physical properties shown in Fig. 7, notably total gas density and thermal pressure (top left-hand and bottom-left-hand panels) are more or less continuous, meaning that the distribution of these quantities is unimodal and well characterized by the median radial value. However, this is not always the case. For instance, gas at a distance of ∼$0.25R_{200c}$ is separated into two different components: a cold phase making up the star-forming body of the galaxy in the form of a thin disc-like morphology, which extends in the plane of the disc to these distances, and a hot phase heated to the virial temperature outside of the galaxy itself, i.e. beyond the disc extent and/or just above or below its vertical extent. Similar features are visible in the H I density – a significant amount of neutral hydrogen close to the disc, which is otherwise ionized in gas of higher temperatures, and entropy – rotationally-supported cold gas has lower entropy compared to the hot gas in the halo.

The distributions of gas temperature (centre left) and entropy (lower right-hand panel) show large scatter in values at fixed distance, namely large variations from one location within the CGM to another. For example, at $0.4R_{200c}$, the temperature of gas varies by over two orders of magnitude; on the other hand, the H I density can vary by ∼nine orders of magnitude, while the spread in metallicity ranges from supersolar to nearly pristine at fixed distance, indicative of extremely-incomplete metal mixing in the CGM as well as the accretion of hardly enriched gas.

Fig. 7 clearly shows that, according to TNG50 and at fixed distance, different parts of the halo, i.e. at different angular positions, exist in very different physical states. This is the case not only for the specific galaxy depicted there, but applies to each MW-like galaxy in TNG50. For temperature and entropy, we again include dotted black curves showing the median trend when only gas hotter than 10^5 K

[6]The FIRE-2 profile is available at $z = 0.25$, and is restricted to the hot halo only, by excluding substructure and low entropy gas ($K < 5$ keV cm^2).

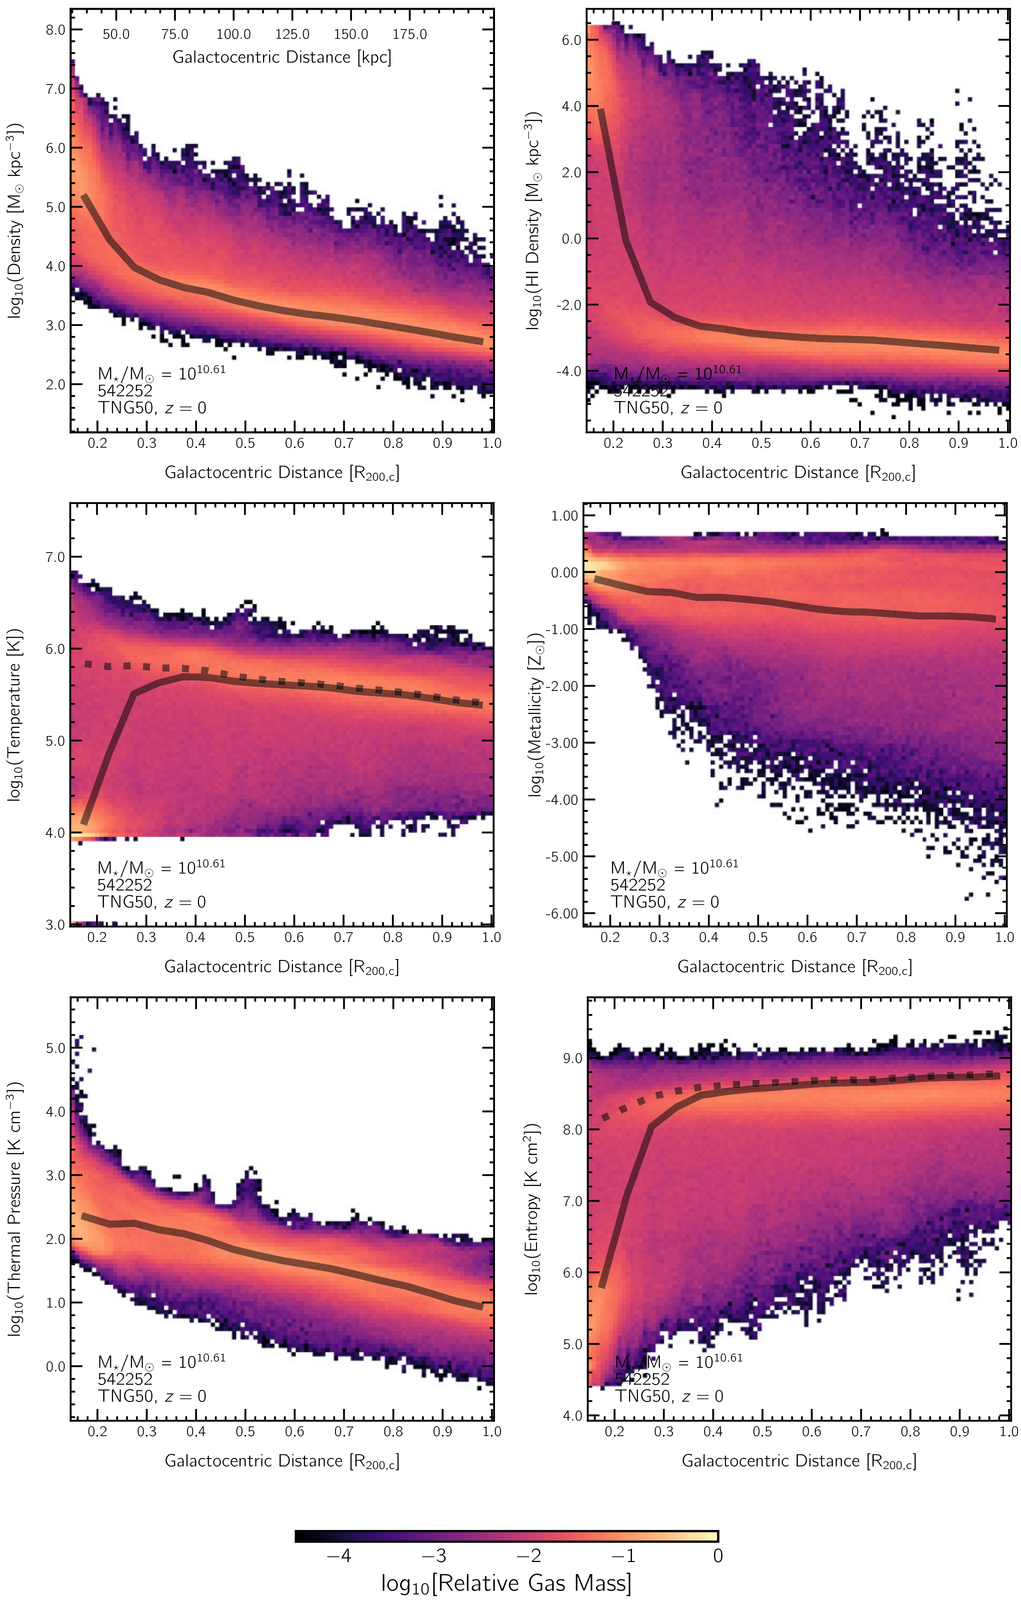

Figure 7. Two-dimensional radial profiles of six physical gas properties for a single TNG50 MW-like galaxy (subhalo 542252), showing the same quantities as Fig. 6. The colours correspond to the relative distribution of gas mass, normalized to the maximum value ($=M_{\rm pixel}/M_{\rm pixel,max}$). In each panel, the black curves show the median profiles of the galaxy's CGM, i.e. the spherically-averaged gas quantity as in Fig. 6, thin curves (there, one curve per galaxy). Satellites are excised. Beyond spherically-averaged radial profiles, we see that gas properties in the CGM can have significant diversity at fixed distance, depending on where they are within the halo, i.e. their angular coordinates. This is particularly true in the inner halo, where distinct (multimodal) gas populations are evident.

is considered. These profiles diverge (i.e. solid versus dotted) in the inner halo as a result of the dominance of cold gas, but asymptotically converge in the outer halo.[7]

3.3 Ranges of variation of TNG50 CGM properties

To quantify the diversity of CGM properties around each galaxy of the TNG50 MW-like sample, Fig. 8 shows histograms of the same six important properties of gas in the CGM: density, H I density, temperature, metallicity, thermal pressure, and entropy. In each panel, the thin curves correspond to individual galaxies, which are coloured by the stellar mass of the galaxy. The four thicker curves show the binned behaviour of haloes split into four quartiles of stellar mass, as indicated by their colours.

In the distribution of CGM gas density (Fig. 8, top-left-hand panel) a dominant peak occurs at $10^3\,M_\odot\,kpc^{-3}$ for most galaxies. A few haloes also show a prominent peak at $10^{6.5}\,M_\odot\,kpc^{-3}$, which however is no longer visible in the binned curves. The first peak corresponds to a density of $\sim 10^{-4.5}\,cm^{-3}$, and is characteristic of densities in the outer halo ($\gtrsim R_{200c}$), while the second peak ($10^{6.5} M_\odot$ $kpc^{-3} \gtrsim 0.1\,cm^{-3}$) traces star-forming gas, a negligible component of the CGM mass budget. No clear monotonic trend with stellar mass is visible in the binned curves. Similar features are apparent with H I density (top-right-hand panel), but in this case, two peaks are observed in the binned curves as well. While the peak at high densities presumably corresponds to largely neutral gas, the peak at lower values corresponds to largely ionized gas.

The distributions of gas temperature (Fig. 8, centre-left-hand panel) show a secondary peak at 10^4 K for all haloes, and a main peak near the virial temperature for haloes of this mass ($\sim 10^6$ K). Gas colder than 10^4 K is not realized in the simulation (see Section 2.1), hence the abrupt cut-off at low temperatures. On the other hand, we place star-forming gas at $\sim 10^3$ K for visualization, separating the different mass bins horizontally. Notably, the median value for the highest stellar mass quartile is zero, and in general, star forming gas is negligible in comparison to the mass of the rest of the CGM. Integrating these distributions, we find that the mass of cold gas is typically about half an order of magnitude smaller than hot gas mass, consistent with Figs 4 and 5. The binned temperature distributions evolve monotonically with mass: haloes with lower stellar mass have lower virial temperatures. The shape of these distributions is similar at cooler temperatures (irrespective of stellar mass), since this traces the physics of the radiative cooling function.

Observations of the MW halo report a temperature component around $\sim 10^{6.3}$ K (Bluem et al. 2022), which is roughly consistent with the second peak observed in the TNG50 temperature distributions. In addition, observational studies infer gas phases as hot as $\sim 10^{6.5}$ K (southern galactic sky Kaaret et al. 2020) and $\sim 10^{6.9}$ K (northern galactic sky Bluem et al. 2022) in the MW halo, albeit the abundance of these phases is uncertain. The temperature of gas within the eROSITA bubble(s) (see Introduction) has been inferred to be about $10^{6.5}$ K with Suzaku (Kataoka et al. 2013) and $10^{6.6-6.7}$ K from the modelling of O VII and O VIII emission lines (Miller et al. 2016). Gas at such high temperatures or even higher exists in a large fraction of TNG50 MW-like galaxies, and is due, in TNG, to galactic-centre bubble-like heating events from the SMBH (Pillepich et al. 2021).

[7]Recall that satellites are excluded from our analysis, which would otherwise be visible as overdensity spikes at the radii where satellites are present.

Distributions of metallicity (Fig. 8, centre-right-hand panel) are distinctly unimodal. They exhibit a peak slightly below solar metallicity ($\sim 10^{-0.4}\,Z_\odot$), with the amount of gas mass dropping sharply on either side: on average, the fractional mass of supersolar gas is $\sim$6 per cent, while low-metallicity gas with $<0.1 Z_\odot$ accounts for $\sim$11 per cent of the mass budget. No clear monotonic trends are predicted by TNG50 as a function of galaxy stellar mass. While metallicity estimates of gas in the MW halo are largely uncertain, Ponti et al. (2022) report a subsolar value of $\sim$0.05–0.1 $Z_\odot$ along a given line of sight. On the other end of the metallicity spectrum, observations of high-velocity clouds reveal the possible existence of highly supersolar clouds with metallicities of $\sim 1.65\,Z_\odot$ (Zech et al. 2008) and $\sim 2.08\,Z_\odot$ (Yao, Shull & Danforth 2011), hinting towards a wide range of metallicites in the MW halo, similar to that of TNG50.

Thermal pressure distributions (Fig. 8, lower left-hand panel) are also similarly unimodal, and peak at $\sim 10^2$ K cm^{-3}. Binned curves show no clear dependence across the four stellar mass quartile bins. We remind the reader here that the pressure for star forming gas is derived from the effective equation of state.

Given that entropy is a linear function of temperature, the corresponding distributions (Fig. 8, lower left-hand panel, lower right-hand panel) behave similarly to those of temperature, as expected. There is a subtle peak at low entropy ($\sim 10^5$ K cm^2), followed by a dominant peak at $\sim 10^9$ K cm^2. On average, only $\sim$20 per cent of gas has an entropy lower than $\sim 10^7$ K cm^2. As with the temperature distributions, the binned curves of haloes with larger stellar mass peak at larger values of entropy.

It is somewhat surprising that multimodal distributions are not seen in the thermal pressure distributions, given the presence of distinct density and temperature components in the CGM gas. This suggests that these different phases are in fact in pressure equilibrium. Similarly, the lack of multimodality in the metallicity distributions is a bit unexpected, since the centre-right-hand panel of Fig. 6 suggests incomplete mixing of metals.

3.3.1 Cross-correlations among CGM phases

To explore the relationship between properties of CGM gas across different components, Fig. 9 shows phase diagrams (i.e. two-dimensional distributions of temperature versus density) for two particular galaxies. We select these two galaxies to be in two different states with respect to their SMBH feedback modes. The first (top row) is in the low-accretion state, and therefore in the kinetic feedback mode of the TNG model: this galaxy exhibits multiple X-ray bubbles in its CGM (as indicated by Pillepich et al. 2021, top-right-hand panel of their figs. 2 and 3). The second galaxy (bottom row) is in the high-accretion state, and so in the thermal or ‘quasar’ feedback mode (Weinberger et al. 2017). Broadly speaking, the former case has had ongoing and recent effective SMBH feedback, while the second case is still dominated, in terms of feedback energy or feedback effects, by the supernovae explosions of stars. The former has a relatively high-virial temperature, $T_{vir} \sim 10^6$ K, while the latter is slightly cooler at $T_{vir} \sim 10^{5.6}$ K, a difference which primarily arises due to the difference in halo virial masses of the two systems.

In the left-hand panels of Fig. 9, we colour the phase diagrams by gas galactocentric distance normalized by halo virial radius. We see that dense gas, near or past the star-formation threshold, is preferentially located close to the centre, as expected. A diagonal gradient is evident, such that moving diagonally from the upper right to towards the lower left of the phase diagram, i.e. towards low density and low temperature, corresponds to moving outwards in

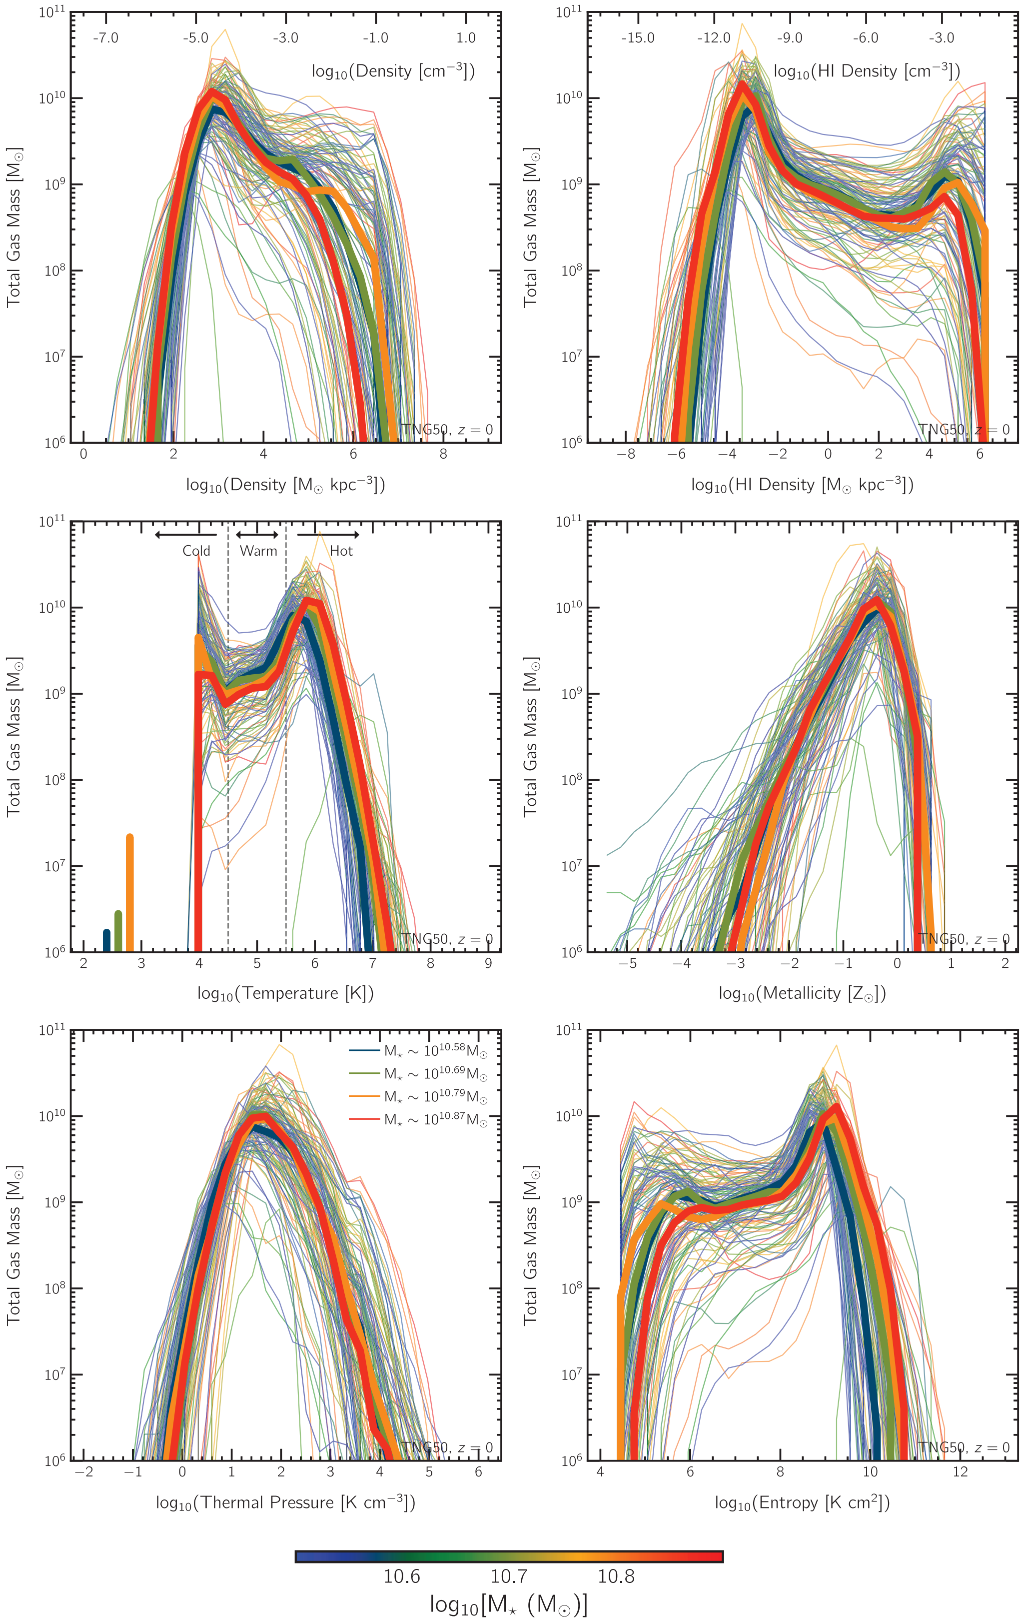

Figure 8. Distributions of the physical properties of gas in the CGM of TNG50 MW-like galaxies at $z = 0$ (in each case), histograms employ 30 bins, and extend between the visible tick marks. Each (thin) curve corresponds to a single galaxy/halo, colour showing its stellar mass, and accounts, as per definition, only of the gas between 0.15 and $1.00 \times R_{\mathrm{vir}}$. The thicker curves show medians across haloes binned into four galaxy stellar mass quartiles. While the metallicity and thermal pressure distributions are unimodal, gas density, H I density, temperature, and entropy are multimodal suggesting that the CGM hosts multiphase gas that forms distinct populations in density, temperature, and entropy. Only temperature and entropy show peaks with strong monotonic trends as a function of stellar mass, with more more massive galaxies hosting gas at higher temperatures and higher entropy in their CGM.

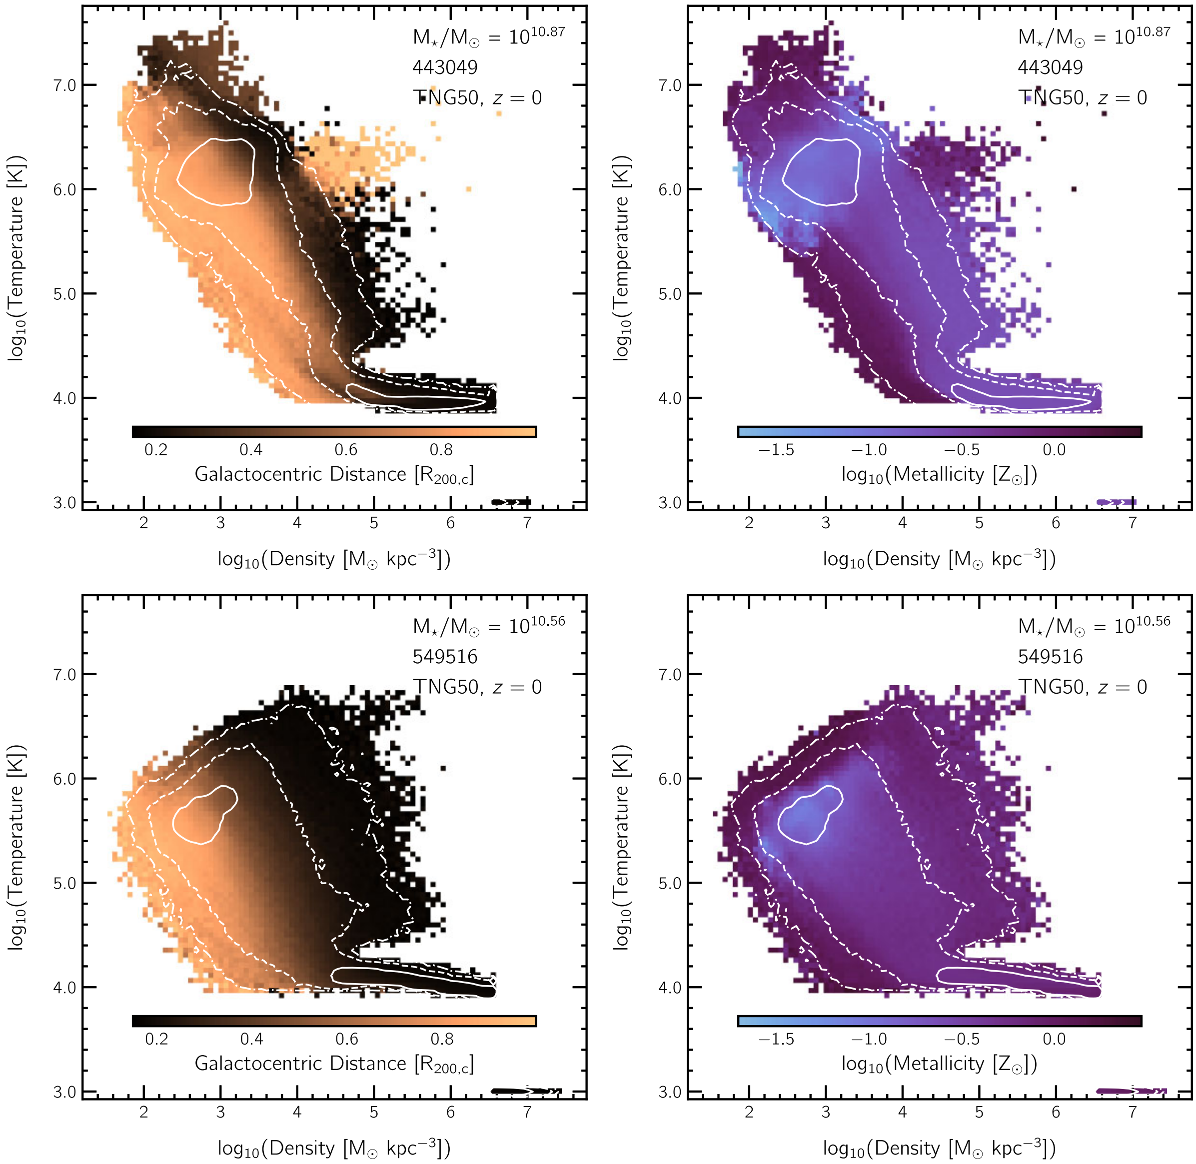

Figure 9. Phase diagrams of temperature versus density for gas in the CGM of two particular TNG50 MW-like galaxies. We select these galaxies as characteristic of two rather different subsets of our sample: the SMBH at $z = 0$ is releasing energy in the low-accretion state kinetic mode (top row, subhalo ID 443049, which exhibits X-ray bubbles in its CGM: Pillepich et al. 2021), or in the high-accretion state thermal mode (bottom row, subhalo ID 549516). Panels on the left show phase diagrams coloured by galactocentric distance, while those on the right are coloured by metallicity. In all cases, contours show where gas mass is actually located in this plane ([solid, dashed, dot–dashed] = [10, 1, 0.1 per cent]), outlining those pixels which contain the respective mass fractions, relative to the maximum. Cold, dense gas is preferentially located in the inner regions of the halo, as opposed to the diffuse, warm-hot phase that dominates at larger distances. The kinetic mode channel of SBMH feedback can heat gas to supervirial temperatures ($\gtrsim 10^7$ K), generating metal-rich components which can span the extent of the halo.

radius. However, in the kinetic AGN feedback example (top), we also find dense (hot) gas at large galactocentric distances, driven away by the intense activity of the SMBH in the kinetic feedback mode. No such feature is present in the halo in the thermal mode AGN feedback.

In the right-hand panels of Fig. 9, we instead colour the gas phases by metallicity. Gas with average metallicities between $\sim 10^{-1}$–$10^{-0.5}$ Z$_\odot$ is predominantly at the virial temperature, and at low densities of $\sim 10^3$ M$_\odot$ kpc^{-3}. At these densities, gas that is hotter than the virial temperature has higher metallicity, which is again presumably linked to feedback processes and outflows from the galaxy. We again note the presence of a gradient that extends diagonally down and to the right from the dark cloud towards the star-forming gas. This likely reflects a cooling channel, i.e. the path that gas takes through the phase diagram as it cools. At the high-density end, we note differences between the two galaxies: in the kinetic SMBH feedback example (top), we see that regions of high-density/high-temperature have relatively higher metallicities than the rest of the high-density regions, a signature of enriched outflows driven by the SMBH-driven outflows (see also Pillepich et al. 2021).

To visualize the relative abundance of gas mass across these phase diagrams, we also draw three white contours on each panel: the solid, dashed, and dot–dashed curves trace those pixels that have [10, 1, 0.1 per cent] mass fractions with respect to the maximum

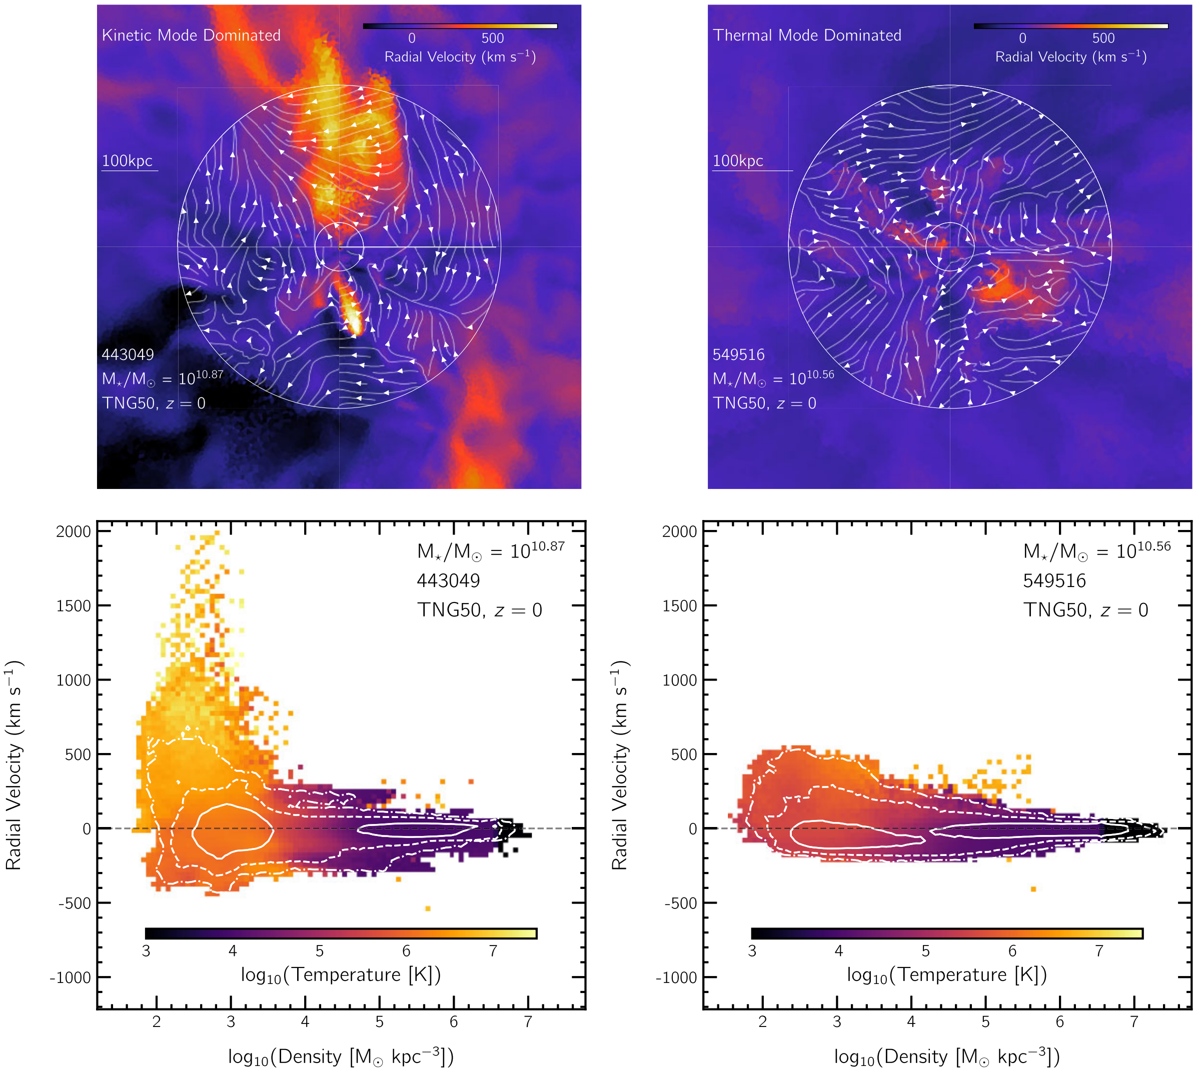

Figure 10. Gas kinematics in the CGM and their relation to SMBH feedback from the galaxy's centre according to TNG50 at $z = 0$. The top panels visualize gas radial velocity (negative denoting inflow and positive denoting outflow), in the frame of reference of the galaxy. We orient the galaxies edge-on, and take an infinitesimally thin slice through the mid-plane of the halo. The left and right columns show two different galaxies, the same as previously shown for the phase diagram analysis. Namely, the TNG50 MW-like galaxy on the left is an example dominated by SMBH feedback in the kinetic mode (low-accretion rate; subhalo ID 443049), whereas the TNG50 MW-like galaxy on the right is experiencing only SMBH feedback from the thermal mode (high-accretion rate; subhalo ID 549516). The bottom row shows phase diagrams of radial velocity versus density, coloured by temperature. Contours show where gas mass is actually located in this plane ([solid, dashed, dot–dashed] = [10, 1, 0.1 per cent]), outlining those pixels which contain the respective mass fractions, relative to the maximum. The galaxy on the right has only modest outflows up to $\sim$500 km s^{-1} with no strong directionality. In contrast, the galaxy on the left has outflows up to $\sim$2000 km s^{-1}, which extend to large distances and are rather asymmetric, and which are dominated by hot $\sim$10^7 K gas.

value in each diagram, respectively. For both galaxies, we note that the majority of the mass is in either (i) dense, cold gas with $T \sim 10^4$ K, or (ii) in diffuse gas close to the virial temperature, which is consistent with the bimodal structure of the one-dimensional temperature distributions examined earlier.

3.4 Kinematics and further connections to SMBH feedback

The preceding analyses reveal several examples of relationships predicted by the TNG model between the properties of CGM gas and properties of the central galaxy. Previous works have shown that the TNG model of kinetic SMBH feedback effectively drives high-velocity outflows in galaxies (Weinberger et al. 2017; Nelson et al. 2019a; Pillepich et al. 2021). Here, we therefore further explore the connection of SMBH feedback with the kinematics of CGM gas in the case of TNG50 MW-like galaxies, and quantify the complexity of the velocity fields in the CGM of MW-like galaxies. Specifically, we distinguish galaxies based on their modes of SMBH feedback in the TNG model: high-accretion thermal mode feedback versus low-accretion kinetic mode feedback.

In Fig. 10, we first provide a visualization to motivate how the mode of SMBH feedback affects gas kinematics, at least in TNG50. The top panels show slices with the galaxies having been rotated to an edge-on orientation of radial velocities for two galaxies. On the left, we choose a prototypical TNG50 MW-like galaxy, where the total (i.e. cumulative in time) energy injection by SMBH feedback has been dominated by the kinetic mode, and the SMBH is accreting gas at a sufficiently low-Eddington ratio at the present day

($z = 0$) that it is also in the kinetic mode. On the right, we choose a representative TNG50 MW-like galaxy, in which the thermal mode has been the dominant channel for SMBH feedback energy injection.

Two main differences are immediately apparent in Fig. 10: the (outflowing) radial velocities are much larger in the first case (left column), and this outflowing gas is preferentially located perpendicular to the disc. In the thermal mode dominated galaxy (right column), SMBH feedback does not produce within the TNG implementation of thermal energy injection, such a high-velocity outflow, no one with such a strong spatial dependence. We note, however, that other implementations of thermal energy injection from SMBHs, different from the TNG model, may be able to drive fast outflows: see the discussion in Pillepich et al. (2021).

As mentioned in Section 2.1, energy injection through both these modes is isotropic at the injection scale. The directionality of the outflows in the kinetic mode instead stems from the interaction of the outflows with gas in the disc and/or CGM (Nelson et al. 2019b; Péroux et al. 2020; Pillepich et al. 2021; Truong et al. 2021b). As discussed in previous works, the kinetic mode of SMBH feedback in the TNG model is an efficient quenching mechanism due to a combination of strong ejective feedback, i.e. outflows together with a preventative heating of the CGM (Nelson et al. 2018a; Terrazas et al. 2020; Zinger et al. 2020). The galaxy shown in the left-hand panel is indeed in the process of quenching and departing from the main sequence, thereby moving into the green valley.

The bottom panels of Fig. 10 show phase diagrams of radial velocity versus density, coloured by temperature, for the same two galaxies. We see that radial velocities are larger in the galaxy undergoing kinetic mode SMBH feedback (left), where outflows up to $\sim$2000 km s^{-1} are evident, although such high speeds are attained only by largely low-density gas ($\lesssim 10^4\,\mathrm{M_\odot\,kpc^{-3}}$). This is in contrast to the second galaxy undergoing thermal mode SMBH feedback, where supernovae and stellar feedback in general are primarily responsible for driving galactic-scale outflows (right), where speeds reach only $\sim$500 km s^{-1}. In addition, gas temperatures are higher in the first case ($\gtrsim 10^7$ K), and this hot gas component outflows the fastest.

In both cases, we note that the majority of cold dense gas on average is neither strongly outflowing nor inflowing, i.e. this is the material which makes up the rotationally-supported gaseous disc centred at zero radial velocity. However, it is believed, both in observations (Di Teodoro et al. 2020) and simulations (Schneider et al. 2020), non-negligible amounts of cold gas exist in outflows, although the mechanism through which cold gas is accelerated, or produced in outflows remains an open question.

In Fig. 11, we quantify halo gas kinematics and inflow/outflow behaviour across our full sample of 132 TNG50 MW-like galaxies. The top panel shows distributions of radial velocity of CGM gas. Each thin curve corresponds to the CGM of an individual galaxy, coloured by mass of the SMBH. The solid thick curves show median values for four quartiles, with the galaxy sample split based on SMBH mass. The dashed thick curves also show median lines for four quartiles, but with the sample split based on the total time-integrated (i.e. cumulative in time) kinetic mode SMBH energy injection of the SMBH.

Most of the distributions are skewed towards positive radial velocities (outflowing gas), although the radial velocities of the bulk of all gas mass are close to zero, i.e. most gas is neither strongly inflowing nor outflowing. The solid median lines show that the galaxies with the least massive SMBHs have smaller high-velocity tails, in comparison to those with more massive SMBHs. However, this trend is not monotonic, i.e. the quartile corresponding to the most massive SMBHs does not have the highest velocity outflows. This arises because the connection between SMBH mass and energy injection in a given mode is not one-to-one. The thick dashed lines make this clear: they show a steady monotonic trend where galaxies whose SMBHs have injected progressively more total energy in the kinetic mode have the highest outflow velocity tails. Although not shown here, such a trend is not seen when galaxies are split into quartiles based on their total energy injection in the thermal quasar mode.

In order to also consider spatial information, the bottom-left-hand panel of Fig. 11 shows three distinct mass flow rates as a function of distance. In solid, dashed, and dash–dot lines, we show the mass outflow-, mass inflow-, and net mass flow-rates, respectively, averaged across subsets of TNG50 MW-like galaxies. As above, we split the sample into quartiles based on SMBH mass (from least to most massive: purple, red, orange, yellow). We find that MW-like galaxies with less massive SMBHs have greater outflow rates at smaller distances (i.e. at the inner boundary of the CGM), with values decreasing towards larger distances. The trend reverses for galaxies with more massive SMBHs, where outflow rates are higher at larger distances. We interpret this as the effect of the transition from thermal to kinetic SMBH feedback mode: in TNG, less massive SMBHs are predominantly in the thermal mode, while more massive SMBHs are dominated by kinetic mode feedback.

Similarly, inflow rates (dashed lines) are larger for galaxies with less massive SMBHs at small galactocentric distances, with the inflow rate reducing towards the virial radius. The increased inflow rate close to the disc for galaxies in the lowest quartile indicates a cooling flow regime, wherein cooling turns more efficient at high densities. For galaxies with more massive SMBHs, however, no strong trends with respect to distance are observed, with the gas inflow rate remaining rather flat through most of the CGM. Finally, the net flow rates (dot-dashed lines) are low in order of a few solar masses per year, and have little dependence on galactocentric distance. Only galaxies with the least massive SMBHs have a net (median) inflow rate at all distances. The other three bins have positive net mass flow rates at almost all distances; galaxies with more massive SMBHs have slighly larger net outflow rates throughout the CGM. Overall, the majority of circumgalactic gas in MW-like haloes is neither strongly outflowing nor inflowing.

The bottom-right-hand panel of Fig. 11 shows mass flow rates as a function of absolute radial velocity, or faster (i.e. cumulative in velocity). Once again in solid, dashed, and dash–dotted lines, we show the mass outflow-, mass inflow-, and net mass flow-rates, respectively. However, in this case, we show median lines for the entire sample (i.e. without splitting into quartiles). Instead, we show the outflow rate at three radii: $\{0.15, 0.5, 1.0\} \times R_{200c}$ – black, red, and orange, respectively. Similar to the top panel, we observe that most gas is moving slowly, with both inflow and outflow rates decaying quickly towards high velocities, at all distances. For instance, at a distance of $0.15R_{200c}$, the inflow rate is $\sim 10\,\mathrm{M_\odot\,yr^{-1}}$ when all inflowing gas is considered (i.e. all gas with $v_r < 0$ km s^{-1}), dropping quickly to $\sim 1\,\mathrm{M_\odot\,yr^{-1}}$, when only gas inflowing faster than 200 km s^{-1} is considered. Similar numbers hold for outflows as well.

For outflows, we see that the rate of slowly outflowing gas is roughly independent of distance. At high speeds ($\gtrsim$600 km s^{-1}), outflow rates begin to depend on distance: such outflows are faster at smaller distances. For instance, the outflow rate for gas with $v_{\rm rad} > 600$ km s^{-1} at a distance of $0.15R_{200c}$ is $\sim 1\,\mathrm{M_\odot\,yr^{-1}}$, but is only $\sim 0.3\,\mathrm{M_\odot}\,yr^{-1}$ at the virial radius. At even higher speeds ($\gtrsim$1200 km s^{-1}), the outflow rate at large distances is negligible.

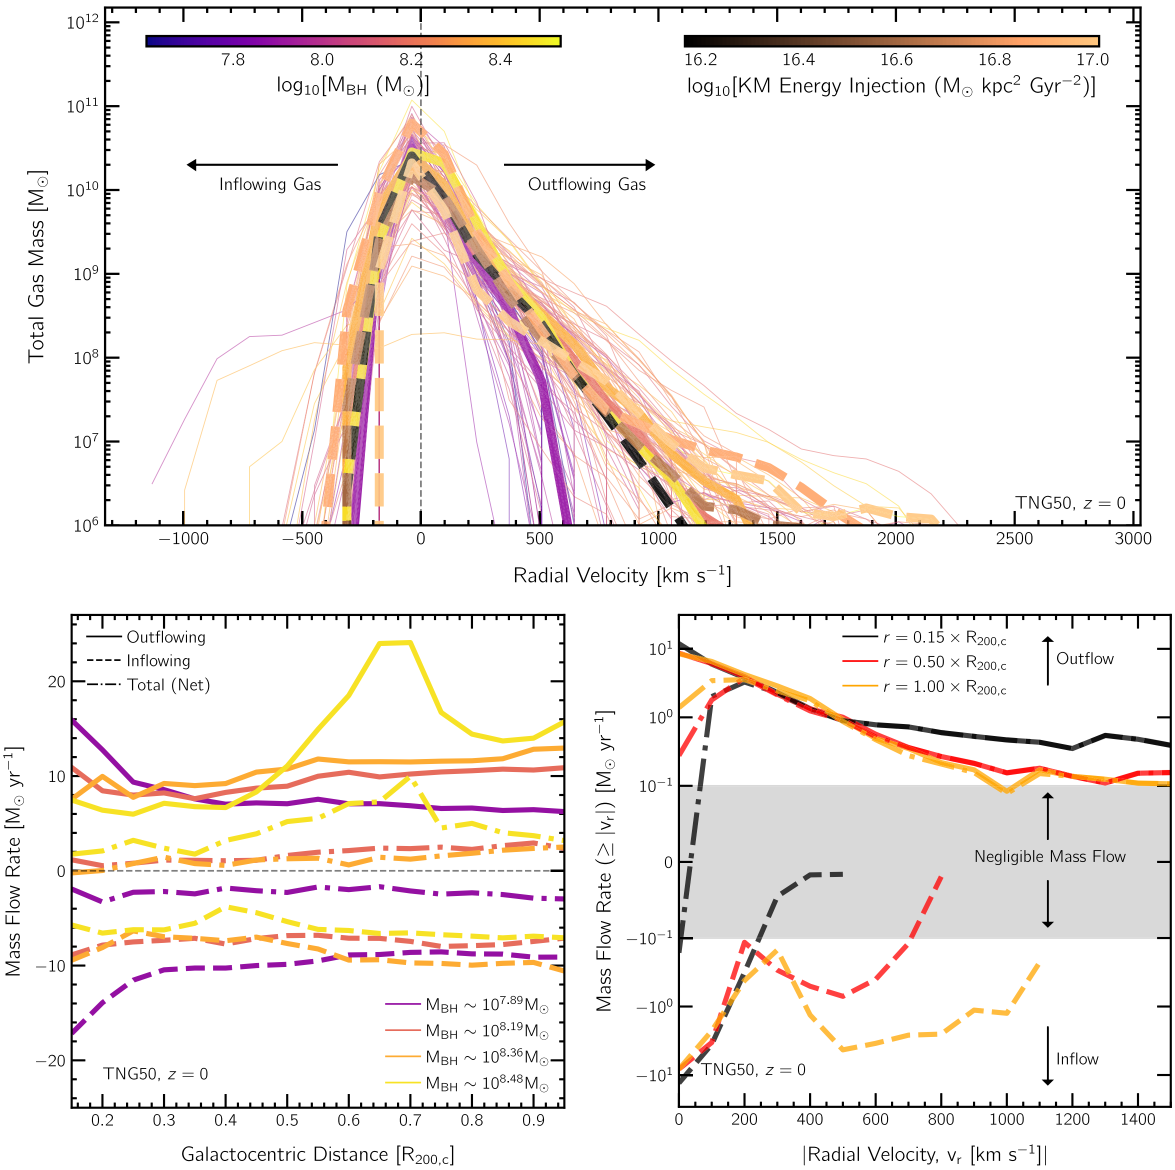

Figure 11. The kinematics and velocity structure of gas flows, and the balance between inflows and outflows in the CGM of $z = 0$ MW-like galaxies in TNG50. In the top panel, we show distributions of gas radial velocity (histograms employ 30 bins, and extend (exactly) between -1100 and $3000\,\mathrm{km\,s^{-1}}$). Thin curves correspond to individual galaxies, coloured by SMBH mass (top left colourbar), while the two sets of thick curves are binned into four quartiles each: solid in bins of SMBH mass (top left colourbar), and dashed in bins of total time-integrated kinetic mode (KM) energy injection from the SMBH (top right colourbar). We see that galaxies with larger total kinetic mode energy injection have tails of higher-velocity outflows, representing the integrated impact of multiple episodes of kinetic-mode energy injection. The bottom-left-hand panel shows the (median) mass outflow rate as a function of galactocentric distance, with galaxies split into four quartiles based on their SMBH mass (four colours). We use solid, dashed, and dot–dashed lines to represent outflowing, inflowing, and all gas, respectively. We find that galaxies with the least massive SMBHs are dominated by inflows at all distances, while the other three bins are dominated by outflows at almost all distances. The bottom-right-hand panel shows mass flow rates of gas moving at a given absolute radial velocity, or faster (i.e. cumulative in velocity). Once again, we use solid, dashed, and dot–dashed lines to represent outflowing, inflowing, and all gas, respectively, averaged across all 132 MW-like galaxies. The black, red, and orange lines refer to outflow rates across shells at radii 0.15, 0.5, and 1.0 R_{200c}, respectively. We see that gas outflows at larger velocities at smaller galactocentric distances, while the opposite is true for inflows.

There is little gas moving at such large velocities so far from the galaxy itself.

Similarly, the rate of gas inflow at small velocities is roughly independent of distance, possibly indicating the existence of significant subcentrifugal rotation throughout the CGM, resulting in deviations from local hydrostatic equilibrium (Oppenheimer 2018). However, inflow rates at large velocities ($\gtrsim 250\,\mathrm{km\,s^{-1}}$) decrease towards smaller radii; only gas at large distances from the centre has large inflow speeds. We interpret this as the interaction of accreting gas with gas in the halo, which slows down due to drag and/or interactions with outflows. Finally, the median net flow rate is almost always positive (i.e. gas is outflowing), except for slowly moving material with $\lesssim 100\,\mathrm{km\,s^{-1}}$, where the net flow rate drops essentially to zero. In conjunction with the fact that slowly moving material dominates the mass budget of the CGM (top panel), we conclude that a majority of gas is in quasi-static equilibrium.

This quantification of the kinematics of CGM gas makes it clear that the gaseous halo of TNG50 MW-like galaxies hosts not only a complex multiphase structure, but also a complex dynamical structure that is closely linked to energetic feedback events from the central galaxy and its SMBH.

3.5 Magnetic fields in the CGM of TNG50 MW-like galaxies

A modelling element that sets apart the TNG simulations from many other similar projects is the inclusion of MHD. Therefore, we close this analysis by quantifying the existence and properties of magnetic fields in the CGM of MW-like galaxies. In the TNG simulations, a uniform (weak) primordial magnetic field (10^{-14} comoving Gauss) is seeded at the start of the simulations ($z = 127$), which is then amplified due to the process of structure formation and feedback (Marinacci et al. 2018). By low redshift, magnetic fields in haloes obtain significantly higher strengths, which could influence the evolution and dynamics of halo gas (Pakmor et al. 2020), possibly by playing a role in suppressing fluid instabilities (e.g. Berlok & Pfrommer 2019; Sparre, Pfrommer & Ehlert 2020). Here we focus on two related properties: magnetic field strength $|B|$ and the pressure ratio β^{-1} of magnetic to thermal pressure components in the gas.

Fig. 12 shows our key results for the magnetic fields in the haloes of MW-like galaxies from TNG50 at $z = 0$. In the top and middle panels, we visualize the magnetic field properties for a single galaxy, which is oriented edge-on. This is the same system shown in eight different visualized quadrants in Figs 2 and 3. In the top, we show slices of magnetic field strength (left) and pressure ratio (right): the magnetic field strength is generally higher closer to the disc/inner region of the CGM, as a result of which the magnetic pressure is also the highest in the same regions. However, the structure at fixed distance is complex, in terms of variation in magnetic field strength.

The middle panels of Fig. 12 show phase diagrams for the same galaxy as the top panels. On the left, we colour by the magnetic field strength, while we colour by the pressure ratio on the right.[8] We see that cold dense gas has the strongest magnetic fields, while diffused hot gas has the weakest. A diagonal gradient is present in the bottom left-hand panel: as one moves diagonally from high-temperature and high-density gas (top right) to lower-temperature and lower-density gas (lower left), the magnetic field strength decreases, tracing the increase in galactocentric distance. Although weak, the reverse trend is present in the bottom right-hand panel: the pressure ratio increases gradually diagonally towards lower-temperature and lower-density gas (lower left), since this is the direction in which thermal pressure tends to decrease.

Finally, the bottom panels of Fig. 12 give the spherically-averaged radial profiles of the same two magnetic-field quantities as a function of galactocentric distance, for all 132 TNG50 MW-like galaxies: thin curves, coloured by galaxy stellar mass as in the previous figures. These are mass-weighted average magnetic field properties at the given galactocentric distance. These plots quantify the visual impressions above: the median magnetic field strength is roughly one and a half orders of magnitude greater at the inner boundary of the CGM as compared to the outer boundary, dropping from $\sim 1\,\mu$G at $0.2R_{200c}$ to $\sim 0.1\,\mu$G at $0.5R_{200c}$, to $\sim 0.03\,\mu$G at the virial radius (in the median). In addition, the magnetic field strength, on average, is larger in the CGM of galaxies with lower stellar masses (at all distances).

We compare ther simulation results for the magnetic field profile around a MW-like galaxy from the Auriga simulations (Au6; van de Voort et al. 2021). While the median value of TNG50 MW-like galaxies is similar to the Au6 profile (dark grey dashed line) close to the inner boundary of the CGM, differences begin to arise at $\sim 0.25R_{200c}$, beyond which the TNG50 median curve is moderately lower, by up to $\sim$0.3 dex. However, the Au6 profile is within the galaxy-to-galaxy variation predicted by TNG50 for MW-like galaxies. While Auriga is also run with AREPO, and utilizes the same MHD solver as TNG, along with the same primordial seed strength, differences in the galaxy physics model, i.e. feedback physics could contribute to the offset of the Au6 profile with respect to the TNG50 median. This is, further, a comparison of a single halo versus the TNG50 sample, and halo to halo diversity exists. In addition, Au6 is simulated at a slightly higher resolution ($m_{\rm gas} \sim 5.4 \times 10^4\,{\rm M}_\odot$, plus added refinement to maintain a spatial resolution of $\sim$1 kpc in the CGM) with respect to TNG50 ($m_{\rm gas} \sim 8.4 \times 10^4\,{\rm M}_\odot$). Slightly stronger magnetic field strengths are therefore expected, given that the late time strengths increase slightly with better numerical resolution Marinacci et al. (2015).

We also show results (Ponnada et al. 2022) from one of the simulations based on the FIRE-2 model, where magnetic fields were also included (MHD+). We note here that the MHD solver utilized in the FIRE-2 simulations is quite different from the one used in TNG, and so is the mass resolution of gas ($m_{\rm gas} \sim 7 \times 10^3\,{\rm M}_\odot$). The magnetic field strength in the FIRE-2 MHD+ m12i run (dark grey dot–dashed line) is offset by ~ -1.5 dex with respect to the TNG50 median at all distances. We also include the curve corresponding to the FIRE-2 CR+ m12i run at $z = 0$, in which cosmic rays were included in addition to magnetic fields (effective diffusion coefficient: $\kappa = 3 \times 10^{29}$). This case (dark grey dotted line) is very similar to the MHD+ run, expect at large galactocentric distances ($\gtrsim 0.65 \times {\rm R}_{200c}$), where the CR+ curve is slightly higher with respect to MHD+.

Owing to their low-strength magnetic fields are difficult to measure observationally within the MW halo. While Strong, Moskalenko & Reimer (2000) report a value of $\sim 6\,\mu$G around the solar neighbourhood ($\sim$8 kpc), and Crutcher et al. (2010) estimate a slightly higher value of $\sim$7.6–11.2 μG close to the galactic disc ($\sim$4 kpc), magnetic field strengths remain poorly constrained in the MW halo. The above two constraints are shown through the black dashed line and grey shaded region, respectively, in the centre-left-hand panel. While we do not show radial profiles of TNG50 at these galactocentric distances, it is visible through the slice visualization (top left-hand panel) that the typical magnetic field strength in the central regions of TNG50 MW-likes is of order $\sim 10\,\mu$G, broadly consistent with available observational constraints.

As a result of lower magnetic field strengths in the CGM, magnetic pressure is generally subdominant in the CGM in comparison to thermal pressure (Nelson et al. 2020), as can be seen in the radial profile of $P_{\rm B}/P_{\rm th}$ (Fig. 12, lower right-hand panel), although there are regions where magnetic pressure dominates, as is visible in the top-right-hand panel. The median TNG50 ratio declines from $\sim$1 at $0.2R_{200c}$ to $\sim$0.01 at the virial radius. An analysis of this magnetic to thermal pressure ratio has also been made with the same set of FIRE-2 simulations discussed above (Hopkins et al. 2020). The pressure ratio in the FIRE-2 MHD+ m12i run at $z = 0$ run (dark grey dot–dashed line) is significantly lower than the median value from TNG50 throughout the CGM: by 2–3 dex, depending on distance. In

[8] As before, contours show where gas mass is actually located in this plane ([solid, dashed, dot–dashed] = [10, 1, 0.1 per cent]), outlining those pixels which contain the respective mass fractions, relative to the maximum.

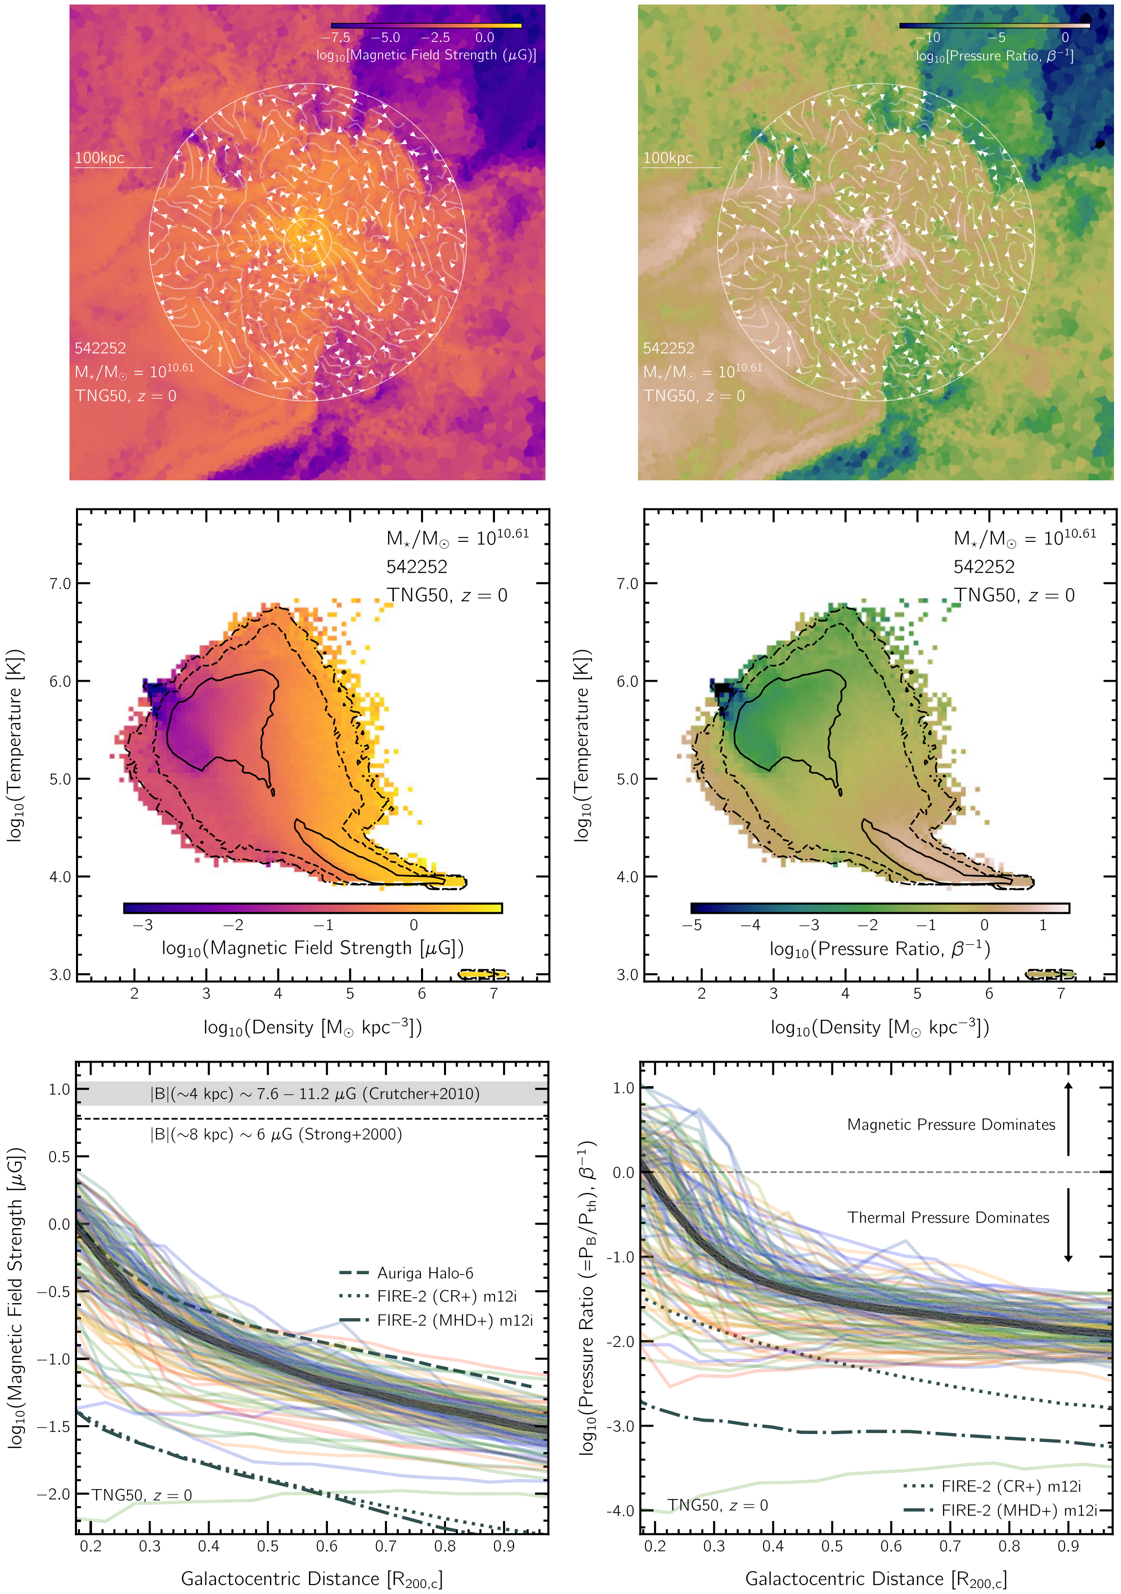

Figure 12. The structure, strength, and importance of magnetic fields in the CGM of MW-like galaxies in TNG50. Top and middle panels: magnetic fields properties of the same MW-like galaxy previously visualized in detail (subhalo ID 542252). In the top, we show slices of magnetic field strength (left) and pressure ratio (right). For these projections, the galaxy is oriented edge-on and streamlines show the x-/z-components of the B-fields, i.e. the local direction of the magnetic fields in the plane of the image. The middle panels show phase diagrams in the temperature-density plane, coloured by the same two B field quantities. Both the magnetic field strength $|B|$ and pressure ratio (β^{-1}) are highest in the inner CGM. Bottom panels: spherically-averaged radial profiles of magnetic field strength and magnetic-to-thermal ratio for all the 132 MW-like galaxies in TNG50. Colour represents stellar mass, as previously.

the case of the FIRE-2 CR+ m12i run (dark grey dotted line), the pressure ratio is closer to the TNG50 result, but still smaller at all distances: a difference of ∼1.2 dex in the inner regions of the halo, and ∼0.4 dex at a distance of $\sim 0.7\mathrm{R}_{200c}$. As in our comparison with Au6 above, the difference between TNG and FIRE-2 is likely due to a combination of differences in sample selection, ISM/feedback models, and numerical techniques.

While recent theoretical studies of the CGM have begun exploring the importance of magnetic fields in the evolution of MW-like galaxies, different simulations currently produce impressively different results. Comparisons with the real MW halo will provide important constraints. The recent observation of magnetic fields in an external $M_\star \sim 10^{10.69}\,\mathrm{M}_\odot$ galaxy at $z \sim 0.36$ by Prochaska et al. (2019) suggests that observational inference of halo magnetic field properties is a promising future direction (Heald et al. 2020).

4 SUMMARY AND CONCLUSIONS

In this paper, we have quantified the CGM at $z = 0$ of MW-like galaxies in TNG50, a cosmological magnetohydrodynamical simulation hosting 132 such galaxies. These have been selected based on their stellar mass ($10^{10.5-10.9}\,\mathrm{M}_\odot$), stellar discyness, and Mpc-scale environment. They are situated within haloes whose M_{200c} is in the range $\sim 10^{11.7-12.5}\,\mathrm{M}_\odot$. Throughout we define the CGM as the region between 0.15 and $1.0R_{200c}$ of each halo, removing satellites.

The key result of our analysis is that, according to TNG50, the physical properties of gas in the CGM – density, temperature, pressure, entropy, H I content, metallicity, and X-ray emission – are diverse both across different MW-like galaxies and across the CGM of the same galaxy, as are the magnetic field properties, and the kinematics of inflows and outflows. The galaxy-to-galaxy diversity is connected, according to the TNG model, to the properties of the galaxy itself and to recent feedback activity, particularly from the central SMBH.

We summarize our specific findings as follows:

(i) The resolved two-dimensional structure of physical gas properties such as density, temperature, and metallicity reveal enormous diversity across the sample of MW-like haloes, as well as within individual haloes, as a function of distance, and due to inhomogeneities on small ∼kpc scales in different regions of the CGM (Figs 1, 2, and 3).

(ii) Integral properties of the CGM such as total gas mass, gas mass in the cold phase, and mass of metals in gas show a large scatter across the sample, but show no significant trend as a function of galaxy stellar mass. Total X-ray luminosity does increase rapidly with stellar mass, with more massive galaxies emitting more. Large galaxy-to-galaxy variation is present, and this scatter is explained by a strong correlation with specific star formation rate (sSFR) of the galaxy, which in turn depends on SMBH feedback, with higher sSFR galaxies hosting most massive and X-ray luminous CGM components (Fig. 4).

(iii) The amount of gas in the CGM varies significantly across the TNG50 MW-like sample: while the median fractional mass of gas in the CGM with respect to $f_\mathrm{b}M_{200c}$ is ∼29 per cent, the 16$^\mathrm{th}$ and 84$^\mathrm{th}$ percentiles are ∼15 to ∼41 per cent, respectively. The median fractions of mass for cold gas, warm gas, hot gas, and metals (all in the CGM) are $\sim 5^{+7}_{-4}$, 6^{+5}_{-4}, 14^{+8}_{-6}, and 17^{+8}_{-10}, respectively. Importantly, the hot gas component ($T > 10^{5.5}$ K) dominates the CGM mass budget of MW-like galaxies (Fig. 5).

(iv) Spherically-averaged radial profiles of temperature and entropy indicate that the (population-wide median) value of these two quantities is smaller in the inner regions of the CGM and increases with distance. However, the spherically-averaged gas density, H I density, metallicity, and thermal pressure decline monotonically outwards (Fig. 6). Two-dimensional radial distributions reveal the existence of multimodal distributions at fixed distance, particularly for temperature and entropy, although the overall homogeneity of thermal pressure indicates that different components are roughly in pressure equilibrium (Fig. 7). According to TNG50, within such a pressure-equilibrated CGM, different phases of gas co-exist, ranging across >4 orders of magnitude in density, temperature, metallicity, and entropy (Fig. 8).

(v) In TNG50, feedback from the central SMBH has an important effect not only on the properties (Fig. 4) but also on the kinematics of gas in the CGM (Fig. 10). Kinetic mode SMBH feedback can produce gas outflows in MW-like galaxies with large velocities ($\gtrsim$500–2000 km s^{-1}) and supervirial temperatures ($>10^{6.5-7}$ K). As a result, galaxies with more massive SMBHs suppress inflow through the CGM, and shift the net flow through the CGM from inflow to outflow. The majority of gas in the CGM is slowly moving ($\lesssim$100 km s^{-1}), tracing a quasi-static halo (Fig. 11).

(vi) Finally, the CGM of TNG50 MW-like galaxies has a rich magnetic field structure: complex on small scales yet coherent across the halo. The magnetic field strength is as high as ∼1 μG in the halo, but with a steep radial profile: largest near the centre and declining with radius. As a result, magnetic pressure dominates over thermal pressure only in the inner regions of the CGM of MW-like galaxies, although there do exist local regions of gas throughout the halo in which magnetic pressure is the dominant component (Fig. 12).

In this work, we have explored several important physical properties of the CGM of MW-like galaxies as predicted by the TNG50 simulation. However, several linked processes and topics have been excluded: for instance, the impact of mergers and satellite galaxies and, in general, the time evolution of the CGM and its state at $z > 0$ for the progenitors of the present-day MW. The origin of cold-phase gas in the CGM, and small-scale cold clouds in particular, remains a topic for future work. While the TNG simulations include magnetic fields and complex turbulent motions, they neglect other non-thermal components, namely cosmic rays, which may play a role in the evolution of the CGM. Future simulations and analyses will explore these questions together with their observable signatures in the gaseous halo of the MW.

ACKNOWLEDGEMENTS

RR and DN acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Group (grant number NE 2441/1-1). AP acknowledges funding by the DFG – Project-ID 138713538 – SFB 881 ('The Milky Way System', subprojects A01 and A06). RR is a Fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). The TNG50 simulation was run with compute time granted by the Gauss Centre for Supercomputing (GCS) under Large-Scale Projects GCS-DWAR on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK), and Nordrhein-Westfalen (MIWF-NRW). This analysis has been carried out on the VERA supercomputer of the Max Planck Institute for

Astronomy (MPIA), operated by the Max Planck Computational Data Facility (MPCDF). This research has made use of NASA's Astrophysics Data System Bibliographic Services. The authors thank the anonymous referee for constructive feedback, and Philip Hopkins and Sam Ponnada for sharing relevant data points.

DATA AVAILABILITY

The TNG simulations, including TNG50, are publicly available and accessible at www.tng-project.org/data (Nelson et al. 2019a). Post-processing catalogues of star formation rates and molecular hydrogen fractions used in this paper are also available on the same website. Data directly related to this publication is available upon reasonable request.

REFERENCES

Anderson M. E., Gaspari M., White S. D. M., Wang W., Dai X., 2015, MNRAS, 449, 3806
Ashley T., Fox A. J., Jenkins E. B., Wakker B. P., Bordoloi R., Lockman F. J., Savage B. D., Karim T., 2020, ApJ, 898, 128
Ashley T., Fox A. J., Cashman F. H., Lockman F. J., Bordoloi R., Jenkins E. B., Wakker B. P., Karim T., 2022, Nat. Astron., 6, 968
Ayromlou M., Nelson D., Yates R. M., Kauffmann G., White S. D. M., 2019, MNRAS, 487, 4313
Berlok T., Pfrommer C., 2019, MNRAS, 489, 3368
Bluem J. et al., 2022, ApJ, 936, 72
Bogdán Á. et al., 2013a, ApJ, 772, 97
Bogdán Á., Forman W. R., Kraft R. P., Jones C., 2013b, ApJ, 772, 98
Bordoloi R. et al., 2017, ApJ, 834, 191
Brüns C. et al., 2005, A&A, 432, 45
Burchett J. N. et al., 2019, ApJ, 877, L20
Butsky I. S., Quinn T. R., 2018, ApJ, 868, 108
Butsky I. S., Fielding D. B., Hayward C. C., Hummels C. B., Quinn T. R., Werk J. K., 2020, ApJ, 903, 77
Byrohl C. et al., 2021, MNRAS, 506, 5129
Cai Z. et al., 2017, ApJ, 839, 131
Chadayammuri U., Bogdan A., Oppenheimer B., Kraft R., Forman W., Jones C., 2022, ApJ, 936, L15
Chen H.-W., Zahedy F. S., Johnson S. D., Pierce R. M., Huang Y.-H., Weiner B. J., Gauthier J.-R., 2018, MNRAS, 479, 2547
Clark S., Bordoloi R., Fox A. J., 2022, MNRAS, 512, 811
Comparat J. et al., 2022, A&A, 666, A156
Corlies L., Peeples M. S., Tumlinson J., O'Shea B. W., Lehner N., Howk J. C., O'Meara J. M., Smith B. D., 2020, ApJ, 896, 125
Crutcher R. M., Wandelt B., Heiles C., Falgarone E., Troland T. H., 2010, ApJ, 725, 466
Damle M. et al., 2022, MNRAS, 512, 3717
Davies J. J., Crain R. A., Oppenheimer B. D., Schaye J., 2020, MNRAS, 491, 4462
de Graaff A., Cai Y.-C., Heymans C., Peacock J. A., 2019, A&A, 624, A48
Di Teodoro E. M., McClure-Griffiths N. M., Lockman F. J., Denbo S. R., Endsley R., Ford H. A., Harrington K., 2018, ApJ, 855, 33
Di Teodoro E. M., McClure-Griffiths N. M., Lockman F. J., Armillotta L., 2020, Nature, 584, 364
Donnari M. et al., 2019, MNRAS, 485, 4817
Dutta A., Sharma P., Nelson D., 2022, MNRAS, 510, 3561
Engler C. et al., 2021, MNRAS, 507, 4211
Esmerian C. J., Kravtsov A. V., Hafen Z., Faucher-Giguère C.-A., Quataert E., Stern J., Kereš D., Wetzel A., 2021, MNRAS, 505, 1841
Faerman Y., Sternberg A., McKee C. F., 2017, ApJ, 835, 52
Faerman Y., Sternberg A., McKee C. F., 2020, ApJ, 893, 82
Fang T., Bullock J., Boylan-Kolchin M., 2013, ApJ, 762, 20
Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416
Ferland G. J. et al., 2013, RMxAA, 49, 137
Fernández X., Joung M. R., Putman M. E., 2012, ApJ, 749, 181
Fielding D., Quataert E., McCourt M., Thompson T. A., 2017, MNRAS, 466, 3810
Fielding D. B. et al., 2020, ApJ, 903, 32
Fox A. J. et al., 2015, ApJ, 799, L7
Genel S. et al., 2014, MNRAS, 445, 175
Gnedin N. Y., Kravtsov A. V., 2011, ApJ, 728, 88
Goulding A. D. et al., 2016, ApJ, 826, 167
Gronke M., Oh S. P., 2020, MNRAS, 492, 1970
Gupta A., Mathur S., Krongold Y., Nicastro F., Galeazzi M., 2012, ApJ, 756, L8
Hafen Z. et al., 2019, MNRAS, 488, 1248
Hani M. H., Ellison S. L., Sparre M., Grand R. J. J., Pakmor R., Gomez F. A., Springel V., 2019, MNRAS, 488, 135
Heald G. et al., 2020, Galaxies, 8, 53
Henley D. B., Shelton R. L., Kwak K., Joung M. R., Mac Low M.-M., 2010, ApJ, 723, 935
Hennawi J. F., Prochaska J. X., Cantalupo S., Arrigoni-Battaia F., 2015, Science, 348, 779
Hodges-Kluck E. J., Miller M. J., Bregman J. N., 2016, ApJ, 822, 21
Hopkins P. F., 2015, MNRAS, 450, 53
Hopkins P. F. et al., 2020, MNRAS, 492, 3465
Huang X., Jiang Y.-f., Davis S. W., 2022, ApJ, 931, 140
Ji S. et al., 2020, MNRAS, 496, 4221
Joung M. R., Bryan G. L., Putman M. E., 2012, ApJ, 745, 148
Kaaret P. et al., 2020, Nat. Astron., 4, 1072
Kataoka J. et al., 2013, ApJ, 779, 57
Kauffmann G., Nelson D., Borthakur S., Heckman T., Hernquist L., Marinacci F., Pakmor R., Pillepich A., 2019, MNRAS, 486, 4686
Kereš D., Hernquist L., 2009, ApJ, 700, L1
Leclercq F. et al., 2017, A&A, 608, A8
Leclercq F. et al., 2020, A&A, 635, A82
Lehner N. et al., 2020, ApJ, 900, 9
Li M., Tonnesen S., 2020, ApJ, 898, 148
Li Z., Hopkins P. F., Squire J., Hummels C., 2020, MNRAS, 492, 1841
Li J.-T., Decourchelle A., Miceli M., Vink J., Bocchino F., 2016, MNRAS, 462, 158
Lim S. H., Barnes D., Vogelsberger M., Mo H. J., Nelson D., Pillepich A., Dolag K., Marinacci F., 2021, MNRAS, 504, 5131
Marinacci F., Vogelsberger M., Mocz P., Pakmor R., 2015, MNRAS, 453, 3999
Marinacci F. et al., 2018, MNRAS, 480, 5113
McCourt M., O'Leary R. M., Madigan A.-M., Quataert E., 2015, MNRAS, 449, 2
McMillan P. J., 2011, MNRAS, 414, 2446
Miller M. J., Bregman J. N., 2015, ApJ, 800, 14
Miller M. J., Hodges-Kluck E. J., Bregman J. N., 2016, ApJ, 818, 112
Naiman J. P. et al., 2018, MNRAS, 477, 1206
Nelson D., Genel S., Vogelsberger M., Springel V., Sijacki D., Torrey P., Hernquist L., 2015, MNRAS, 448, 59
Nelson D. et al., 2018a, MNRAS, 475, 624
Nelson D. et al., 2018b, MNRAS, 477, 450
Nelson D. et al., 2019a, Comput. Astrophys. Cosmol., 6, 2
Nelson D. et al., 2019b, MNRAS, 490, 3234
Nelson D. et al., 2020, MNRAS, 498, 2391
Nicastro F., Senatore F., Krongold Y., Mathur S., Elvis M., 2016, ApJ, 828, L12
Oppenheimer B. D., 2018, MNRAS, 480, 2963
Oppenheimer B. D. et al., 2020, MNRAS, 491, 2939
Pakmor R., Marinacci F., Springel V., 2014, ApJ, 783, L20
Pakmor R. et al., 2020, MNRAS, 498, 3125
Peek J. E. G. et al., 2011, ApJS, 194, 20
Peeples M. S., Werk J. K., Tumlinson J., Oppenheimer B. D., Prochaska J. X., Katz N., Weinberg D. H., 2014, ApJ, 786, 54
Peeples M. S. et al., 2019, ApJ, 873, 129
Péroux C., Nelson D., van de Voort F., Pillepich A., Marinacci F., Vogelsberger M., Hernquist L., 2020, MNRAS, 499, 2462
Pillepich A. et al., 2018a, MNRAS, 473, 4077

Pillepich A. et al., 2018b, MNRAS, 475, 648
Pillepich A. et al., 2019, MNRAS, 490, 3196
Pillepich A., Nelson D., Truong N., Weinberger R., Martin-Navarro I., Springel V., Faber S. M., Hernquist L., 2021, MNRAS, 508, 4667
Planck Collaboration et al., 2016, A&A, 594, A13
Ponnada S. B. et al., 2022, MNRAS, 516, 4417
Ponti G. et al., 2022, preprint (arXiv:2210.03133)
Popping G. et al., 2019, ApJ, 882, 137
Predehl P. et al., 2020, Nature, 588, 227
Prochaska J. X. et al., 2017, ApJ, 837, 169
Prochaska J. X. et al., 2019, Science, 366, 231
Putman M. E., Bland-Hawthorn J., Veilleux S., Gibson B. K., Freeman K. C., Maloney P. R., 2003, ApJ, 597, 948
Putman M. E., Peek J. E. G., Joung M. R., 2012, ARA&A, 50, 491
Qu Z., Bregman J. N., 2019, ApJ, 880, 89
Qu Z., Bregman J. N., Hodges-Kluck E., Li J.-T., Lindley R., 2020, ApJ, 894, 142
Qu Z., Lindley R., Bregman J. N., 2022, ApJ, 924, 86
Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, MNRAS, 430, 2427
Salem M., Besla G., Bryan G., Putman M., van der Marel R. P., Tonnesen S., 2015, ApJ, 815, 77
Sanchez N. N., Werk J. K., Tremmel M., Pontzen A., Christensen C., Quinn T., Cruz A., 2019, ApJ, 882, 8
Schneider E. E., Ostriker E. C., Robertson B. E., Thompson T. A., 2020, ApJ, 895, 43
Segers M. C., Oppenheimer B. D., Schaye J., Richings A. J., 2017, MNRAS, 471, 1026
Sharma P., McCourt M., Quataert E., Parrish I. J., 2012, MNRAS, 420, 3174
Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, MNRAS, 452, 575
Smith R. K., Brickhouse N. S., Liedahl D. A., Raymond J. C., 2001, ApJ, 556, L91
Sokołowska A., Mayer L., Babul A., Madau P., Shen S., 2016, ApJ, 819, 21
Sotillo-Ramos D. et al., 2022, MNRAS, 516, 5404
Sparre M., Pfrommer C., Ehlert K., 2020, MNRAS, 499, 4261
Springel V., 2010, MNRAS, 401, 791
Springel V., Hernquist L., 2003, MNRAS, 339, 289
Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726
Springel V. et al., 2018, MNRAS, 475, 676
Stern J., Hennawi J. F., Prochaska J. X., Werk J. K., 2016, ApJ, 830, 87
Stern J., Fielding D., Faucher-Giguère C.-A., Quataert E., 2019, MNRAS, 488, 2549
Stocke J. T., Keeney B. A., Danforth C. W., Shull J. M., Froning C. S., Green J. C., Penton S. V., Savage B. D., 2013, ApJ, 763, 148
Strong A. W., Moskalenko I. V., Reimer O., 2000, ApJ, 537, 763
Su M., Slatyer T. R., Finkbeiner D. P., 2010, ApJ, 724, 1044
Terrazas B. A. et al., 2020, MNRAS, 493, 1888
Truong N. et al., 2020, MNRAS, 494, 549
Truong N., Pillepich A., Werner N., 2021a, MNRAS, 501, 2210
Truong N., Pillepich A., Nelson D., Werner N., Hernquist L., 2021b, MNRAS, 508, 1563
Tumlinson J., Peeples M. S., Werk J. K., 2017, ARA&A, 55, 389
van de Voort F., Bieri R., Pakmor R., Gómez F. A., Grand R. J. J., Marinacci F., 2021, MNRAS, 501, 4888
Vogelsberger M. et al., 2014a, MNRAS, 444, 1518
Vogelsberger M. et al., 2014b, Nature, 509, 177
Voit G. M., Meece G., Li Y., O'Shea B. W., Bryan G. L., Donahue M., 2017, ApJ, 845, 80
Wakker B. P., van Woerden H., 1997, ARA&A, 35, 217
Weinberger R. et al., 2017, MNRAS, 465, 3291
Werk J. K. et al., 2016, ApJ, 833, 54
Werk J. K. et al., 2019, ApJ, 887, 89
Yao Y., Shull J. M., Danforth C. W., 2011, ApJ, 728, L16
Yun K. et al., 2019, MNRAS, 483, 1042
Zech W. F., Lehner N., Howk J. C., Dixon W. V. D., Brown T. M., 2008, ApJ, 679, 460
Zhu G. et al., 2014, MNRAS, 439, 3139
Zinger E. et al., 2020, MNRAS, 499, 768

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 3

The Circumgalactic Medium of Milky Way-like Galaxies in the TNG50 Simulation – II: Cold, Dense Gas Clouds and High-Velocity Cloud Analogs

3.1 Statement of contribution

- Scientific Analysis: My contribution was central. In addition, the analysis benefited greatly from feedback provided by collaborators.
- Figures: I independently produced the figures based on the planned scientific analysis. The process also greatly benefited from iterative input by collaborators.
- Writing: I primarily wrote most of the manuscript text. The drafts were further refined based on feedback from collaborators.
- Code and Simulation Development: This paper made use of the publicly available TNG50 simulation. The code to analyze the simulation data and generate plots was written by me, but it also benefited from input provided by colleagues.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **522,** 1535–1555 (2023) https://doi.org/10.1093/mnras/stad951
Advance Access publication 2023 March 29

The circumgalactic medium of Milky Way-like galaxies in the TNG50 simulation – II. Cold, dense gas clouds and high-velocity cloud analogs

Rahul Ramesh ,[1★] Dylan Nelson [1] and Annalisa Pillepich [2]

[1] *Universität Heidelberg, Zentrum für Astronomie, Institut für theoretische Astrophysik, Albert-Ueberle-Str 2, D-69120 Heidelberg, Germany*
[2] *Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

Accepted 2023 March 24. Received 2023 March 10; in original form 2023 February 6

ABSTRACT

We use the TNG50 simulation of the IllustrisTNG project to study cold, dense clouds of gas in the circumgalactic media (CGM) of Milky Way-like galaxies. We find that their CGM is typically filled with of order one hundred (thousand) reasonably (marginally) resolved clouds, possible analogs of high-velocity clouds (HVCs). There is a large variation in cloud abundance from galaxy to galaxy, and the physical properties of clouds that we explore – mass, size, metallicity, pressure, and kinematics – are also diverse. We quantify the distributions of cloud properties and cloud-background contrasts, providing cosmological inputs for idealized simulations. Clouds characteristically have subsolar metallicities, diverse shapes, small overdensities ($\chi = n_{\rm cold}/\, n_{\rm hot} \lesssim 10$), are mostly inflowing, and have sub-virial rotation. At TNG50 resolution, resolved clouds have median masses of $\sim 10^6\,{\rm M}_\odot$ and sizes of $\sim$10 kpc. Larger clouds are well converged numerically, while the abundance of the smallest clouds increases with resolution, as expected. In TNG50 MW-like haloes, clouds are slightly (severely) underpressurized relative to their surroundings with respect to total (thermal) pressure, implying that magnetic fields may be important. Clouds are not distributed uniformly throughout the CGM but are clustered around other clouds, often near baryon-rich satellite galaxies. This suggests that at least some clouds originate from satellites, via direct ram-pressure stripping or otherwise. Finally, we compare with observations of intermediate and high velocity clouds from the real Milky Way halo. TNG50 shows a similar cloud velocity distribution as observations and predicts a significant population of currently difficult-to-detect low velocity clouds.

Key words: galaxies: haloes.

1 INTRODUCTION

The circumgalactic medium (CGM), the halo of gas surrounding galaxies, is believed to be critically linked to their formation and evolution. While the bulk of the CGM is home to a volume-filling diffuse warm-hot gas phase, the CGM is also commonly multiphase. It can host small clouds of cold, dense gas (see Donahue & Voit 2022 for a recent review of the CGM).

Historically observed through their HI emission, compact gas clouds ($\lesssim$1–10 kpc) are observed in the Milky Way (MW) halo at large velocities with respect to the local standard of rest (LSR) and have hence been named high-velocity clouds (HVCs; e.g. Muller, Oort & Raimond 1963; Wakker 1991; Wakker & van Woerden 1997). More recently, such clouds have been further differentiated based on their velocities, into so called low-velocity clouds, intermediate-velocity clouds (IVCs), and very high-velocity clouds (VHVCs; e.g. Haffner, Reynolds & Tufte 2001; Peek et al. 2009; Lehner & Howk 2011). While early studies mainly targeted the MW halo, which is believed to contain many thousands of such clouds (e.g Putman et al. 2002; Moss et al. 2013), more recent explorations have begun identifying clouds around external galaxies as well (e.g Gim et al. 2021).

Although these clouds have been observed for many decades now, there remain multiple open questions. For example, their origin is highly debated: while a fraction of these clouds may be related to the stripping of gas as satellites infall into the potential minimum of their host halo (e.g. Olano 2008), theory also suggests that clouds can form via condensation of hot halo gas (e.g. Binney, Nipoti & Fraternali 2009; Joung, Bryan & Putman 2012; Fraternali et al. 2015) and 'fountain' flows of gas in and around galaxies (e.g. Fraternali & Binney 2006). Recently, Lehner et al. (2022) and Marasco et al. (2022), using a sample of (observed) IVCs and HVCs, showed that both diffuse 'rain-like' inflows and collimated outflows are present in their sample, adding weight to the galactic fountain scenario. They, however, do note that only $\sim$30 per cent of their clouds are outflowing, which is slightly lower than expected with the galactic fountain scenario, although such a bias may just stem from the collimated geometry of outflows.

Since HVCs were first observed through their 21cm line emission, it was assumed that these clouds are typically dominated by cold gas and are pristine with respect to their metallicity content. However, recent observations of metal-line absorption, along the line of sight (LOS) to quasars and stars, have shown that these clouds may indeed contain non-negligible amounts of metals. For instance, using data collected by the far-ultraviolet spectroscopic explorer, Sembach et al. (2000) observed OVI absorption in HVCs, while Savage et al. (2000) reported the correlation between MgII absorption and known locations of HVCs using the *Hubble Space Telescope*. Lehner,

★ E-mail: rahul.ramesh@stud.uni-heidelberg.de

Keenan & Sembach (2001) showed the presence of CII, OI, SII, SIII, and SIV along a sightline that intercepts a known HVC. Using observations with the Cosmic Origin Spectrograph, Richter et al. (2017) noted the presence of SiIII in HVCs. Further, the metallicity distributions across a sample of HVCs can be large, with values ranging from highly subsolar to supersolar (e.g. Wakker 2001; Fox et al. 2016); on the other hand, the range of metallicities for IVCs is relatively narrower (Lehner et al. 1999; Wakker 2001), possibly indicating varied and different origins. Finally, observations of Hα emission from these clouds suggest the presence of slightly warmer gas, potentially in interface layers between the cold clouds and the background hot halo (Tufte, Reynolds & Haffner 1998; Haffner et al. 2001).

While distances to HVCs are generally unconstrained, given that they are predominantly observed through HI emission, more recent studies have begun providing (upper limit) distance estimates for a large number of clouds. For instance, absorption features along the LOS of stars have been used to estimate a distance of $\leq$10 kpc to Complex C (Wakker et al. 2007), $\leq$15 kpc to Complex GCP (Wakker et al. 2008), and $\sim$4.4 kpc to Complex WD (Peek et al. 2016). In addition to absorption studies, measurements of magnetic field strengths around clouds may be useful in estimating their distances (Grønnow et al. 2017). A consequence of the (earlier) lack of distance estimates was that basic properties of clouds like their masses and sizes were poorly understood. Current work now suggests that the aforementioned properties of clouds show large diversity, varying from large HI masses of $\sim 10^7\,\mathrm{M}_\odot$ (Complex C, Thom et al. 2008) down to $\lesssim 10^5\,\mathrm{M}_\odot$ (Wakker 2001; Adams, Giovanelli & Haynes 2013). Despite employing these novel techniques to estimate distances to clouds, an unbiased view of HVCs in the MW is currently not available, such that the typical distance to HVCs is not known, and neither are the mean mass or size or e.g. the mass/size distributions.

Another mystery surrounding these cold CGM clouds is their expected lifetime, i.e. long-term survival in the face of fluid instabilities and mixing processes. A large variety of numerical 'cloud crushing' simulations have studied these questions. Early works typically suggested that these clouds are short-lived, either because they are destroyed (Klein, McKee & Colella 1994; Schneider & Robertson 2017) or broken into smaller fragments (Mellema, Kurk & Röttgering 2002) by shocks and/or hot winds driven by supernovae. That is, the cloud-shredding time-scale is of order the time required to accelerate to their relatively large velocities (Zhang et al. 2017). However, more recent studies propose that they may actually not be as unstable as previously thought. For instance, the Kelvin-Helmholtz (KH) instability can create a turbulent mixing layer between the cloud and the surrounding halo gas, giving rise to a warm gas interface layer that cools rapidly (Nelson et al. 2020), thereby contributing a sizeable amount of cold gas to the cloud (Gronke & Oh 2018; Fielding et al. 2020). In particular, Gronke & Oh (2020) suggest that radiative cooling may be more important than, and win against, the KH instability.

In addition, magnetic fields may play an important role in stabilizing these clouds, either through the associated magnetic pressure that counterbalances the thermal pressure of the ambient hot medium (Nelson et al. 2020), through magnetic tension that suppresses buoyant oscillations of the condensing gas (Ji, Oh & McCourt 2018), or possibly by enhancing the Rayleigh–Taylor instability that leads to larger rates of condensation (Grønnow et al. 2022). Similar to magnetic fields, cosmic rays may also contribute pressure support, thereby providing additional stability to these cold gas clouds (for e.g. Butsky & Quinn 2018). Furthermore, under the right conditions, cold clouds of gas may coagulate, i.e. small fragments may coalesce into a larger mass (Gronke & Oh 2022). However, it is to be noted that the outcomes of these simulations generally depend upon the specifics of the clouds versus background properties and setup, the physical processes included, and numerical resolution (e.g. Jennings et al. 2022). As of today, cloud survival remains an open topic.

Idealized numerical studies of cloud survival have a fundamental limitation: they are all non-cosmological, 'wind tunnel' simulations, i.e they assume the existence of a pre-formed cloud of a given composition and evolve this cloud in the presence of a hot-ambient medium. As a result, these simulations do not account for the complexity and structure of a realistic CGM, its diversity across the galaxy population, nor its evolution across cosmic time. In addition, such numerical experiments cannot answer questions that are input assumptions, namely, the origin and properties of clouds, including their mass and size distributions, as well as their composition, and relative dynamics with respect to the background CGM.

In this paper, we use the TNG50 simulation of the IllustrisTNG project to investigate the existence and properties of cold clouds around MW-like galaxies. TNG50 is a cosmological uniform-volume simulation that has been shown to be able to realize small-scale, cold gas structures in the CGM of high-mass elliptical galaxies, as traced by neutral HI and MgII (Nelson et al. 2020). As a magnetohydrodynamical simulation over a large cosmological volume, TNG50 is able to account for the physically plausible, complex, and diverse CGM of MW-like galaxies, together with their embedded cloud populations. Here, in particular, we focus on the $z = 0$ CGM of 132 MW-like galaxies whose global gaseous-halo properties have already been extensively characterized by Ramesh, Nelson & Pillepich (2023), focusing on the physical properties of their cold clouds.

This paper is organized as follows: in Section 2, we provide a brief description of the IllustrisTNG simulations and TNG50, the sample selection process, and the algorithm we employ to identify clouds. In Section 3, we present the results of this work: the abundance and global statistics of cold clouds in the CGM, including their intrinsic physical properties and distribution through the halo. We demonstrate a suggestive correlation between the location of clouds and satellite galaxies with relatively large baryon fractions. Finally, in Section 4, we summarize our main results and conclude.

2 METHODS

2.1 The TNG50 simulation

In this paper, we use the TNG50-1 (hereafter, TNG50) simulation (Nelson et al. 2019b; Pillepich et al. 2019) of the IllustrisTNG project (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018b; Springel et al. 2018). These are a set of cosmological magnetohydrodynamical simulations run with the code AREPO (Springel 2010). IllustrisTNG builds on its predecessor, the original Illustris simulation (Genel et al. 2014; Vogelsberger et al. 2014a, b; Sijacki et al. 2015), employing a modified model for galaxy physics (Weinberger et al. 2017; Pillepich et al. 2018a), along with the addition of magnetic fields (Pakmor, Marinacci & Springel 2014), an important physical quantity that was previously absent in Illustris.

TNG50 simulates a volume of $\sim$(50 cMpc)^3 at an average baryonic mass resolution of $\sim 8 \times 10^4\,\mathrm{M}_\odot$. It remains the highest resolution cosmological simulation that exists at this volume and the largest volume cosmological simulation that exists at this resolution. As explained in detail in Pillepich et al. (2018a), the TNG simulations include recipes for a variety of physical processes that are believed to play a critical role in galaxy formation and evolution. We refer the reader to that paper for all details of the TNG model.

We briefly mention that the TNG simulations employ the Springel & Hernquist (2003) two-phase subgrid model to stochastically convert star-forming gas ($n_H \geq 0.1$ cm^{-3}) to stars. As a result, the temperature of star-forming gas is 'effective', i.e. not physical. For our analysis, we always set the temperature of star-forming gas to its cold phase value, 10^3 K,[1] which dominates by mass ($\geq$90) per cent.

While the amount of neutral hydrogen content of gas is directly tracked and output by the simulation, the fraction of this component in atomic hydrogen is not. We use the Gnedin & Kravtsov (2011) H_2 model to estimate the fraction of neutral hydrogen in H_2, thereby arriving at an estimate of neutral atomic hydrogen by subtracting the two values, following Popping et al. (2019).

The TNG simulations adopt a cosmology consistent with the Planck 2015 analysis (Planck Collaboration XIII 2016), with: $\Omega_\Lambda = 0.6911$, $\Omega_m = 0.3089$, $\Omega_b = 0.0486$, and $h = 0.6774$.

2.2 The Milky Way-like galaxy sample

In this work, we use the same sample of MW-like galaxies as Ramesh et al. (2023)[2]: these are a set of 132 galaxies that are (i) centrals, i.e. they lie at the potential minimum of their host friends-of-friends (Davis et al. 1985) haloes; (ii) reside in haloes that are not overly massive (virial mass, $M_{200,c} < 10^{13}$ M$_\odot$); (iii) have a stellar mass, measured within a 3D aperture of 30kpc, in the range $10^{10.5}$ M$_\odot$ to $10^{10.9}$ M$_\odot$; (iv) are reasonably well isolated (no other galaxy having $M_\star > 10^{10.5}$ M$_\odot$ within a distance of 500 kpc); and (v) are discy, either through visual inspection of stellar light maps or based on a constraint on the minor-to-major axis ratio of the stellar mass distribution ($s < 0.45$). We refer the reader to Pillepich et al. (in preparation) for the motivation behind the choice of these criteria and the extent to which such a sample captures our real MW.

Following Ramesh et al. (2023), we (i) define the circumgalactic medium (CGM) as the region bounded by $[0.15, 1.0] \times R_{200c}$ of the corresponding halo, (ii) exclude gas that is gravitationally bound to satellite galaxies (i.e. all galaxies in a halo that are not centrals), as identified by the substructure identification algorithm SUBFIND (Springel et al. 2001), and (iii) consider all gas in the simulation volume, i.e. we do not restrict our selection of gas cells based on their friends-of-friends halo membership. Unless otherwise stated, the stellar mass of a galaxy is defined as the sum of the mass of all stars within an aperture of 30 pkpc. Star formation rates are defined by the stellar mass formed within an aperture of 30 pkpc, over the last billion years (Donnari et al. 2019; Pillepich et al. 2019).

2.3 Cold gas cloud identification algorithm

A consistent way to identify clouds in simulations run with AREPO is via spatially contiguous sets of Voronoi cells (following Nelson et al. 2020). In particular, 'natural neighbours' of a Voronoi cell are those that share a face, i.e. they directly touch. An ensemble of naturally connected Voronoi cells is a union of convex polyhedra, and this geometrical structure is well suited to represent clouds of arbitrary size and shape. Our algorithm consists of two steps:

(i) We first identify 'cold' gas as those cells whose temperature $T < 10^{4.5}$ K. We note that our algorithm is not overly sensitive to this chosen threshold value. For instance, varying the threshold by ∼0.2 dex only changes the average number, and size, of clouds by ∼10 per cent. This suggests that the different temperature cutoffs identify the same clouds, but include different amounts of interface gas around their cores (see Nelson et al. 2020).

(ii) Once cold gas cells are identified, we group them into clouds by looking for contiguous sets of Voronoi cells. For the main results of this work, we consider only those clouds that contain at least 10 Voronoi member cells. Various numerical experiments have concluded that clouds need to be resolved by at least few tens to hundreds of resolution elements to adequately capture their growth and evolution (e.g. Klein et al. 1994; Nakamura et al. 2006; Yirak, Frank & Cunningham 2010; Goldsmith & Pittard 2016; Pittard & Goldsmith 2016). We therefore consider a lower limit of 10 member cells per cloud. This avoids issues with low number statistics, while restricting our analyses to marginally well-resolved clouds. We note that most of our results are *qualitatively* similar even if a slightly lower threshold of member cells per cloud is adopted. Cases that differ qualitatively are shown explicitly in the corresponding panels.

We run the algorithm for all gas within R_{200c}, with satellite gas excluded. For each halo, the algorithm typically returns one 'massive' cloud of mass $\gtrsim 10^{8.5}$M$_\odot$ that lies close to the centre of the halo, i.e. the galaxy itself. We exclude this object from our analysis. In addition, unless otherwise stated, we exclusively consider clouds that lie outside $0.15 \times R_{200c}$, which is our inner boundary for the CGM.

3 RESULTS

3.1 AREPO, TNG50, and CGM cold clouds

To provide an idea of the cold clouds that are present in the TNG50 simulation in the CGM of MW-like galaxies, Fig. 1 shows a visualization of an individual cloud identified by our algorithm (Section 2.3). This cloud has a mass of $M_{\rm cloud} \sim 10^{6.4}$ M$_\odot$ and is composed of 33 Voronoi cells. Each panel shows a slice of the Voronoi mesh in a small region centred on the chosen single small cloud. The Voronoi cells that belong to this cloud are outlined with white/black lines. Note that only a subset of the 33 Voronoi cells are visible in these panels, since cells displaced along the direction perpendicular to the screen are not visible in the slice.

In the top-left panel, we colour by (3D) gas density: at the centre of the cloud, gas is more dense compared to its outskirts, and the density drops with increasing distance from the cloud centre, into the background region surrounding this cloud. Since AREPO dynamically (de-)refines the mesh to ensure that the mass of all gas cells is roughly equal, the Voronoi cells at the centre of the cloud are naturally smaller, resulting in higher spatial resolution.

The top-right panel shows the temperature of gas: the centre of this cloud is dominated by cold gas (T $\lesssim 10^{4.5}$K) and is embedded in a pre-dominantly hot medium (T $\gtrsim 10^{5.5}$K). Although the numerical resolution in TNG50 at these small scales is limited, we see that the cloud and background regions are separated by an intermediate warm-phase ($10^{5.5}$K $\gtrsim$ T $\gtrsim 10^{4.5}$K), which may play a role in increasing the longevity of such clouds (e.g. Gronke & Oh 2018; Fielding et al. 2020; Nelson et al. 2020; Abruzzo, Fielding & Bryan 2022).

In the bottom-left panel, we colour by the thermal pressure of gas. Gas in this cloud is thermally underpressurized with respect to the

[1]This is the temperature invoked in the TNG model for the cold phase of the subgrid ISM model (Springel & Hernquist 2003) and does not necessarily reflect the true gas temperature in star-forming gas complexes, clouds, or cores. For our purposes, any value below $T < 10^4$ K will not significantly impact the present results.

[2]This is a subset of the TNG50 'MW/M31 sample' of galaxies presented in Pillepich et al. (in preparation), where these galaxies are discussed.

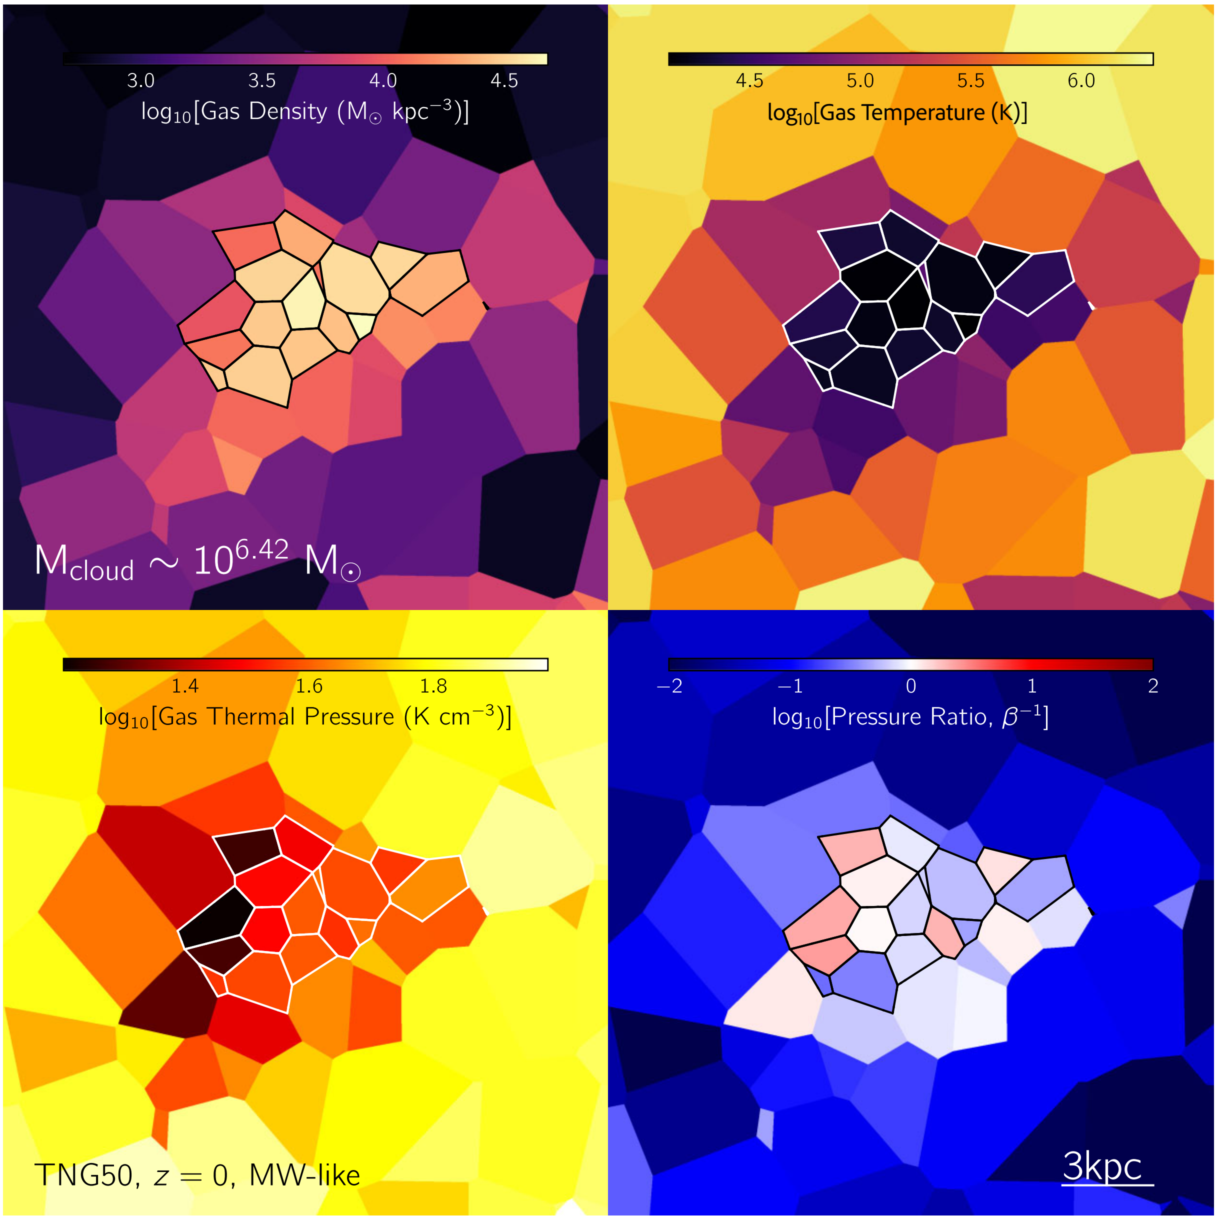

Figure 1. A visualization of the physical properties of a single cloud ($M_{\rm cloud} \sim 10^{6.4}$ M$_{\odot}$) from the CGM of a MW-like galaxy in the TNG50 simulation. Each panel shows a slice of the Voronoi mesh centred on the same cloud. Voronoi cells that are members of the cloud are outlined by either black or white lines. In cyclic order, from the top left to bottom left, cells are coloured by their density, temperature, ratio of magnetic to thermal pressure, and thermal pressure. The central region of this cloud is dense and cold and relative to the local environment; this region is overdense, cooler, thermally underpressurized, and has $\beta \sim 1$ with magnetic and thermal pressures in rough balance.

background. The distribution of thermal pressure is asymmetric due to the asymmetric distribution of temperature and density.

Lastly, in the bottom-right panel, we show the ratio of magnetic to thermal pressures, β^{-1}. The majority of the cloud has $\beta \sim 1$, indicating that magnetic and thermal pressure are in rough equipartition. While certain regions of this cloud are dominated by magnetic pressure, the thermal component dominates in others. This is in contrast to clouds in much larger group-mass haloes in TNG50, where magnetic pressure strongly dominates within cold clouds likely due to the higher ambient densities (Nelson et al. 2020).

The properties of the TNG50 cloud showcased in Fig. 1 are typical, as we expand upon in the following sections.

3.2 Location and number of cold clouds around MW analogs

Zooming out from a single cloud, we visualize the entire distribution of clouds in one halo from our sample of TNG50 MW-like galaxies. Fig. 2 shows the gas column density in the background, integrated over an extent of $\pm R_{200c}$ along the LOS direction. The two concentric white circles are drawn at radii $[0.15, 1.0] \times R_{200c}$, the chosen inner and outer boundaries of the CGM. In the foreground, the distribution of clouds is shown: clouds composed of less than 10 Voronoi member cells (which we do not generally include in our analyses) are drawn as unfilled white circles, with their radii scaled by the size of the cloud; more massive clouds, i.e. the clouds we consider in this work, are shown instead directly by their (projected) convex hulls, as translucent filled shapes, to provide a better sense of their

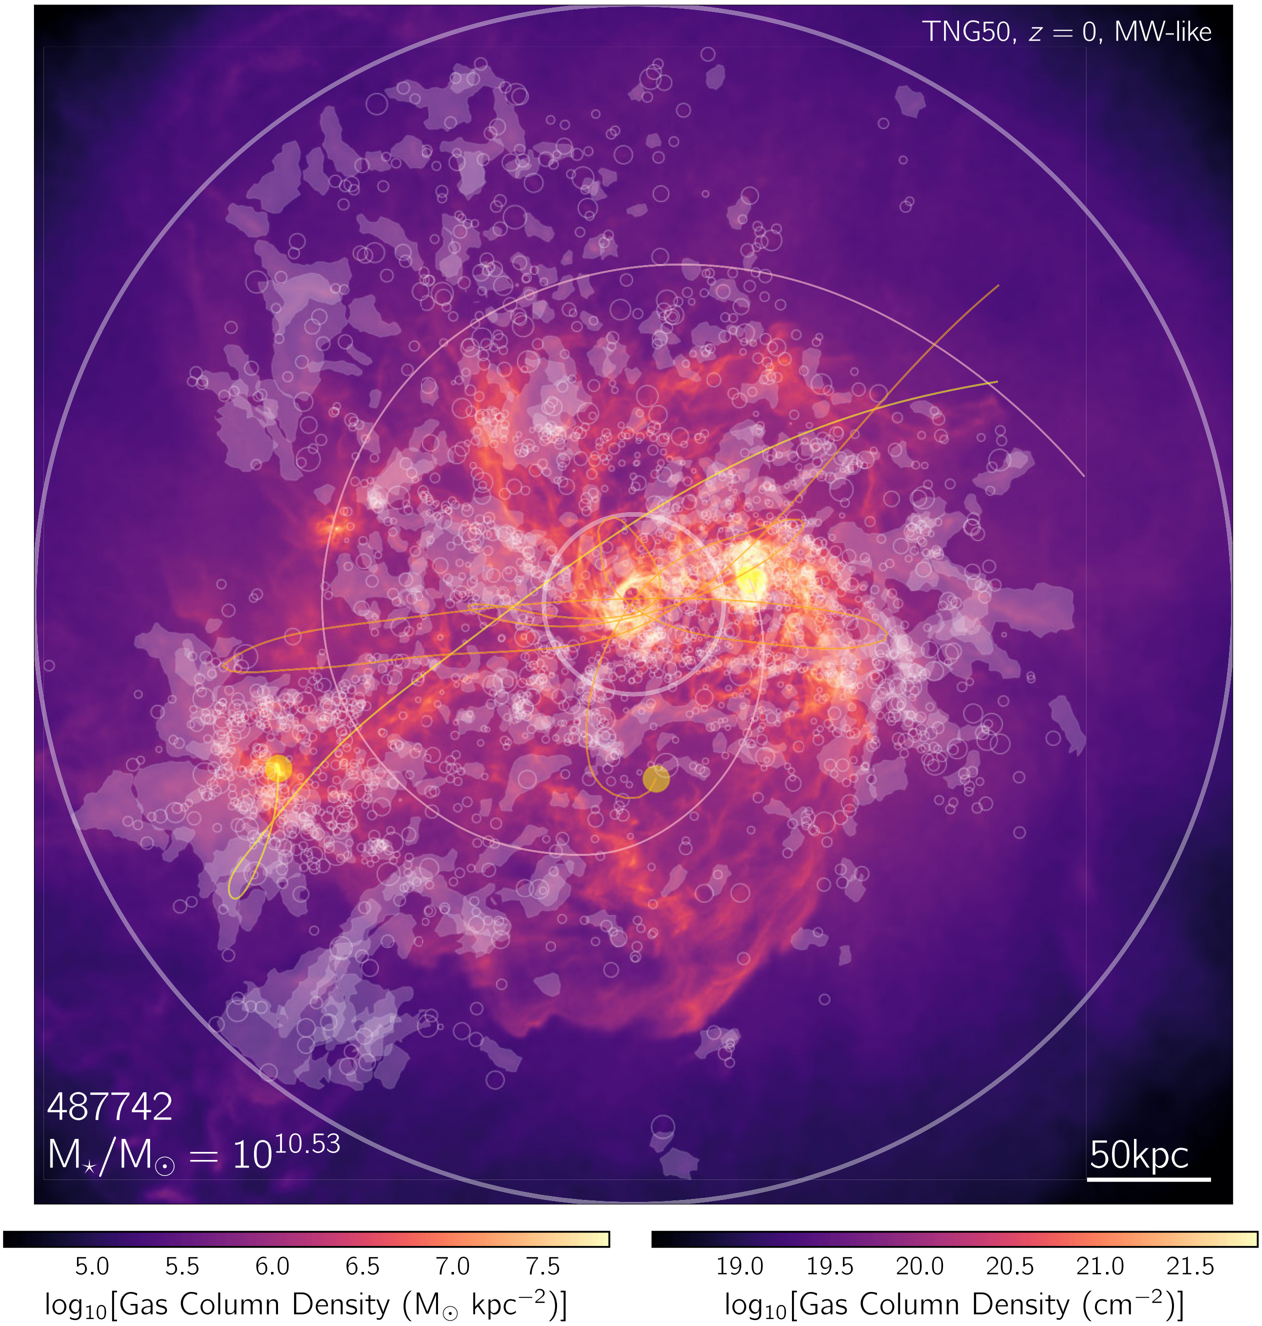

Figure 2. A visualization of the distribution of clouds around one of the TNG50 MW-like galaxies from our sample (subhalo 487742). Background colour shows gas column density, while the two concentric white circles are drawn at radii $[0.15, 1.0] \times R_{200c}$. Unfilled white circles in the foreground correspond to unresolved clouds, i.e. those composed of less than 10 gas cells and thus not included in our main analyses. Translucent shapes correspond to more massive clouds, which comprises our 'fiducial' sample. The yellow circles show the location of satellite galaxies with large baryon ratios (>10 per cent), while the connected curves show the past trajectory of these satellites. Clouds are preferentially seen to lie close to satellites or close to the their past trajectories. We interpret the former as gas that has been 'freshly' stripped, and the latter as gas that was stripped in the past.

complex structure. Note that regions where shapes appear 'brighter' correspond to overlapping clouds along the LOS. It is clear that (i) an enormous number of cold clouds are present in the CGM of MW-like galaxies, and (ii) CGM cold gas clouds come in a variety of shapes and sizes.

The yellow filled circles correspond to the positions of satellite galaxies with large baryon ratios (>10 per cent; at $z = 0$), and the curves connected to each trace their past orbital trajectories. A large number of clouds lie close to these satellites or close to their past orbits.[3] However, there are some clouds that are close to neither, and such clouds could be linked to satellites that are no longer baryon rich (but were at some point in the past), or could just have drifted

[3]We encourage the reader to explore the significant diversity of cloud populations across the sample in our online infinite gallery.

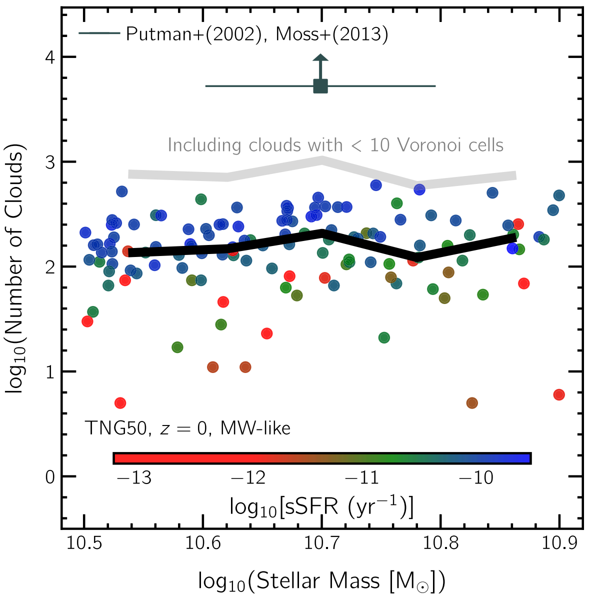

Figure 3. Number of clouds in the CGM around each TNG50 MW-like galaxy as a function of its stellar mass, with scatter points coloured by the sSFR of the galaxy. On average, the number of clouds show a flat trend with stellar mass of the galaxy, albeit with a large scatter at fixed galaxy stellar mass. A trend in the colours of the points is, however, seen: galaxies with higher sSFRs have a larger number of clouds in their CGMs. The black (grey) solid curve depicts median results for our fiducial (alternative) definition of clouds, i.e. with a minimum number of 10 (1) cells each.

away after having been stripped. This figure conveys the concept that we expand upon in the next sections (originally discussed in Nelson et al. 2020): clouds cluster around satellites. Satellite galaxies clearly make an important contribution to the cold gas contents of the CGM (Rohr et al., in preparation). In addition, many clouds also exist close to other clouds, as opposed to being uniformly distributed through the halo: clouds cluster around themselves. Both these correlations are key results of this work, which we later quantify in Section 3.6.

As discussed in Ramesh et al. (2023), the properties of halo gas across our sample of MW-like galaxies exhibit a large diversity. To explore the effect of this galaxy-to-galaxy diversity on the properties of clouds, Fig. 3 plots the number of CGM clouds as a function of stellar mass of the galaxy. The circles, coloured by the specific star formation rate (sSFR) of the galaxy, denote individual haloes, while the median across galaxies is shown in the black solid curve.

The CGM of TNG50 MW-like galaxies typically contain of order one hundred reasonably resolved clouds and of order one thousand marginally resolved clouds. The median shows no significant trend as a function of stellar mass, although a large scatter is evident, with numbers varying between a few clouds per halo to a few hundred. However, a strong trend is apparent in the vertical direction across the scatter: galaxies with higher values of sSFR preferentially host a greater number of clouds in their haloes. As discussed extensively in Ramesh et al. (2023), such a trend with properties related to *cold gas* may arise due to two factors: either (i) because outflows generated by the central supermassive black hole (SMBH) both suppress star formation in the galaxy and heat up CGM gas (e.g. Weinberger et al. 2017; Zinger et al. 2020), or (ii) because of a physical connection between the flow of gas through the CGM and the SF activity of the galaxy. The latter case suggests that these clouds may play a role in sustaining star formation in the galaxy by replenishing the required 'fuel', as is commonly proposed for HVCs in the MW halo (e.g. Lepine & Duvert 1994). In the former case, it is possible that high-velocity outflows generated by the kinetic mode of SMBH feedback in TNG50 destroy clouds that are already present (analogous to clouds being shredded by supernova feedback, e.g. Schneider & Robertson 2017), but it may also be the case that clouds are prevented from forming in the first place.

A robust comparison with observations of the real MW halo is not possible. However, we take a first step with the grey square, which shows a lower limit for the number of observed HVCs. To arrive at this estimate, we stack the catalogs presented in Putman et al. (2002) and Moss et al. (2013). These contain a total of 1956 and 1693 clouds, respectively, with an overlap of 1021 clouds: the stacked sample therefore contains a total of 2628 unique clouds. Since both these surveys observe only the southern sky, we assume that the northern sky contains an equal number of clouds, and multiply the sample size by two, yielding a lower limit of 5256 clouds. For the stellar mass, we assume a value of $M_{\star,\mathrm{MW}} \sim (5 \pm 1) \times 10^{10}$ M$_\odot$ (Bland-Hawthorn & Gerhard 2016).

This observed value is more than an order of magnitude above our most populous halo. However, several caveats exist in this comparison: the observations only report HVCs, i.e. only those clouds that satisfy a given velocity threshold, while we do not here apply any such cut in TNG50. Most clearly, the number of clouds in simulations depends on numerical resolution, and we expect more clouds at better resolution (Nelson et al. 2020, Fig. 4, and Ramesh & Nelson in preparation). The number of identified clouds is greater if we relax the minimum number of cells threshold. In this case, the grey curve shows the median relation for the case where all clouds with at least one member Voronoi cell are considered, i.e. if poorly resolved clouds are also included. As seen, the sample size increases by a factor of ∼10, with a largely similar trend with respect to stellar mass. We undertake a more realistic comparison with observations in Section 3.7.

3.3 Physical properties of clouds in TNG50 MW-like galaxies

Fig. 4 quantifies the physical properties of our TNG50 cold clouds. In the top-left panel, we show the probability distribution function (PDF) of cold cloud masses. The distributions for individual MW-like galaxies are shown with thin curves, coloured by the mass of the central SMBH. The thick black curve shows the median of the entire galaxy sample, while the purple and yellow curves show the median of the lowest and highest octiles, dividing the sample into eight percentiles based on SMBH mass.

The cloud mass distribution peaks around $10^{6.1}$ M$_\odot$, which is slightly more than 10 times the average mass of baryon resolution elements (gas cells) in TNG50. The PDF drops sharply towards the left of this peak, as a result of our cloud definition and the resolution limit of the simulation. More massive cold clouds are rarer. No strong trend with respect to SMBH mass is seen in case of distributions of cloud masses. For reference, in grey, we also show the median distribution for the case where clouds down to one member cell are considered, offsetting this distribution vertically downwards better visibility. The shape of the grey curve largely resembles the one in black, however, with the peak shifted towards lower masses.

While masses of HVCs in the MW halo are poorly constrained owing to a lack of distance estimates, recent studies suggest that the typical HI mass of clouds varies from $\lesssim 10^5$M$_\odot$ (Wakker 2001; Adams et al. 2013) to $\sim 10^7$M$_\odot$ (Thom et al. 2008). With TNG50, we can thus start to study clouds for a similar mass regime as in the

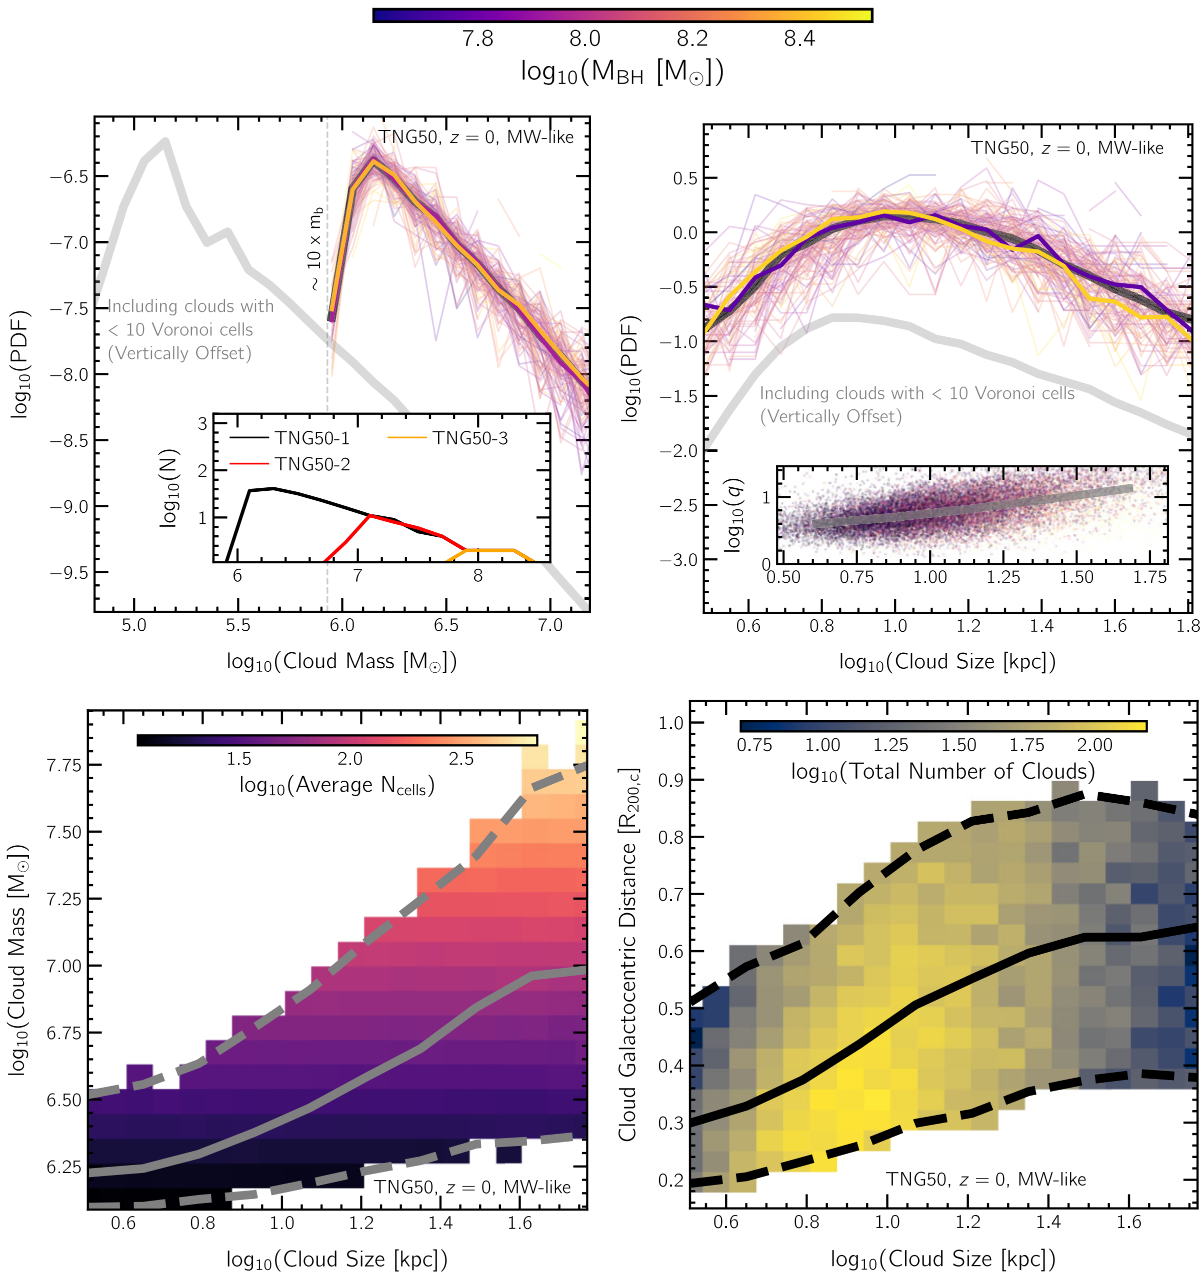

Figure 4. Properties of the clouds in the CGM of TNG50 MW-like galaxies: the top-left and top-right panels show the cloud mass and cloud size distributions, respectively. In each panel, the thin curves correspond to individual haloes, colour-coded by central SMBH mass. Solid black curves show the median PDF across the sample, and purple and yellow curves correspond to median distributions of the two extreme octile regions, splitting by SMBH mass. The top-left panel includes an inset where the mass function of clouds is compared to lower resolution versions (TNG50-2 and TNG50-3) of TNG50 (i.e. TNG50-1), demonstrating good resolution convergence above the resolution limit threshold of each simulation. The top-right panel shows an inset corresponding to the relation between the major-to-minor axis ratio (q) and cloud size. In the bottom-left panel, we show the cloud mass–size relation, with the background coloured by the average number of gas cells per gas cloud, and the bottom-right panel shows the cloud distance–size relation, with the coloured instead showing the total number of clouds in each given pixel. The most abundant clouds have masses of $\sim 10^6\,\mathrm{M}_\odot$ and sizes of $\sim$10 kpc. Lower mass clouds, with smaller sizes, tend to reside at smaller galactocentric distances.

real MW, with the caution of limited resolution towards the low mass end.

In the inset, we assess numerical resolution convergence. To do so, we compare the cloud mass distribution from TNG50-1 (i.e. TNG50, black) with its lower resolution counterparts, TNG50-2 (red; 8 times lower mass resolution) and TNG50-3 (orange; 64 times lower mass resolution). The curves correspond to the median across the sample of MW-like galaxies, matched between pairs of runs. Importantly, we see good convergence between the different runs above the resolution limit: namely, the number density of 'massive' clouds is roughly independent of resolution. In each case, the mass function peaks, turns over, and then drops rapidly at roughly 10 times the resolution limit of the simulation, corresponding to a minimum of 10 member gas cells per cloud. As expected, this peak shifts towards higher masses for the lower resolution runs. This is directly analogous to the halo mass function in any cosmological simulation, where resolution convergence demands that the space number density (and properties) of haloes *above the resolution limit* are in good agreement (e.g Boylan-Kolchin et al. 2009; Prada et al. 2012; Springel et al. 2021). Clouds (or haloes) can only exist above a given mass resolution threshold, and smaller structures are simply absent at lower resolution.

The top-right panel of Fig. 4 shows the cloud size distribution. To estimate sizes, we fit the vertices of Voronoi cells of clouds to an ellipsoid and consider the size to be the geometric mean of the the lengths of the three axes. As before, we construct PDFs for each galaxy, all of which are shown with thin curves, coloured by SMBH mass. The thick black curve corresponds to the median across the sample, while the purple and yellow curves are the median of the two extreme octiles in SMBH mass.

The cloud size distributions peak around ∼10 kpc, and the median PDF drops monotonically on either side. A weak trend of cloud sizes is visible as a function of SMBH mass: galaxies with less massive SMBHs have a slightly greater fraction of clouds in the low-size regime and a slightly lower fraction of clouds in the high-size regime, as compared to galaxies with more massive SMBHs at their centres. We believe that this is linked to the radial distance-distributions of clouds in these haloes, which we return to in Section 3.5. When the sample is split based on other quantities such as stellar mass or halo mass, a much weaker trend is observed. In grey, we also show the median distribution for the case where clouds with less than 10 member Voronoi cells are included, with the distribution vertically offset for better visibility. The shape of the grey curve largely resembles the one in black, however, with the peak shifted towards smaller sizes.

In the inset, we show the relation between the major-to-minor axis ratio (q) and cloud size (x-axis) for the entire sample of clouds across all galaxies. Each point corresponds to an individual cloud, with colour scaled in accordance to the mass of the cloud: the darkest points show clouds of mass $\sim 10^6 M_\odot$, while the lightest correspond to $\sim 10^7 M_\odot$. The grey curve shows the median. A clear trend is apparent, wherein small cold clouds are more 'spherical', having lower values of q, in comparison to their more massive counterparts. The spherical nature of small clouds could be a result of poor resolution and/or be linked to the fact that AREPO forces gas cells to be 'round' (Springel 2010). A comparison with higher resolution simulations is required to truly assess the shapes of small clouds. Observationally, data suggest that HVCs typically have sizes of $\lesssim$10 kpc (e.g. Thom et al. 2008), although the lack of distance measurements makes physical size inference challenging.

The bottom-left panel of Fig. 4 shows the relation between cloud mass and cloud size. We again stack all CGM clouds across the full sample of TNG50 MW-like galaxies. The median trend is shown with the solid grey curve, and the 16th and 84th percentile regions with dashed lines. Background colour encodes the average number of member Voronoi cells per cloud, and these pixels are clipped outside the 16th and 84th percentile regions for visual clarity. On average, CGM cold clouds with larger sizes are also more massive and are comprised of more gas cells, i.e. they are better resolved.

The bottom-right panel of Fig. 4 shows the relation between the galactocentric distance of clouds (normalized by the virial radius) and their physical size. As before, we stack all clouds across the sample. The solid curve shows the median relation, and dashed curves correspond to the 16th and 84th percentile regions. Background colour shows the number of clouds in each corresponding bin, and pixels are clipped outside the percentile regions. On average, clouds get larger in size (and hence in mass, on average) with increasing distance. Most clouds are present in the inner half of the halo, and have sizes of ∼10 kpc, consistent with the top-right panel. We suspect that this could be a result of increased ram-pressure stripping at smaller galactocentric distances, as a result of which 'big' clouds are fragmented into smaller objects, thereby giving rise to a larger number of clouds in the inner halo, each of which are less massive in comparison to more distant counterparts.

In Fig. 5, we consider several additional thermodynamical and physical properties of cold clouds in the CGM of TNG50 MW analogs: temperature, density, metallicity, pressure, and radial and rotational velocities. As before, we construct individual PDFs for each galaxy but show here only the median of these PDFs for clarity. We further split each PDF into three components based on the galactocentric distance of the clouds: the inner halo (solid curves; $0.15 < r/R_{200,c} \leq 0.4$), central halo (dashed curves; $0.4 < r/R_{200,c} \leq 0.7$), and outer halo (dotted curves; $0.7 < r/R_{200,c} \leq 1.0$). In the upper two rows, we consider two values for each cloud: the mean, and the 90th percentile of the property, computed from all gas that comprises the cloud. The distributions of the former are shown in black curves, while the latter are shown in red. For these upper four panels, in insets, we also show the relation between the mean and 90th percentile values for the stacked set of cold clouds across the entire sample. Each point corresponds to an individual cloud, with colour scaled in accordance to the mass of the cloud: the darkest points correspond to clouds of mass $\sim 10^6 M_\odot$, while the lightest correspond to $\sim 10^7 M_\odot$.

In the top-left panel of Fig. 5, we show distributions of CGM cloud temperatures. As by construction, we aim to identify *cold* gas structures ($T < 10^{4.5}$ K; see Section 2.3) and we expect the cloud-wide average temperatures to be below this threshold value. Clouds in the inner halo have the 'coldest' mean temperatures, with the distribution peaking at $\sim 10^{4.15}$ K. Clouds farther away from the centre are progressively more warm, with distributions shifting horizontally by ∼0.1 and 0.2 dex for the central- and outer-halo, respectively. For each of the three regions of the halo, the distributions corresponding to the 90th percentile values are skewed towards warmer temperatures, suggesting that clouds likely posses significant inhomogeneities in their inner temperature structure, consistent with Nelson et al. (2020). This behaviour is also well captured in the inset, which shows a large spread in the mean-90th percentile plane.

In the top-right panel, we show distributions of the mass densities of the cold clouds. Clouds in the inner regions of the halo are most dense, with a peak in the density distribution at $\sim 10^{4.75}$ $M_\odot$ kpc^{-3}. Clouds farther away from the centre are progressively more rarified, with peaks shifted by roughly −0.4 and −0.8 dex for the central and outer halo distributions, respectively. In all three regions of the halo, distributions of mean and 90th percentiles are largely similar in shape, albeit the latter is skewed towards higher density values. Similar to the previous panel, the inset points towards an inhomogeneity in the inner density structure of clouds, albeit not as pronounced as their temperature structures.

In the centre-left panel, we show distributions of cloud metallicity. The dashed vertical line placed at an x-axis value of 0 demarcates supersolar clouds from their subsolar counterparts. Throughout the MW-like haloes, clouds with (mean) supersolar metallicity (i.e. clouds with $Z \geq Z_\odot$) are subdominant and only account for ∼10.5 per cent of the population of all clouds. In the inner halo, the median cloud metallicity is $\sim 10^{-0.3}$ $Z_\odot$. In the central and outer regions of the halo, metallicity PDFs peak at marginally higher values but are equally broad as the inner halo, stretching $\gtrsim$1 dex between the two extremes. The median PDFs show little difference between the mean

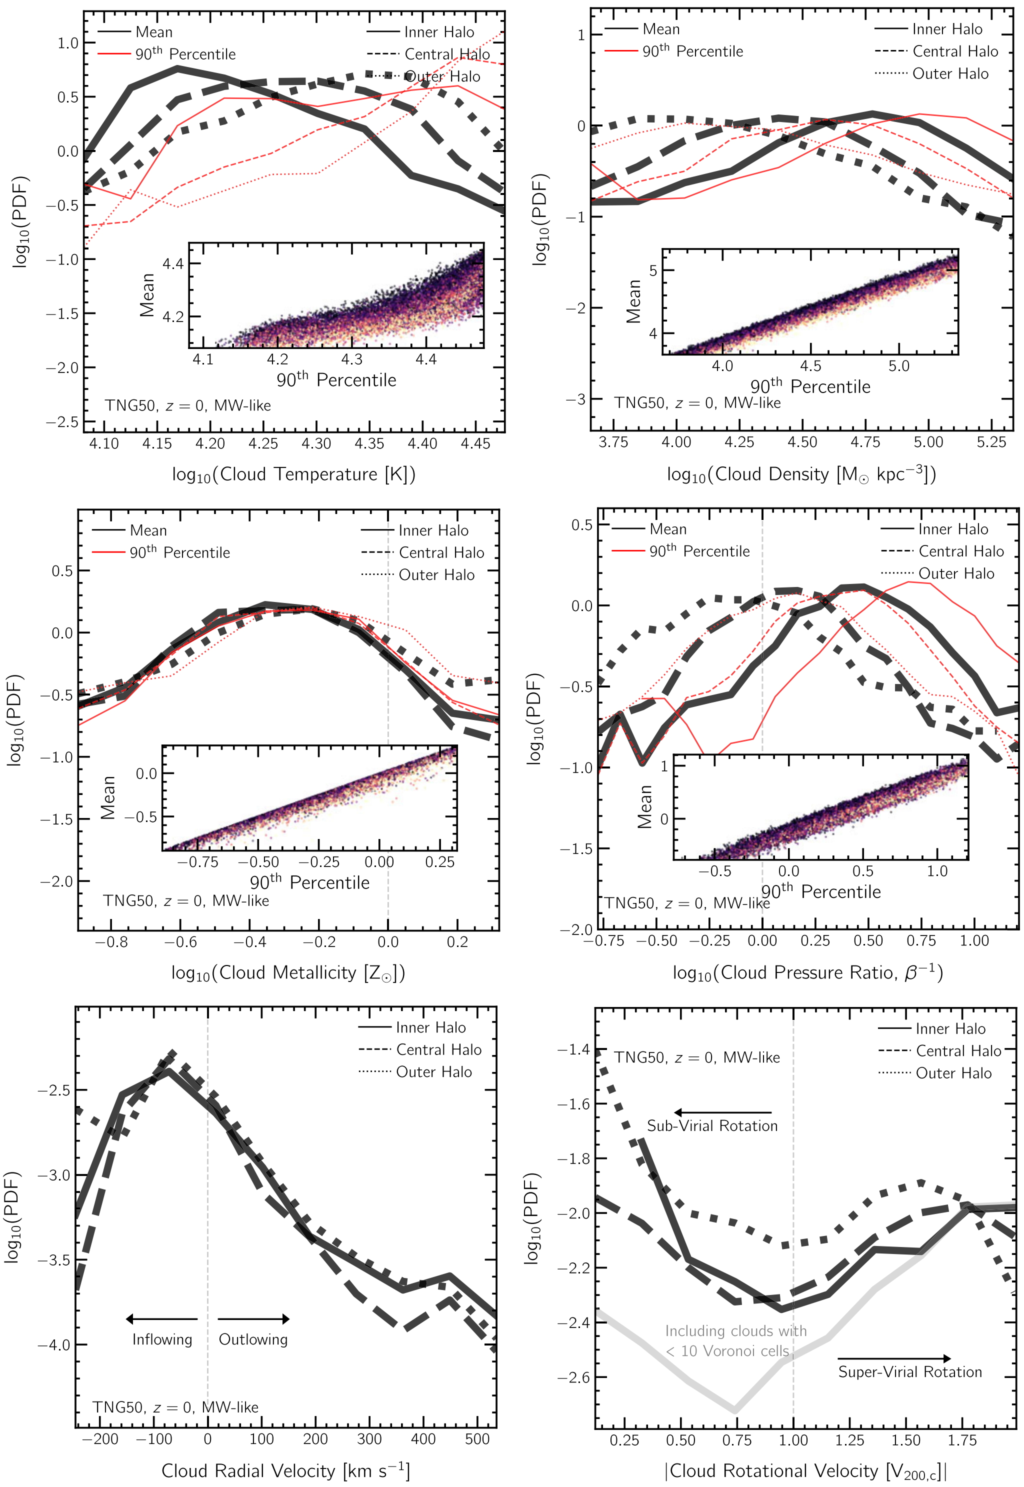

Figure 5. Distributions of physical properties of clouds in the CGM of TNG50 MW-like galaxies. In each panel, we show the median behaviour across the full sample of galaxies, split into three regions of the halo: inner halo (solid curves; $0.15 < r/R_{200,c} \leq 0.4$), central halo (dashed curves; $0.4 < r/R_{200,c} \leq 0.7$), and outer halo (dotted curves; $0.7 < r/R_{200,c} \leq 1.0$). The two upper panels show the distributions of temperature (left) and density (right): clouds closer to the centre of the halo are cooler and denser than those farther away. The centre-left and -right panels show the distribution of cloud metallicity and magnetic to thermal pressure ratio (β^{-1}), respectively: while clouds throughout the halo mainly possess subsolar metallicity and pressure ratios close to unity, clouds in the inner halo are the most likely to be highly enriched and dominated by magentic pressure. The bottom-left and -right panels show distributions of the radial velocity (in $\mathrm{km\,s^{-1}}$) and rotational velocity (normalized by the virial velocity), respectively: a smaller fraction of clouds in the inner halo have dominant rotational motion and a greater fraction are either strongly inflowing or outflowing. Finally, the upper two rows also compare values derived in two ways: the mean value across the gas within each cloud (black; as in the lower panels), and the 90$^{\text{th}}$ percentile value of the gas properties within each cloud (red). The insets show the relation between these two values, demonstrating that TNG50 clouds exhibit inner inhomogeneities in their metal content and temperature-, density- and pressure-structures.

and 90th percentiles for cloud metallicities, although a deviation between the two values is seen in the inset for massive clouds, i.e. well resolved clouds are typically not homogeneous with respect to their metallicity content.

On the observational end, a large range of cloud metallicities have been inferred. For instance, Richter et al. (2001) estimate the metallicities of Complex C and IV Arch to be $\sim$0.1 and $\sim$1 $Z_{\odot}$, respectively, while Collins, Shull & Giroux (2003) and Tripp et al. (2003) propose a slightly higher upper limit of $\sim$0.3–0.6 $Z_{\odot}$ for Complex C. Zech et al. (2008) and Yao, Shull & Danforth (2011) report the possible existence of highly supersolar clouds with metallicities of $\sim$1.65 and $\sim$2.08 $Z_{\odot}$, respectively, although the uncertainties on these measurements are quite large. On the other end of the metallicity spectrum, Tripp & Song (2012) report that the metallicity of HVC gas in the direction of the gaseous 'outer arm' is typically subsolar, possibly around $\sim$0.2–0.5 $Z_{\odot}$. The metal content predicted by TNG50 for the cold clouds around MW-like galaxies encompasses all these observationally inferred metallicity values.

The centre-right panel of Fig. 5 shows distributions of the magnetic to thermal pressure ratio, β^{-1}, of clouds. Overall, $\sim$65.7 per cent of cold clouds in TNG50 MW-like haloes are dominated by magnetic rather than thermal pressure (i.e. $\beta^{-1} > 1$). This is in marked contrast to the case of more massive haloes ($M_{200,c} \gtrsim 10^{13} M_{\odot}$) in the same TNG50 simulation, where magnetic pressure is clearly more dominant with respect to its thermal component for almost the entire population of clouds (Nelson et al. 2020). We suspect that this is a direct effect of the reduced magnetic field strengths in the haloes of MW-like galaxies (see also Faerman & Werk 2023), as compared to their more massive counterparts (Marinacci et al. 2018). The PDF for the inner halo around TNG50 MW-like galaxies has a maxima at $\beta^{-1} \sim 2.8$ and is skewed towards positive values of $\log_{10}(\beta^{-1})$. However, magnetic pressure does not dominate to this extent for clouds in the central and outer regions of the halo. This is likely a consequence of the radial dependence of magnetic field strengths in MW-like haloes (Ramesh et al. 2023). Finally, unlike the case of metallicity, a difference between the distributions of the mean and 90th percentile values is visible: a horizontal offset by $\sim$0.3 dex is present in all three regions of the halo. The inset shows that a larger spread in the mean-90th percentile plane is present, suggesting that some regions of clouds are somewhat dominated more by magnetic pressure than others, similar to what was seen in Fig. 1.

Finally, in the lower panels of Fig. 5, we quantify how the cold clouds move through the CGM of MW-like galaxies in terms of their radial velocity, i.e. the component of the velocity vector that is oriented parallel to the position vector, is shown on the left, and rotational velocity, i.e. the component of the velocity vector along the plane orthogonal to the position vector, on the right.

Across the whole cold cloud sample, TNG50 predicts a relatively robust mix of outflowing versus inflowing cold clouds. Overall, outflowing clouds are subdominant, accounting for $\sim$27.5 per cent of the population. A dependence of radial velocity distributions on distance is present: although all three PDFs peak around the same value, the tails portray different behaviours at the high-velocity end. The inner halo has a slightly larger fraction of clouds moving radially at larger velocities, both in the inflowing and outflowing directions. The central and outer haloes have similar fractions of clouds at the high inflow velocity end. However, the central halo hosts a relatively larger fraction of cold clouds outflowing at large velocities in comparison to the outer halo.

Direct comparison with kinematic results from observations is difficult, since only the velocity component along the LOS to the observer is accessible and not the radial velocity with respect to the centre of the MW. However, Moss et al. (2013) do provide velocities transformed into the frame of reference of the galactic centre, and we note that $\sim$41 per cent of clouds in their catalog are outflowing with respect to the centre of the MW.

The rotational velocities of the clouds in the lower-right panel of Fig. 5 are normalized by the virial velocity of the corresponding halo. In the inner halo, according to TNG50, a smaller fraction of cold clouds are rotating at velocities in excess of the virial velocity, as compared to the other two regions of the halo. This is consistent with the previous panel that shows that a larger fraction of clouds in the inner halo have higher radial velocities: their motion is likely dominated by feedback mechanisms, and less by gravity. Cold clouds at larger distances are more likely to be rotating at supervirial velocities, as compared to the inner halo. In all regions of the halo, clouds with subvirial rotational velocities are dominant. However, no clear monotonic trend with respect to distance is present at small values of rotational velocities.

While all panels except the bottom-right remain qualitatively unaffected if the threshold on minimum number of cells per cloud is reduced (not shown explicitly), considering clouds with less than 10 member cells tilts the balance between subvirial and supervirial rotation: to demonstrate this, in the solid grey curve, we show the median relation for clouds in the inner halo had all clouds down to one Voronoi cell per cloud been considered. Clearly, the (relative) fraction of clouds with subvirial rotation is lower in comparison to supervirial counterparts. Similar trends exist for the central and outer regions of the halo as well (not shown). Higher resolution simulations are required to determine if this behaviour is a consequence of limited resolution, or if clouds of lower masses (i.e $\lesssim 10^6$ $M_{\odot}$) have physically different kinematics and rotation.

3.4 Relationship between kinematics and physical properties

Kinematics and physical cloud properties are likely related. Observational studies have noted a correlation between metallicity of clouds and their radial velocities (e.g. Richter et al. 2001), such that metal-rich clouds are preferentially outflowing, and thus likely originate from the galaxy. On the other hand, metal-poor clouds are instead preferentially associated with inflows, and so are likely of non-galactic origin and possibly tracing relatively less enriched cosmological gas accretion (Nelson et al. 2013).

In Fig. 6, we explore these correlations for TNG50 clouds. In both panels, we stack all clouds across our sample of MW-like galaxies and split them into three bins based on distance: inner halo ($0.15 < r/R_{200,c} \leq 0.4$), central halo ($0.4 < r/R_{200,c} \leq 0.7$), and outer halo ($0.7 < r/R_{200,c} \leq 1.0$). Median values for each of these bins are shown through solid curves, and coloured by the mean distance corresponding to each bin, while shaded regions show the 16th–84th and 5th–95th percentile regions for the entire sample.

The left-hand panel shows the radial velocities of clouds as a function of their (mean) metallicities. A monotonic trend is seen, wherein radial velocities increase with increasing metallicity. Interestingly, all three medians cross a value of 0 km s^{-1} as they reach supersolar metallicity. Medians corresponding to larger distances are more steep in the positive radial velocity regime and are less steep in the negative radial velocity regime. In the right-hand panel, we show the rotational velocity of clouds (normalized by the virial velocity) as a function of cloud metallicity. All three median curves monotonically shift to higher metallicites at lower rotational velocities. That is, cold clouds with more radial and less rotational motions are more metal enriched. TNG50 therefore suggests that highly enriched cold clouds are dominated by outflows and feedback-driven motions rather than

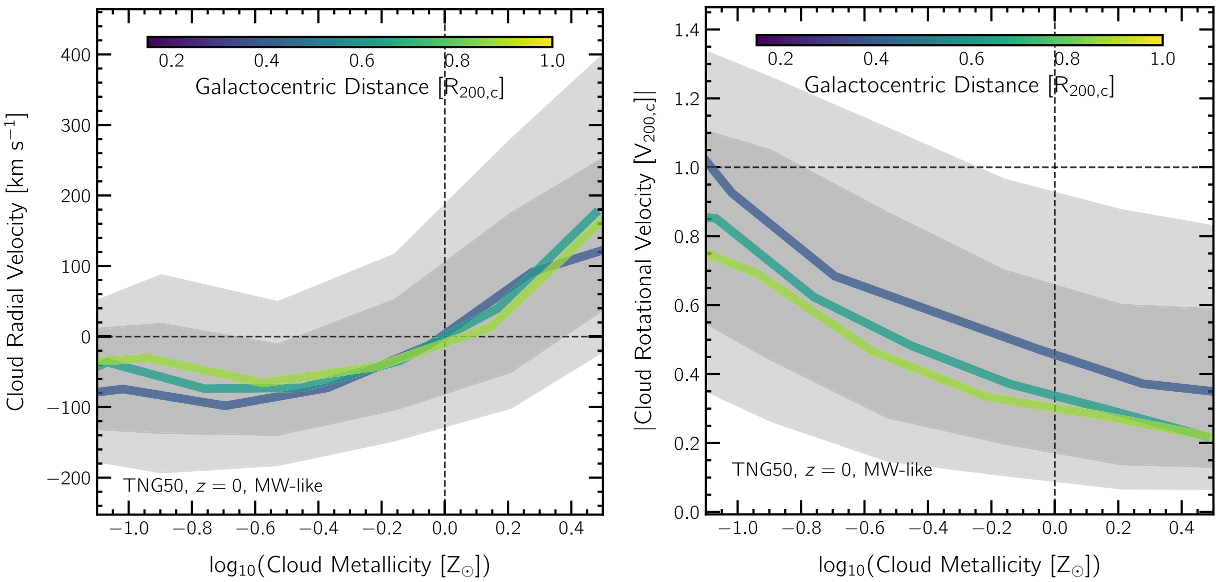

Figure 6. Relation between kinematics of clouds and metallicity in the CGM of TNG50 MW-like galaxies. In the left-hand (right-hand) panel, we show the radial (rotational) velocity of clouds as a function of their metallicity. In both panels, the solid curves represent median values for three different mass bins, while the shaded regions correspond to percentile regions of the entire sample. On average, clouds with subsolar metallicity are seen to be inflowing, with more pristine clouds inflowing at larger velocities, while clouds with superclouds metallicity are preferentially outflowing.

having gravitationally induced dynamics only, while the opposite holds for their more pristine counterparts.

3.5 Spatial distribution of clouds throughout MW-like haloes

Having quantified the basic properties of CGM cold clouds in TNG50 MW-like galaxies, we now turn to the distribution of clouds throughout the haloes and how their properties with respect to their local surroundings depend on galactocentric distance.

In Fig. 7, we show the distributions of the positions of cold clouds: the main panel focuses on the 3D radial distance, while the two insets show distributions of longitude (l) and latitude (b). PDFs of individual MW-like haloes are shown as thin curves, coloured by the mass of the SMBH at the centre of each galaxy. The median of all these PDFs is shown with the thick black curve. In the main panel, we also show medians or MW-like galaxies with over and undermassive SMBHs (two extreme octiles, thick yellow and purple curves).

Cold clouds within the inner boundary of the CGM, i.e. $< 0.15 \times \mathrm{R}_{200c}$, are rare – most cold gas mass in those regions makes up the galactic disc itself, which is excluded by construction by our cold cloud finding algorithm (Section 2.3). On average across the galaxy sample, TNG50 clouds are most commonly found around $0.25 - 0.4 \times \mathrm{R}_{200c}$ and become progressively more rare at larger distances, as already noted for the galaxy of Fig. 2.

Individual-galaxy curves are clearly 'noisy', not as a result of low-number statistics but rather because clouds cluster in specific locations. In addition, a strong trend is present with respect to SMBH mass: galaxies with less massive SMBHs at their centres have a larger fraction of their clouds at smaller distances, while the inverse is true for galaxies with more massive SMBHs. Similar to the trend with respect to sSFR in Fig. 3, we speculate this could be a result of kinetic winds of more massive SMBHs, which (i) drive (cold) gas away from the centre (e.g. Zinger et al. 2020; Ayromlou, Nelson & Pillepich 2022) and/or (ii) destroy clouds or inhibit their formation at small galactocentric distances, either directly via hydrodynamical interactions or indirectly by heating. HVCs within the real MW are thought to have a similar distance trend, with clouds reducing in number towards outer regions of the halo (Olano 2008, although this inference requires theoretical distance modelling).

The left inset of Fig. 7 shows the distribution of cold clouds as a function of galactic longitude. As before, individual curves depicting individual galaxies show a large degree of diversity, primarily due to cold clouds being grouped around each other at specific longitudes. Such a feature is also seen in observations of the MW halo: for instance, in the all-sky NHI map of Westmeier (2018), clouds are not uniformly distributed in longitude but rather are more concentrated at a subset of longitude ranges.

The inset on the right similarly shows the distribution of clouds as a function of galactic latitude. We see a preference for a greater abundance of clouds at low latitudes, indicating alignment with the plane of the gaseous galactic disk. Specifically, [50, 75, and 90] per cent of all cold clouds are within latitudes of [30, 49, and 66] degrees. A similar trend has also been observed in the MW halo, with a strong concentration of HVCs close to the galactic plane (Moss et al. 2013).

3.5.1 *TNG50 cold clouds in relation to their ambient CGM*

In Fig. 8, we show how cold cloud properties compare to their local background environment, and how they vary with distance. Since properties of gas, in general, vary with distance throughout the halo of TNG50 MW-like galaxies (Ramesh et al. 2023), we focus on differences of properties of clouds with respect to the properties of the local surrounding: these are labelled as $\Delta_{\mathrm{cl,bk}}$, i.e. they are the (mean) value of a given property of a cloud minus the (mean) property of the surrounding. For the background, we consider two different definitions of layers of gas around each cloud[4]: (i) the layer immediately

[4]Each is identified geometrically using the connectivity of the Voronoi gas tessellation and the naturally connected Voronoi neighbours of cloud cells. Our background layers therefore have a one-cell thickness roughly equivalent to the numerical resolution and reflect the irregular shapes of clouds.

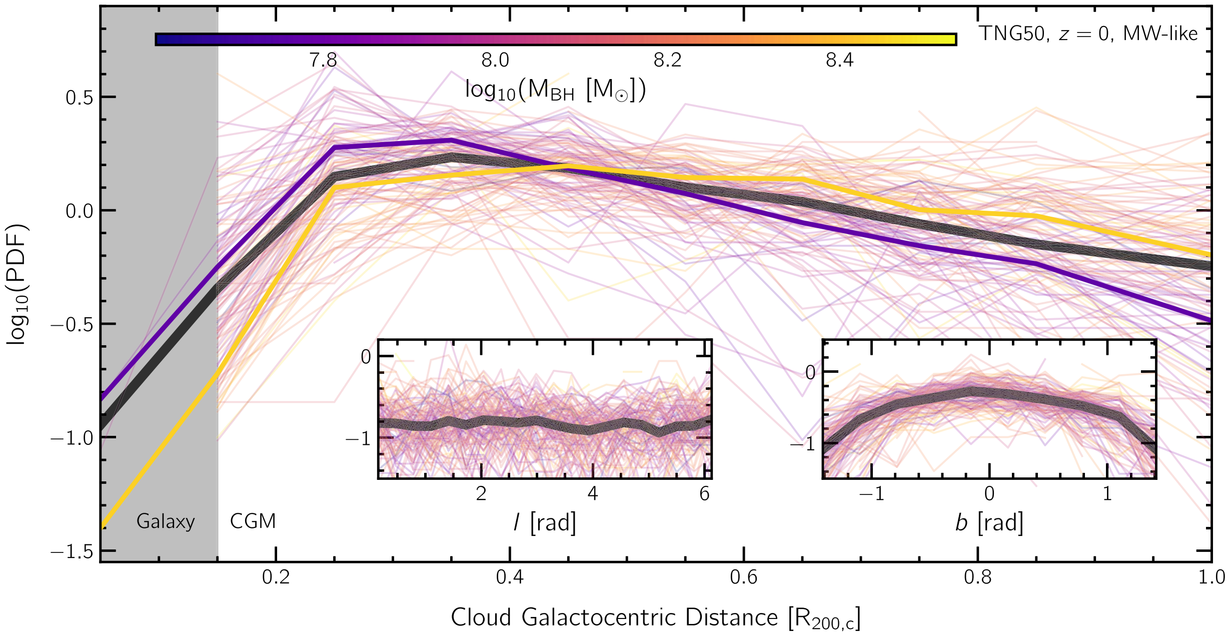

Figure 7. The spatial distributions of clouds in the CGM of TNG50 MW-like galaxies. The main panel shows distributions of the galactocentric distance of clouds (normalized by the virial radius of the halo, $R_{200,c}$). Cold clouds are found throughout the CGM, particularly at intermediate radii. Individual galaxies are shown with thin curves, coloured by the mass of the central SMBH, while the median of the sample is shown in the thicker black curve. In the main panel, we also show medians of the two extreme octile regions. A trend with SMBH mass is present, wherein galaxies with less massive SMBHs have a greater number of clouds at smaller galactocentric distances. The left (right) inset shows the distribution of longitude (latitude) of the cloud population. While the medians in each case are more or less smooth, the individual curves are noisy, suggesting that clouds are clustered in specific locations.

surrounding the cloud, i.e. the 'intermediate layer', shown throughout in black and (ii) the next more outwards layer, which surrounds the intermediate layer, i.e. the true 'background layer', shown in red. As in previous figures, in each panel, we construct PDFs for individual galaxies and show their median but split into three different distance regimes: inner halo (solid curves; $0.15 < r/R_{200,c} \leq 0.4$), central halo (dashed curves; $0.4 < r/R_{200,c} \leq 0.7$), and outer halo (dotted curves; $0.7 < r/R_{200,c} \leq 1.0$).

In the upper-left panel of Fig. 8, we show distributions of temperature contrasts between clouds and their surroundings. At all distances, clouds are cooler than both their intermediate and background layers, which is expected since our cloud-finding algorithm specifically selects cold gas to be part of clouds. A distance trend in values of temperature-contrasts is present: in the inner halo, on average, clouds are cooler than their intermediate layer by ~1.4 dex, which reduces to ~1.1 dex in the central halo, and further to ~0.8 dex in the outer halo. Contrasts with respect to the background layer are more pronounced: in all three regions of the halo, these distributions are offset by roughly −0.3 dex with respect to the distributions corresponding to the intermediate layer. This is consistent with earlier studies of TNG50 cold clouds (Nelson et al. 2020) and Fig. 1, wherein an intermediate mixing layer of warm gas is believed to be sandwiched between cold clouds and their hot backgrounds.

The upper-right panel shows a fundamental quantity of cold clouds in hot media: the density contrast, often written as $\chi = n_{\rm cold}/n_{\rm hot}$ (e.g. Scannapieco & Brüggen 2015). At all distances, TNG50 cold clouds are denser than both their surrounding intermediate and background layers. A halocentric distance trend is present, wherein the overdensity of clouds with respect to their surroundings is greater at smaller galactocentric distances: with respect to their intermediate layers, cold clouds are denser than the surroundings by ~0.65 dex in the inner halo and by ~0.55 and 0.45 dex in the central and outer regions of the halo, respectively. Density contrasts with respect to the background layers are larger: ~0.95, 0.9, and 0.85 dex in the inner, central, and outer halo, respectively. However, at all distances, distributions corresponding to contrasts with respect to the background are again broader versus the intermediate layer. These findings are relevant in the context of previous, idealized numerical experiments. In fact, wind tunnel simulations generally assume high values of density contrasts, typically from ~10 to ~1000 (e.g. Schneider & Robertson 2017). Within the CGM of TNG50 MW-like galaxies, we find such large density contrasts rarely: they make up the tail of the distribution only. Our result suggests that typical cold clouds in MW-like haloes may have much weaker overdensities with respect to the background CGM, which would impact mixing efficiencies and overall survivability.

In the centre-left panel of Fig. 8, we show the relationship between cloud total pressure and background total pressure, i.e. summing magnetic plus thermal pressure. On average, clouds at all distances are slightly underpressurized with respect to their surroundings, with a greater contrast with respect to the background layer in comparison to the intermediate layer. A trend with distance is apparent, whereby distributions corresponding to smaller distances are skewed towards more negative values. Similar trends are observed in the contrast distributions with respect to the background layer, albeit at slightly more enhanced values. In addition to a horizontal offset of the peaks, the distributions corresponding to the background layers are broader than their intermediate layer counterparts at all distances.

Although clouds, on average, are only weakly underpressurized with respect to their surroundings when total pressure is considered; the pressure contrast is more striking when only the thermal component is considered. To illustrate this, we show the distributions of

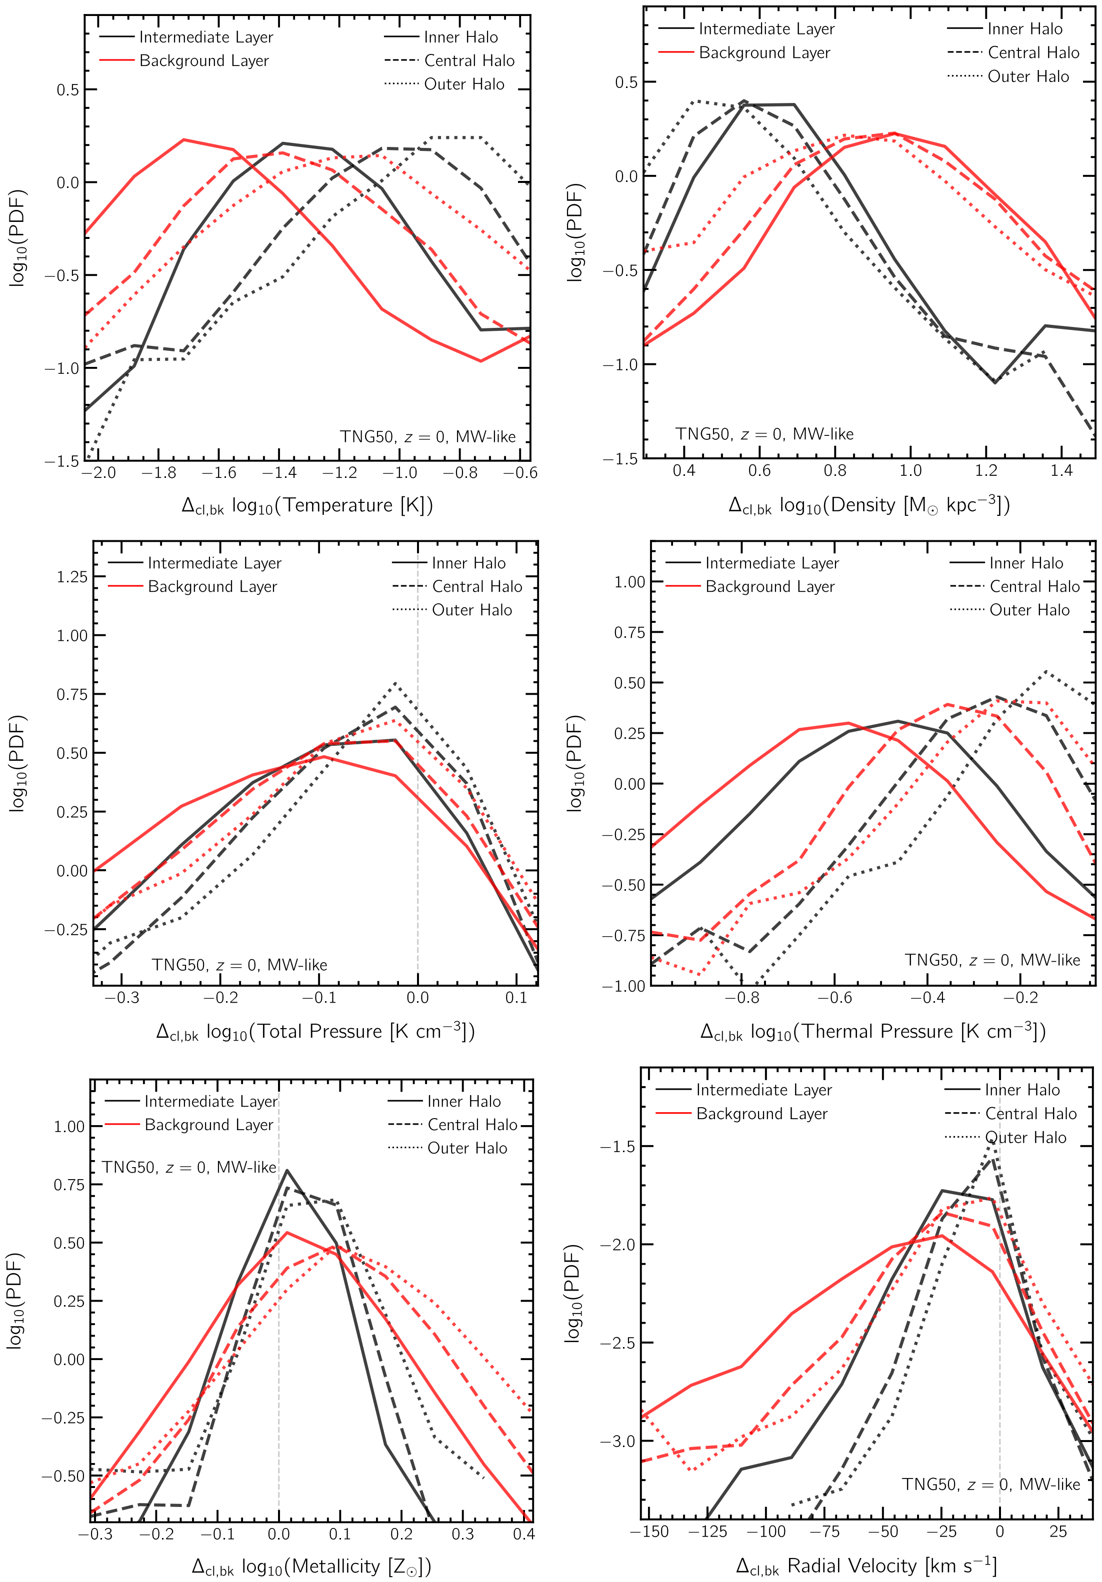

Figure 8. Distributions of contrast in properties of TNG50 MW-like clouds with respect to their intermediate and background layers, i.e. with respect to the local, ambient CGM gas. In each panel, distributions of contrasts corresponding to the former are shown in black, while the latter are shown in red. We show median PDFs for three distance regimes: inner halo (solid curves; $0.15 < r/R_{200,c} \leq 0.4$), central halo (dashed; $0.4 < r/R_{200,c} \leq 0.7$), and outer halo (dotted; $0.7 < r/R_{200,c} \leq 1.0$). Contrasts of temperature are shown in the top-left panel, density in the top-right, total pressure in the centre-left, thermal pressure in the centre-right, metallicity in the bottom-left, and radial velocity in the bottom-right. In all cases, contrasts are more enhanced with respect to the background as compared to the intermediate layer, both in terms of the median and the width of the distributions.

thermal pressure contrasts between clouds and their surroundings in the centre-right panel of Fig. 8. The median contrasts are much larger, peaking at roughly −0.45 and −0.6 dex with respect to the intermediate and background layers, respectively, in the inner halo, and with a similar distance trend as the previous panel. It is clear that magnetic pressure is an important component for these clouds, modifying cloud–wind interaction, and without which they would need to contract further to reach pressure equilibrium (Li et al. 2020; Nelson et al. 2020; Sparre, Pfrommer & Ehlert 2020).

In the lower-left panel of Fig. 8, we show PDFs of cloud-background metallicity contrast. On average, clouds are slightly more enriched than both their intermediate and background layers, at all distances. This may naturally arise since gas with metals cools faster (i.e. through metal line cooling), and is hence more likely to be cold. For both the intermediate and background layers, a very weak trend of this metallicity contrast with distance is present: in the inner halo, the median difference of cloud metallicity with respect to both the surrounding layers is overall negligible (∼0.01 dex) whereas it is somewhat larger in the central and outer halo (by ∼0.04 and 0.07 dex). Additionally, at all distances, the distributions of contrasts with respect to the background are more broad in comparison to contrasts with respect to the intermediate layer, i.e. clouds are more homogeneous with their immediate surroundings than with gas further away, as one would intuitively expect. All distributions are further skewed towards positive contrasts, i.e. clouds are much more likely to be overmetallic with respect to their surroundings than otherwise.

Finally, the lower-right panel of Fig. 8 shows the kinematic connection between cold clouds and the ambient media. In particular, we show the distribution of cloud-background relative radial velocity. On average, clouds at all distances are (weakly) inflowing with respect to their surrounding gas layers. Distance trends are again noticeable, with clouds at smaller galactocentric distances inflowing faster than their surroundings as compared to more distant clouds. Radial velocity contrasts with respect to the intermediate layer peak at few km s^{-1} in all three regions of the halo, but the widths of these distributions are large: tens of km s^{-1} and are greater at smaller galactocentric distances. With respect to the background layer, a distance trend is noticeable in both the median and width of distributions: a median contrast of roughly −25 km s^{-1} in the inner and central regions of the halo reduces to roughly −5 km s^{-1} in the outer halo, in addition to widths of distributions being larger at smaller galactocentric distances.

Such velocity differences are a combined result of inflow/outflow motion of the cold cloud and the inflow/outflow of the background. There are four different possibilities:

(i) Cloud outflowing and intermediate layer (background layer) outflowing: This accounts for ∼26 per cent of all cases. The median contrast of radial velocity in this case is roughly −7.2 (−15.4) km s^{-1}, i.e. when both clouds and the surroundings are outflowing, clouds are outflowing slower with respect to the surrounding. This suggests ongoing acceleration.

(ii) Cloud inflowing and intermediate layer (background layer) outflowing: This accounts for ∼4 per cent (11 per cent) of all cases. The median contrast of radial velocity here is approximately −24.7 (−60.6) km s^{-1}, i.e when clouds are inflowing amidst outflowing surrounding gas, there is a large velocity contrast present. These relatively rare cases are likely accreting clouds that are being hit by outflowing galactic-scale winds.

(iii) Cloud inflowing and intermediate layer (background layer) inflowing: This accounts for ∼69 per cent (62 per cent) of all cases. The median value of radial velocity contrast in this case is roughly −10.9 (−22.9) km s^{-1}, i.e. when both clouds and the surrounding are inflowing, clouds are inflowing faster than their surroundings. This suggests ballistic acceleration enabled by cloud overdensities $\chi > 1$.

(iv) Cloud outflowing and intermediate layer (background layer) inflowing: This accounts for ∼0.5 per cent (1.1 per cent) of all cases. The median radial velocity contrast here is ∼10.8 (21.5) km s^{-1}, and this is the only case in which clouds, on average, are outflowing faster than their surroundings. These rare cases reflect outflow-driven, or *in situ* outflow formed clouds, which are no longer comoving with any bulk outflow and so can interact with the ambient CGM.

The distribution of radial velocity contrasts can thus be summarized as follows: (a) Most clouds ($\gtrsim$90 per cent) are flowing in the same direction as their surroundings and with typically small velocity contrasts (few km s^{-1}). This reflects the peak of the radial velocity contrast distributions; (b) Clouds flowing in the opposite direction with respect to their surroundings typically have much larger velocity contrasts (few tens of km s^{-1}) and populate the tails of the radial velocity contrast distributions. Although this needs to be investigated in more detail, we speculate that the latter is less common because cold clouds are possibly destroyed as a result of enhanced instabilities when there is a large velocity contrast with respect to the surrounding, for example, as a result of increased mass loss driven by the KH instability (e.g. Sander & Hensler 2021). However, if the cold clouds that we identify are pre-dominantly related to cold gas stripping of satellites, the former case would naturally be more common since cold gas that is stripped falls with roughly the same velocity as the satellite shortly after being stripped.

3.6 Spatial clustering of CGM cold clouds

As we have seen (Fig. 2), cold gas clouds in the CGM of TNG50 MW-like galaxies are not distributed uniformly throughout the halo but rather appear to be clustered in certain locations, often near satellite galaxies or tails of their stripped gas.

We quantify this phenomenon in Fig. 9, where we study the minimum distance between a cold cloud and its nearest satellite.[5] The left-hand panel shows PDFs of the distance of clouds to their nearest satellite galaxy. As before, we construct individual PDFs for each halo and then the median across haloes.

We consider three types of satellites based on their (current) baryon fractions, i.e. ratio of baryon mass to total subhalo mass: >10 per cent, 1–10 per cent, and 0 per cent in black solid, dashed, and dotted curves, respectively. Since the first category of satellites are typically less frequent (3–4 on average per halo) in comparison to the latter two categories, we down sample the number of the latter two to avoid any biases. We do so through the following procedure: if a halo contains N satellites with baryon fraction >10 per cent, we randomly select N satellites of the latter two categories from the available pool. To make sure that this random-picking does not bias our results, we repeat the process 100 times for each galaxy.

[5]Note that clouds related to gas that has been stripped in the distant past would naturally end up with relatively large distance values, despite physically originating from the corresponding satellite. In the future, we will overcome this limitation by using the Monte Carlo tracers to first identify the associated satellite, and then estimate the closest distance to the past trajectory of that satellite.

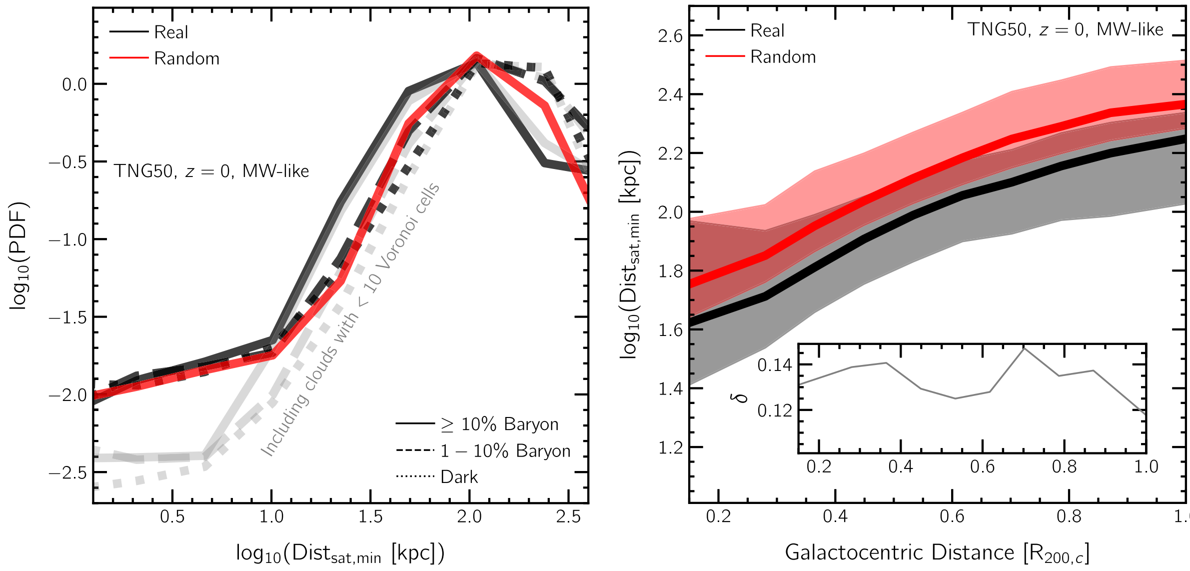

Figure 9. The clustering of cold clouds around satellite galaxies in the CGM of TNG50 MW-like galaxies. In the left-hand panel, we show the distribution of distances between clouds and their nearest satellite (black lines). We split each PDF into different categories based on the (current) baryon fractions of those satellites: >10 per cent, 1–10 per cent, and 0 per cent in solid, dashed, and dotted curves, respectively. As a comparison point, we show the expected distribution after randomizing the position of the nearest satellite (red curve, see text). Clouds are more frequently found close to satellites than would be expected in the random case, this being the case at least for satellites with higher baryon fractions. The right-hand panel shows the relation between the distance to the nearest satellite with a baryon fraction >10 per cent (y-axis) and the galactocentric distance of the cloud (x-axis), for both the real and the random case. The solid curves show the median, while the shaded bands correspond to the 16$^{\text{th}}$ and 84$^{\text{th}}$ percentile regions of the stacked sample of clouds across all 132 MW-like galaxies. The inset shows the difference between the two median curves, i.e. random minus real. At all distances, on average, the correlation between positions of clouds and their nearest satellites is stronger than what is expected if their relative positions were random.

Overall, we find that cold clouds tend to be closer to baryon rich satellites than either dark satellites or (relatively) baryon poor satellites, suggesting a physical i.e. origin link. A vertical offset in the PDFs is seen at $\lesssim$100 kpc as one transitions towards lower values of baryon fractions of satellites. Clouds are thus more likely to lie close to satellites with larger baryon fractions.

A similar observation of cloud-satellite clustering exists in the MW halo, where there is a high concentration of HVCs around the Magellanic Stream and so also close to the Large and Small Magellanic Clouds (Moss et al. 2013). Idealized simulations of gas-rich dwarf galaxy stripping also show similar signatures (Mayer et al. 2006).

To confirm that this correlation between positions of cold clouds and satellites with large baryon fractions is robust, we carry out a random shuffling experiment as follows. Once the nearest satellite to a particular cloud is identified, we randomize the position of that satellite and re-compute the minimum distance between the cloud and the randomized position of the satellite. To avoid any biases, we repeat this procedure 100 times. If the position of a cloud with the associated satellite was truly random, this procedure would have minimal effect. However, if the connection between a cloud and a satellite is ‘real’, the randomization would erase any signature. We show the median PDF corresponding to the random case in red. A clear offset is once again visible, suggesting that a physical correlation is indeed present.

In the right-hand panel of Fig. 9, we show how the distance to the nearest baryon-rich satellite depends on the galactocentric distance of the cloud. The solid lines show the median, while the shaded regions correspond to the 16$^{\text{th}}$ and 84$^{\text{th}}$ percentile regions. In black, we show the trend corresponding to the case emerging from the simulation (‘real’ case). A median distance to the nearest baryon-rich satellite of $\sim$40 kpc at the inner boundary of the CGM (i.e. 0.15R$_{200c}$) increases to $\sim$80 kpc at 0.5R$_{200c}$ and further to $\sim$180 kpc at the virial radius. The width of the percentile regions simultaneously decreases, from $\sim$0.5 dex at 0.15R$_{200c}$ to $\sim$0.3 dex at the virial radius.

Intuitively, we expect a similar qualitative trend even for clouds/satellites distributed randomly since shells (of equal width) at smaller galactocentric distances have smaller volumes. To normalize out this volume effect, we also include the corresponding relations for the random case. As expected, a similar trend with distance is seen, although the red median is clearly offset vertically above the black median curve. The difference between these two medians is shown in the inset, which estimates the true strength of the radial dependence of clustering, i.e. with the volume-scaling effects removed. The offset is rather independent of distance, varying between $\sim$0.12 and 0.15 dex. Thus, at all distances, a weak overcorrelation is seen between the positions of clouds and >10 per cent baryon fraction satellites with respect to the random case.

While the main results of this panel are largely unaffected by the lower-limit threshold for the minimum number of cells per cloud, we mention a subtle difference that is present for the case where clouds with less than 10 member cells are included: the corresponding distributions are shown in grey curves. While the grey and black distributions merge for Dist$_{\text{sat, min}}$ values larger than $\sim$10 kpc, a vertical offset is present in the solid and dashed curves with respect to the dotted one at smaller distances. Although the difference is small, we suspect that it could correspond to gas that has been freshly stripped and is hence present in the distributions corresponding to baryonic satellites but absent in the dark case. While this could simply be a result of poor resolution, it may also be the case that these tiny (unresolved) clouds act as seeds of dense, cold gas that trigger thermal instability, eventually giving rise to larger clouds as

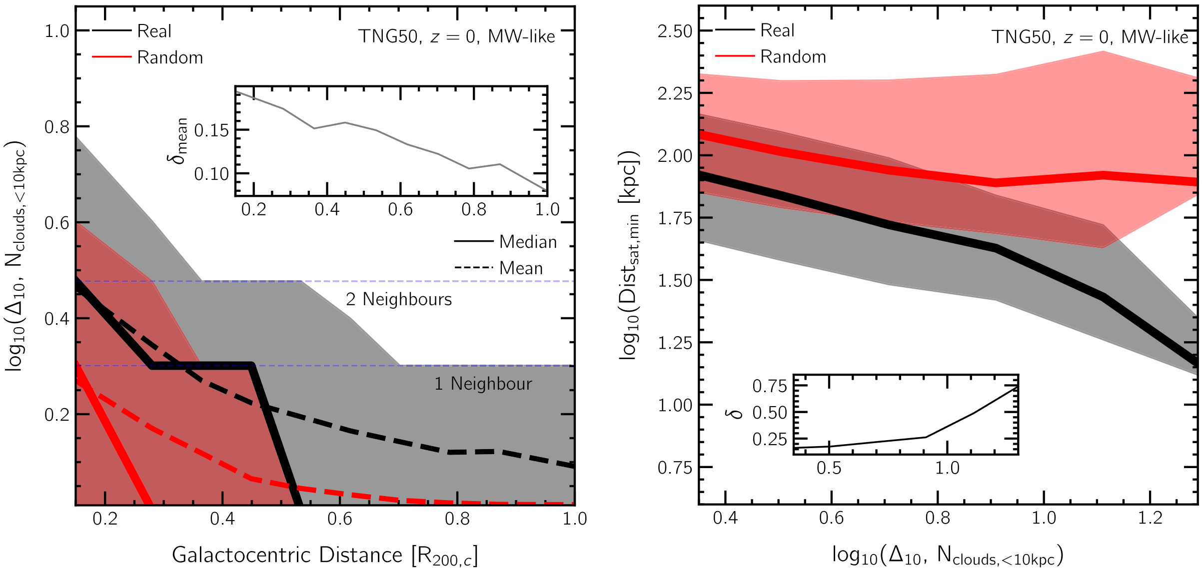

Figure 10. The spatial clustering of cold clouds around other cold clouds in the CGM of TNG50 MW-like galaxies. The left-hand panel shows the trend between Δ_{10}, the number of clouds within spheres of radius 10 kpc centred on each identified cloud, and galactocentric distance. Solid curves correspond to median values, dashed curves to mean values, and shaded regions to 16th and 84th percentiles. Black curves show the actual outcome of the simulation, whereas the red curves show what a random (shuffled) spatial distribution would look like. The inset shows the difference between the two mean values. In the right-hand panel, we show the trend between Δ_{10} and the minimum distance of clouds to the nearest satellite with baryon fraction >10 per cent, for both the real case (black), and the case where satellite positions are randomized (red). The inset shows the difference between the two medians. Clouds are typically overclustered around other clouds with respect to what a random scenario would predict, with stronger overclustering at smaller galactocentric distances. Higher values of Δ_{10} are seen when clouds lie closer to satellite galaxies with 'high' baryon fractions.

gas condenses around them (Nelson et al. 2020; Dutta, Sharma & Nelson 2022).

We next turn to the possibility that cold clouds may be clustered around other cold clouds rather than being distributed uniformly throughout the halo. It is believed that such clustering can increase the longevity of clouds through the process of drafting (Williams & Shelton 2022). We quantify this clustering through the Δ_{10} parameter, which we define as follows: for every identified cloud, Δ_{10} is the number of clouds that lie within a sphere of radius 10 kpc, including itself. A value of one thus implies that there are no neighbouring clouds within this sphere, a value of two corresponds to one neighbour, and so on. As extended clouds are much less frequent than smaller ones in this statistical approach, we neglect the issue that arises because clouds are in fact extended objects.

The left-hand panel of Fig. 10 shows the trend of Δ_{10} with galactocentric distance. Solid curves correspond to median values, dashed curves to mean values, and shaded regions to 16th and 84th percentiles. We begin by discussing the black curves, which represent the actual outcome from the simulation. A median value of two neighbours (i.e. $\log_{10}(\Delta_{10}) = \log_{10}(3)$) in the innermost regions of the halo ($\lesssim 0.3R_{200c}$) reduces to one neighbour (i.e. $\log_{10}(\Delta_{10}) = \log_{10}(2)$) between $0.25R_{200c}$ and $0.5R_{200c}$ and to zero neighbours at farther distances. The mean, however, does not portray such a 'step-like' behaviour and shows a steady monotonic decrease of $\log_{10}(\Delta_{10})$ with distance, reducing from $\sim$0.5 at $0.15R_{200c}$ to $\sim$0.1 at the virial radius, i.e. at all distances, the mean number of neighbours is more than zero and is greater at smaller galactocentric distances.

As before, such a qualitative trend with distance is expected even for a random distribution of clouds due to available volume decreasing towards the halo centre. To remove this effect, we randomize the positions of clouds while keeping their radial number density profile fixed and re-calculate Δ_{10} for the randomized positions. The corresponding trend is shown in red. If the positions of clouds with respect to other clouds was already random, this procedure would not have a significant impact. However, in case clouds are truly clustered around their neighbours, a difference would emerge. Indeed, in both the mean and the median, the randomized scenario shows a smaller value of Δ_{10} than the true case at all distances. The inset shows the difference between the two mean values (δ_{mean}), characterizing the strength of the true radial dependence of clustering. Overclustering with respect to the random case is strongest at smaller galactocentric distances: a δ_{mean} of $\sim$0.19 at $0.15R_{200c}$ reduces to $\sim$0.15 at $0.5R_{200c}$, before dropping steeply to $\sim$0.08 at the virial radius.

While we use the Δ_{10} metric to study clustering, other statistics are equally well suited. For example, the two-point correlation function, or e.g. the distance to the 10th nearest cloud, as a measure of overclustering. We have considered both, and they provide qualitatively similar results, demonstrating an overclustering of cold clouds with respect to the random scenario.

In the right-hand panel of Fig. 10, we show that Δ_{10} is linked to the minimum distance to the nearest satellite with a baryon fraction of >10 per cent. Median values are shown with solid curves, and 16th and 84th percentiles with shaded regions. The black curve shows the true signal, while the red shows a test where the position of satellites are randomized, as discussed above. The black median curve shows a sharp monotonic drop with increasing Δ_{10}: when a cloud has only one neighbour within 10 kpc, the average distance to the nearest baryon-rich satellite is $\sim$100 kpc. When a cloud instead has $\gtrsim$10 neighbours in close proximity, the median distance to the nearest baryon-rich satellite drops to $\lesssim$20 kpc.

The random case portrays a qualitatively similar trend, albeit with an offset of $\sim$0.2 dex at low values of Δ_{10}, and a shallower drop towards higher values of Δ_{10}. Most importantly, the Δ_{10}-trend of the offset between the two median curves is shown in the inset. An offset of 0.2 dex at $\log_{10}(\Delta 10) \sim 0.3$ rises sharply to 0.75 dex at $\log_{10}(\Delta 10) \sim 1.4$. Thus, when clouds are strongly clustered with neighbouring clouds, they are more likely to lie close to a >10 per cent baryon fraction satellite, as opposed to being randomly positioned with respect to such satellites.

3.7 Comparison with observations of the Milky Way

We conclude our investigation with a number of direct comparisons with observed data of IVCs and HVCs in the MW halo. This important connection is enabled by the cosmological context of the TNG50 MW-like galaxies and is unavailable in single cloud and other idealized numerical simulations.

In what follows, we place a hypothetical observer at a random point in the galactic plane, at a distance of 8.34 kpc away from the galactic centre. This observer is considered to be in perfect circular motion around the galactic centre at a velocity of 240 km s^{-1}. This observer is consistent with the known solar location and motion in the real MW (Reid et al. 2014). Since observations do not enforce a minimum radial distance when identifying clouds, we here relax our lower limit for the inner boundary of the CGM, i.e. we include all clouds present within the virial radius of the halo, barring the one massive cloud that is the galaxy itself.

In the top-left panel of Fig. 11, we show PDFs of the LOS velocity of cold clouds. The different coloured regions signify common definitions used to classify clouds: IVCs are those with (absolute) LOS velocities in the range 40–90 km s^{-1}, HVCs with 90–170 km s^{-1}, and VHVCs with >170 km s^{-1} (e.g. Lehner et al. 2022), although some authors refer to all clouds with (absolute) LOS velocities >90 km s^{-1} as HVCs (e.g. Wakker 2001). The red curves show the PDF of a sample of HVCs of the MW presented in Moss et al. (2013). The solid black curve shows the median PDF of HVCs across the whole TNG50 MW-like sample of simulated galaxies. Overall, the agreement is striking. Both curves indicate that the abundance of clouds is smaller for higher velocity clouds. Note that the TNG50 result is the stacked outcome, averaging across all galaxies of our sample, and these have a diversity of properties including stellar disc lengths (see Pillepich et al. in preparation). In the black dashed line, we show the median PDF of all clouds in TNG50, irrespective of velocity. We predict that the abundance of 'low velocity' clouds with (absolute) LOS velocities <40 km s^{-1} – not generally accessible in observations—is roughly independent of velocity, and that these clouds are more abundant than higher velocity clouds.

To check if the motion of these clouds is dominated by gravity, or not, we compare the velocity distribution of clouds to that of hypothetical test particles whose motion is purely determined by gravity. To do so, for each cloud, we compute the (circular) velocity that is required at that distance for the centrifugal force to perfectly balance the force of gravity. The PDFs of these test particles are shown in grey: those consistent with HVC velocities as solid curves and for all velocities in dashed curves. While the shape of the grey dashed curve is similar to that of the black dashed curve in the velocity range ∼[−90, 90] km s^{-1}, i.e. clouds that are not HVCs (or VHVCs), the two curves diverge at higher velocities. That is, gravity alone cannot account for the high velocity tails, suggesting that further astrophysical processes are relevant for the kinematics of HVCs.

In the top-right panel of Fig. 11, we show the relation between cloud mass and the (3D) distance to the observer. The different scatter points show estimates of IVCs (red) and HVCs (orange) from Wakker (2001) for the MW. Note that most of these points are either lower or upper limits, which we denote with arrows in the relevant direction. Since all these clouds were observed through their HI emission, Wakker (2001) use factors of 1.2 and 1.39 to account for the masses of ionized hydrogen and helium, respectively, to arrive at a better estimate of total mass, from their initially inferred HI mass. However, more recent studies suggest that the ionized component of clouds may account for a larger mass fraction (e.g. Lehner & Howk 2011), and thus the masses quoted by Wakker (2001) are likely underestimated. As is, most clouds in the sample have a mass $\lesssim 10^5 M_\odot$, i.e. below the resolution limit of TNG50.

Although these data points follow the expected $M \propto d^2$ dependence (e.g. Wakker & van Woerden 1997), we suspect this to be largely due to the prevalence of lower limits for the distance estimates of a large fraction of clouds, especially for less massive clouds ($\lesssim 10^5 M_\odot$). This artificially results in too small cloud mass estimates. If one were to instead assume a uniform distribution of distances in the range ∼[5, 12] kpc, as motivated by absorption line measurements that are insensitive to mass of clouds (e.g. Lehner et al. 2012, 2022), a different mass distribution would arise. Indeed, for those clouds more massive than $10^5 M_\odot$ in Wakker (2001)'s sample, where distance estimates are a mix of lower-limits, upper-limits, and tighter constraints, there seems to be no strong dependence of mass with distance, at least within the distance brackets available.

For comparison, we show TNG50 IVCs (HVCs) with grey (black) points, and their median with the grey (black) curve. Consistent with the few observational constraints available, TNG50 predicts a weak dependence on distance for clouds above $\sim 10^6 M_\odot$, and a very similar relation for both IVCs and HVCs, out to distances within which clouds in the MW are typically observed. A noteworthy point from this plot is the dearth of TNG50 clouds at small distances. While this is likely because all cold gas in this region is contiguous with the galactic disc, it is possible that the lack of these 'small-scale fountain flows' is a limitation of the simplified stellar feedback driven galactic-wind model of TNG (Section 2; Springel & Hernquist 2003). Further exploration with alternate, more explicit stellar feedback models (e.g. Smith, Sijacki & Shen 2018; Hopkins et al. 2020; Hu et al. 2022) would be essential to comment on small distance clouds. We demarcate the large region inaccessible to TNG50 with the grey region in the lower right corner. Moving to small cloud masses simply requires higher numerical resolution, while moving closer to the disc-halo interface requires more sophisticated models for ISM and stellar feedback physics.

In the bottom-panel of Fig. 11, we show an all-sky map of gas LOS velocity, including only high-velocity (>90 km s^{-1}) cold gas ($T < 10^{4.5}$K) in emission. We show a single TNG50 MW-like galaxy, the same halo from Fig. 2, as seen by our hypothetical observer. To be as realistic as possible, we include all halo gas within the virial radius, including gas that is gravitationally bound to satellites, since sky maps of the MW halo contain such components. Colours indicate the velocity at which gas is moving along the LOS to the observer.

A gallery of such projections are shown in Fig. 12, where the top-left panel is observational data of the real MW, from the HI4PI survey (Westmeier 2018). The other five panels are five randomly selected galaxies from our TNG50 WW-like sample. The most interesting is the top-right panel, which shows a halo that contains an SMC/LMC-like pair. This map exhibits a degree of qualitative similarity to the true data: along with a noticeably patchy distribution of gas throughout the halo, a Magellanic-like stream at [l, b] ∼ [0, −60] deg is present. Although all the other TNG50 all-sky projection lack

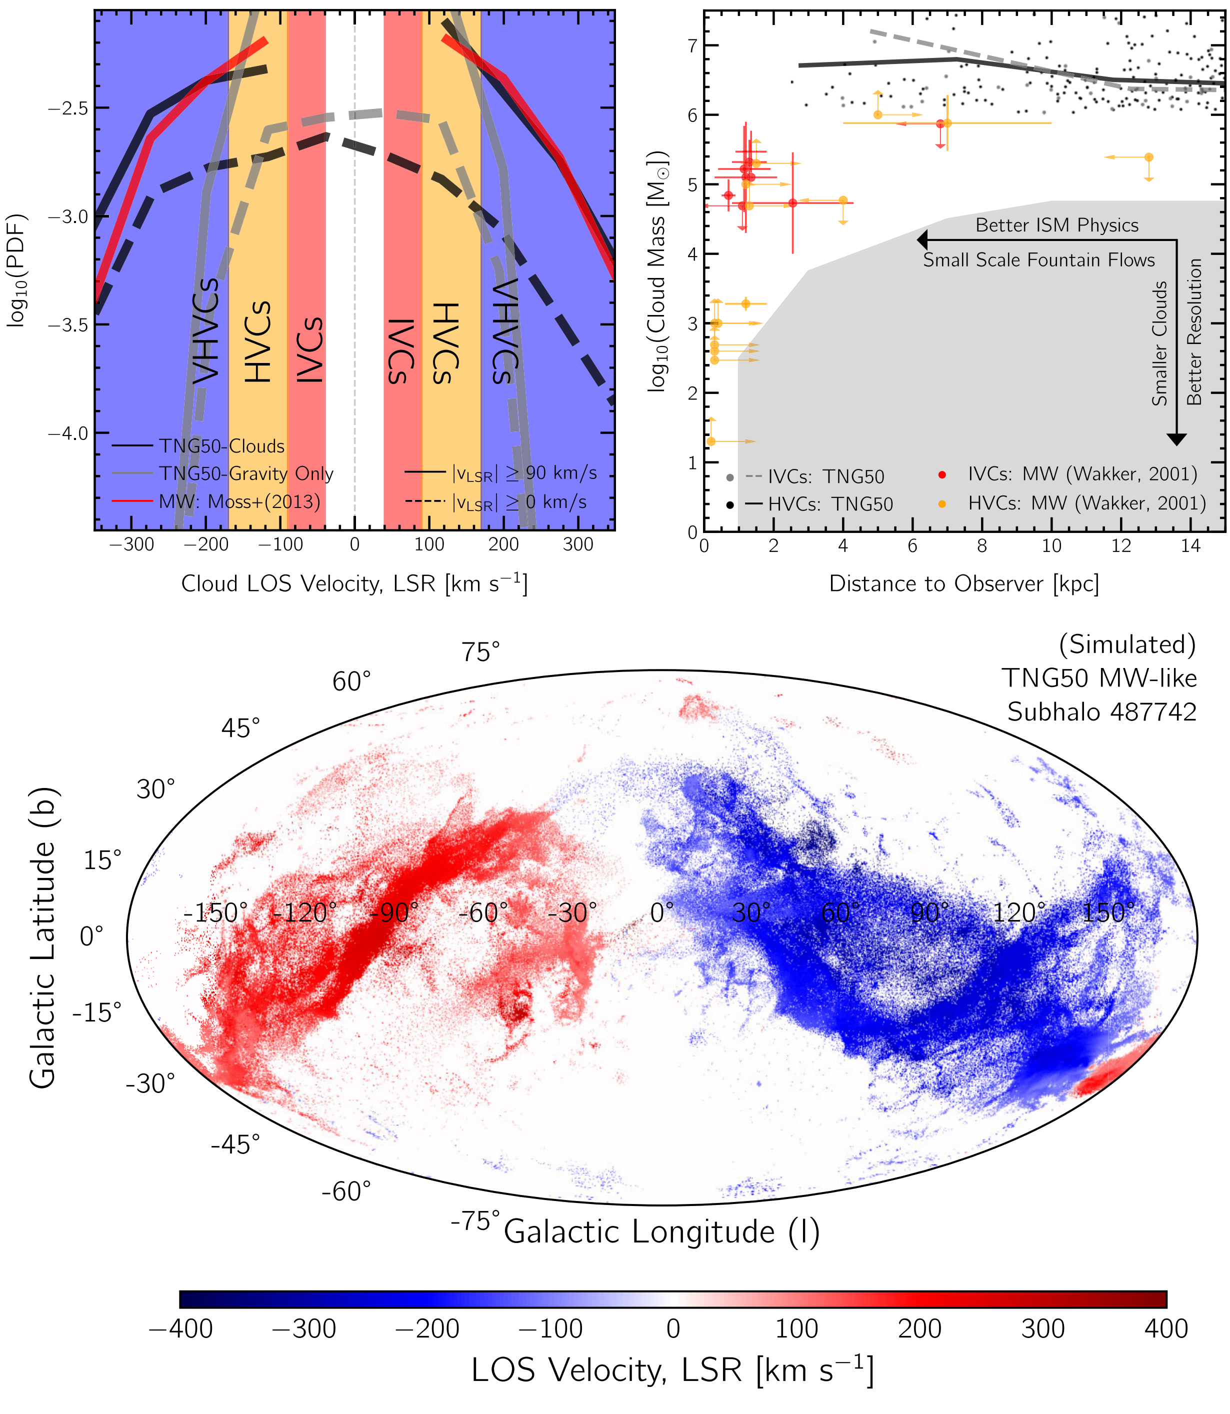

Figure 11. Comparison between the clouds in TNG50 MW-like galaxies with observational data of the real MW. In the top-left panel, we show median PDFs for the LOS velocity in the frame of reference of the LSR. Shaded bands in the background correspond to commonly used definitions for clouds. Red curves correspond to distributions from a catalog of observed HVCs in the MW halo (Moss et al. 2013), while black curves show the equivalent distributions from TNG50 across all 132 MW-like galaxies. We also include the distribution for hypothetical test particles (grey) under purely gravitational motion, demonstrating that the kinematics of HVCs are more complex than gravity alone. In the top-right panel, we show the relation between cloud mass and observer-centric distance. Orange and red points show observational data (Wakker 2001), while black and grey points and median curves are from TNG50, for HVCs and IVCs, respectively. The bottom panel shows an all-sky Aitoff projection of a TNG50 MW-like galaxy (subhalo 487742) from a hypothetical observer at the solar location, with colours corresponding to the LOS velocity of cold gas, in the frame of the LSR. The distribution of neutral hydrogen in the CGM has a complex morphological and kinematic structure.

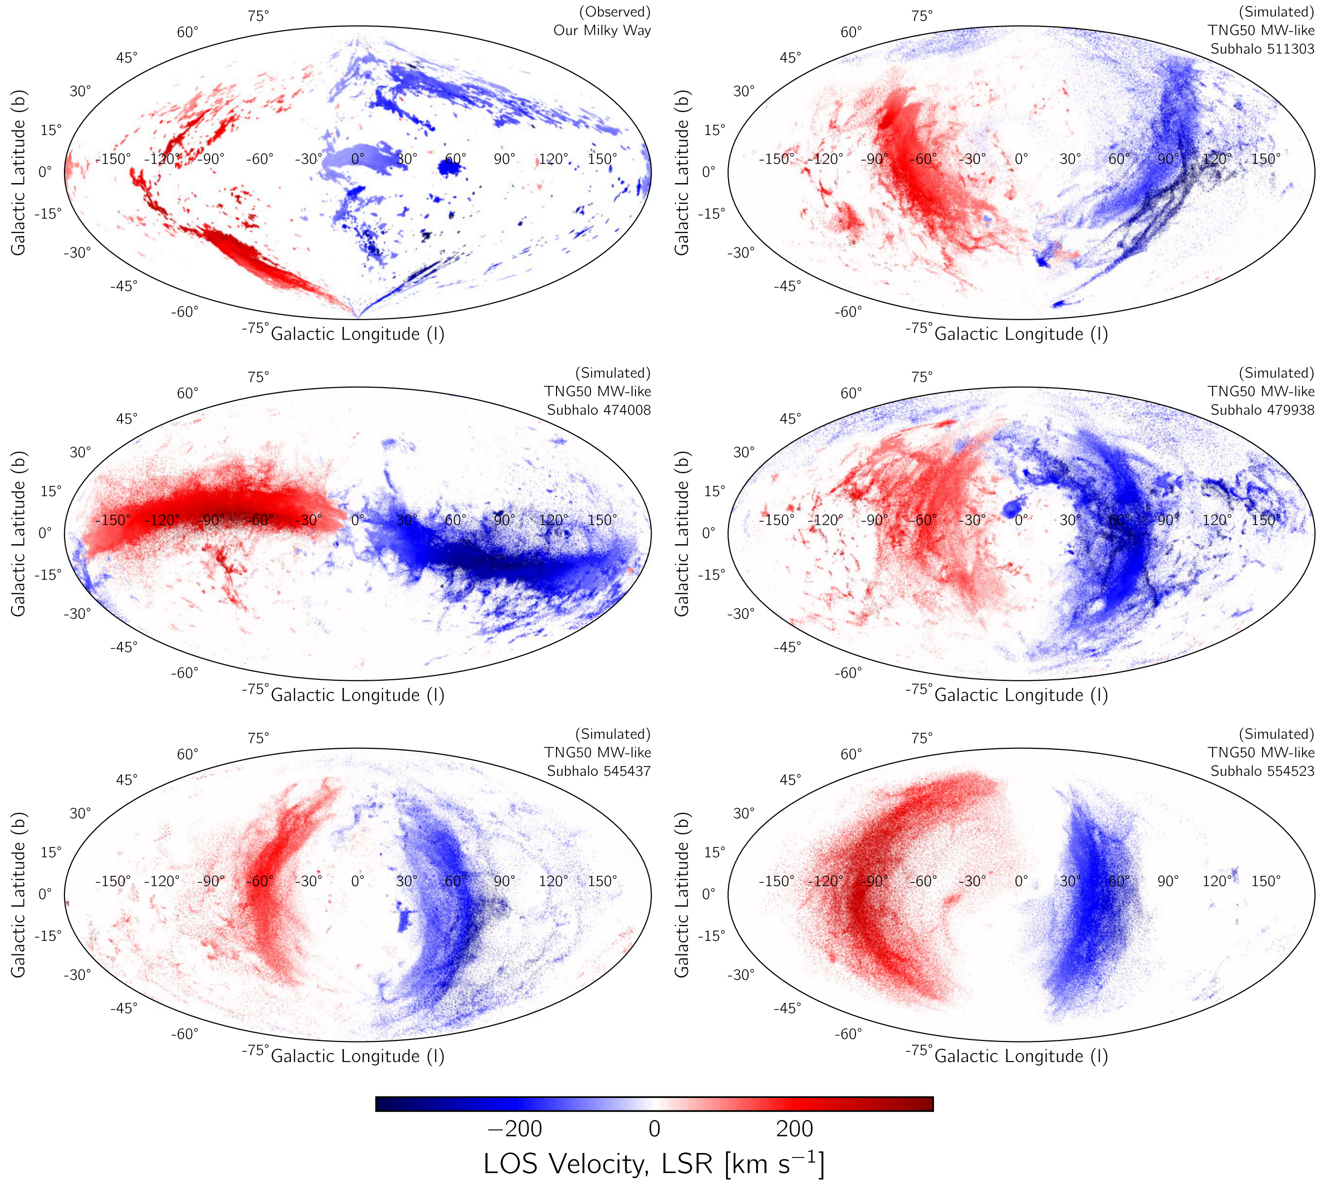

Figure 12. Aitoff projections of cold gas, coloured by LOS velocity, similar to the bottom panel of Fig. 11. In the top-left panel, we show the all-sky projection from the HI4PI survey (Westmeier 2018) for the true MW halo, while all other panels correspond to projections from five randomly selected MW-like galaxies from our TNG50 sample. The most interesting is the top-right panel (subhalo 511303), which contains a SMC/LMC-like pair. Similar to the top-left panel, a Magellanic-like stream around $[l, b] \sim [0, -60]$ deg is present in this case. Overall, the distribution of cold gas through the halo is both spatially and kinematically complex, as well as diverse.

a Magellanic-like stream, a trend in colours is apparent across all these maps, with gas at negative longitudes preferentially outflowing, while gas at positive longitudes is preferentially inflowing. Gas distributions are clearly unique in each TNG50 halo, highlighting the diversity across the sample.

4 SUMMARY AND CONCLUSIONS

In this paper, we have studied the existence, distribution, and physical properties of cold, dense clouds of gas in the circumgalactic medium (CGM) of a sample of 132 MW-like galaxies in the TNG50 simulation at $z = 0$. Our motivation to study such objects stems from the plethora of open questions surrounding HVCs in the real MW halo. TNG50 offers a combination of resolution and volume to begin exploring such clouds in a cosmological context over a wide sample of galaxies, bridging the gap to small-scale, idealized numerical simulations of cold cloud evolution and survival. We summarize our main findings as follows:

(i) MW-like galaxies in TNG50 typically contain of order one hundred (thousand) reasonably (marginally) resolved cold clouds in their gaseous haloes. While the number of clouds shows no significant trend with the stellar mass of the galaxy, the scatter correlates with the sSFR. This suggests that (a) active galatic nucleus feedback quenches star formation and destroys clouds or prevents clouds from forming, and/or (b) the flow of cool gas through the circumgalactic medium (CGM) is physically connected to the fuelling of galactic star formation (Fig. 3).

(ii) Clouds show a large variation in their mass, although most clouds in our sample have a mass close to $\sim 10^6$ M$_\odot$, corresponding to the chosen cloud definition. More massive clouds are larger, with cloud sizes ranging from $\sim$ a few hundred pc to $\sim$ a few tens of kpc. Smaller clouds tend to be found in the inner halo. Clouds also span a wide range of shapes, and smaller clouds are more spherical than their more massive counterparts (Fig. 4).

(iii) With respect to cloud properties, most clouds ($\sim$90 per cent) have subsolar metallicities. However, clouds with metallicity as high as $\gtrsim 2\,Z_\odot$ exist. Most clouds have $\beta \sim 1$, indicating a balance of thermal and magnetic pressure. Magnetic pressure is larger in 2/3 of clouds, although most clouds outside the inner halo ($> 0.4 \times$ R$_{200c}$) are thermally dominated (in contrast to those in more massive haloes in TNG50; Nelson et al. 2020). Clouds are typically inhomogeneous in their metallicity content, and temperature-, density-, and pressure-structures; inner inhomogenities in their temperature structure are the most significant (Fig. 5).

(iv) While most clouds have relatively small radial velocities (of order $\sim$10 km s^{-1}), clouds tracing fast inflows and fast outflows are both present, and these are more prevalent at smaller galactocentric distances. Across the entire halo, inflowing clouds dominate ($\sim$73 per cent across all MW-like galaxies). Overall, clouds tend to be dominated by subvirial rotation. Metallicity correlates strongly with radial velocity: rapidly outflowing cold clouds are the most metal rich, whereas rapidly inflowing clouds are the least enriched, hinting at different physical origins (Figs 5 and 6).

(v) We compare the physical properties of clouds to their surrounding gas, defining local 'intermediate' (i.e. interface) and 'background' layers. On average, cold clouds in the CGM of MW-like galaxies are more metal-rich, denser, cooler, and preferentially inflowing with respective to their backgrounds. However, metallicity and velocity contrasts are small, of order 0.05 dex and 10 km s^{-1}, respectively, on average. We find a typical overdensity of $\chi \lesssim 10$, which is larger at smaller halocentric distances but much smaller than often assumed in idealized cloud simulations. While most clouds are only slightly underpressurized with respect to their surroundings when total (magnetic plus thermal) pressure is considered, they are significantly thermally underpressurized. This suggests magnetic fields may be an important pressure component in cold clouds in the CGM of MW-like galaxies (Fig. 8).

(vi) Cold clouds are not uniformly distributed throughout the halo but are strongly clustered. At all distances, clouds have more neighbouring clouds (within 10 kpc) than would be expected for a random distribution. This overclustering is greater towards the halo centre. We also find a clear clustering of cold clouds around satellite galaxies with large ($\gtrsim$10 per cent) baryon fractions. This suggests a stripping origin for at least a part of the cold cloud population (Figs 9 and 10).

(vii) Finally, we qualitatively compare results from TNG50 with observations of HVCs in the MW halo. The observed LOS velocity distribution of clouds is remarkably consistent with the average MW-like galaxy in TNG50: HVC abundance drops with increasing velocity. We show that the kinematics of cold clouds are not consistent with gravitational motion alone, suggesting that astrophysical feedback processes influence the motion of cold gas in the CGM. For clouds above $\sim 10^6$ M$_\odot$, no trend of mass with distance is seen in TNG50, which is consistent with the limited number of HVC observations available at this mass range. TNG50 predicts that (currently poorly constrained) 'low velocity' clouds are the most abundant and that their abundance is roughly independent of LOS velocity (Fig. 11).

This work is our first attempt to bridge studies of clouds using idealized, small-scale, controlled numerical experiments including wind tunnel or 'cloud crushing' simulations with those using MW-like galaxies realized through large-volume, cosmological, galaxy formation simulations. Building upon the study of Nelson et al. (2020) that found large abundances of cold clouds in high-mass TNG50 haloes, we have focused specifically on galaxies that resemble our own MW.

However, even with TNG50, we have here considered physical phenomena right at the edge of available numerical resolution. We cannot yet demonstrate that the abundance of small, low-mass clouds ($\lesssim$100's pc, $\lesssim 10^5$ M$_\odot$) is realistic, nor converged. Future work requires simulating the circumgalactic medium of MW-like galaxies at significantly higher resolution and with currently missing physics such as radiation transport and cosmic ray pressure components. This will enable the study of even smaller-scale structures while also better resolving the physical properties and evolutionary origins of the HVC-like cloud structures already present in TNG50.

ACKNOWLEDGEMENTS

RR and DN acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Group (grant number NE 2441/1-1). AP acknowledges funding by the DFG – Project-ID 138713538 – SFB 881 ('The Milky Way System', subprojects A01 and A06). RR is a Fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). RR and DN thank J. Christopher Howk for helpful discussions on the observational aspects of clouds. The authors also thank Ralf Klessen for valuable inputs. The authors thank the anonymous referee for constructive feedback that has greatly helped improve the quality of this work. The TNG50 simulation was run with compute time granted by the Gauss Centre for Supercomputing (GCS) under Large-Scale Projects GCS-DWAR on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK) and Nordrhein-Westfalen (MIWF). This analysis has been carried out on the VERA supercomputer of the Max Planck Institute for Astronomy (MPIA), operated by the Max Planck Computational Data Facility (MPCDF). This research has made use of NASA's Astrophysics Data System Bibliographic Services.

DATA AVAILABILITY

The IllustrisTNG simulations, including TNG50, are publicly available and accessible at www.tng-project.org/data (Nelson et al. 2019a). New data products for the MW/M31-like sample are now on the same website (Pillepich et al. in preparation). The 'cosmological cloud catalog' produced and used in this work is publicly released at www.tng-project.org/ramesh23b. Other data related to this publication are available upon reasonable request.

REFERENCES

Abruzzo M.W., Fielding D. B., Bryan G. L., 2022, preprint (arXiv:2210.15679)
Adams E. A. K., Giovanelli R., Haynes M. P., 2013, ApJ, 768, 77
Ayromlou M., Nelson D., Pillepich A., 2022, preprint (arXiv:2211.07659)
Binney J., Nipoti C., Fraternali F., 2009, MNRAS, 397, 1804
Bland-Hawthorn J., Gerhard O., 2016, ARA&A, 54, 529
Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, MNRAS, 398, 1150
Butsky I. S., Quinn T. R., 2018, ApJ, 868, 108
Collins J. A., Shull J. M., Giroux M. L., 2003, ApJ, 585, 336
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Donahue M., Voit G. M., 2022, Phys. Rep., 973, 1
Donnari M. et al., 2019, MNRAS, 485, 4817
Dutta A., Sharma P., Nelson D., 2022, MNRAS, 510, 3561
Faerman Y., Werk J. K., 2023, preprint (arXiv:2302.00692)
Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, ApJ, 894, L24
Fox A. J. et al., 2016, ApJ, 816, L11
Fraternali F., Binney J. J., 2006, MNRAS, 366, 449
Fraternali F., Marasco A., Armillotta L., Marinacci F., 2015, MNRAS, 447, L70
Genel S. et al., 2014, MNRAS, 445, 175
Gim H. B. et al., 2021, ApJ, 922, 69
Gnedin N. Y., Kravtsov A. V., 2011, ApJ, 728, 88
Goldsmith K. J. A., Pittard J. M., 2016, MNRAS, 461, 578
Gronke M., Oh S. P., 2018, MNRAS, 480, L111
Gronke M., Oh S. P., 2020, MNRAS, 492, 1970
Gronke M., Oh S. P., 2022, preprint (arXiv:2209.00732)
Grønnow A., Tepper-García T., Bland-Hawthorn J., McClure-Griffiths N. M., 2017, ApJ, 845, 69
Grønnow A., Tepper-García T., Bland-Hawthorn J., Fraternali F., 2022, MNRAS, 509, 5756
Haffner L. M., Reynolds R. J., Tufte S. L., 2001, ApJ, 556, L33
Hopkins P. F., Grudić M. Y., Wetzel A., Kereš D., Faucher-Giguère C.-A., Ma X., Murray N., Butcher N., 2020, MNRAS, 491, 3702
Hu C.-Y. et al., 2022, preprint (arXiv:2208.10528)
Jennings F., Beckmann R. S., Sijacki D., Dubois Y., 2022, MNRAS
Ji S., Oh S. P., McCourt M., 2018, MNRAS, 476, 852
Joung M. R., Bryan G. L., Putman M. E., 2012, ApJ, 745, 148
Klein R. I., McKee C. F., Colella P., 1994, ApJ, 420, 213
Lehner N., Howk J. C., 2011, Science, 334, 955
Lehner N., Sembach K. R., Lambert D. L., Ryans R. S. I., Keenan F. P., 1999, A&A, 352, 257
Lehner N., Keenan F. P., Sembach K. R., 2001, MNRAS, 323, 904
Lehner N., Howk J. C., Thom C., Fox A. J., Tumlinson J., Tripp T. M., Meiring J. D., 2012, MNRAS, 424, 2896
Lehner N., Howk J. C., Marasco A., Fraternali F., 2022, MNRAS, 513, 3228
Lepine J. R. D., Duvert G., 1994, A&A, 286, 60
Li Z., Hopkins P. F., Squire J., Hummels C., 2020, MNRAS, 492, 1841
Marasco A., Fraternali F., Lehner N., Howk J. C., 2022, MNRAS, 515, 4176
Marinacci F. et al., 2018, MNRAS, 480, 5113
Mayer L., Mastropietro C., Wadsley J., Stadel J., Moore B., 2006, MNRAS, 369, 1021
Mellema G., Kurk J. D., Röttgering H. J. A., 2002, A&A, 395, L13
Moss V. A., McClure-Griffiths N. M., Murphy T., Pisano D. J., Kummerfeld J. K., Curran J. R., 2013, ApJS, 209, 12
Muller C. A., Oort J. H., Raimond E., 1963, Acad. Sci. Paris C. R., 257, 1661
Naiman J. P. et al., 2018, MNRAS, 477, 1206
Nakamura F., McKee C. F., Klein R. I., Fisher R. T., 2006, ApJS, 164, 477
Nelson D., Vogelsberger M., Genel S., Sijacki D., Kereš D., Springel V., Hernquist L., 2013, MNRAS, 429, 3353
Nelson D. et al., 2018, MNRAS, 475, 624
Nelson D. et al., 2019a, Comput. Astrophys. Cosmol., 6, 2
Nelson D. et al., 2019b, MNRAS, 490, 3234
Nelson D. et al., 2020, MNRAS, 498, 2391
Olano C. A., 2008, A&A, 485, 457
Pakmor R., Marinacci F., Springel V., 2014, ApJ, 783, L20
Peek J. E. G., Heiles C., Putman M. E., Douglas K., 2009, ApJ, 692, 827
Peek J. E. G., Bordoloi R., Sana H., Roman-Duval J., Tumlinson J., Zheng Y., 2016, ApJ, 828, L20
Pillepich A. et al., 2018a, MNRAS, 473, 4077
Pillepich A. et al., 2018b, MNRAS, 475, 648
Pillepich A. et al., 2019, MNRAS, 490, 3196
Pittard J. M., Goldsmith K. J. A., 2016, MNRAS, 458, 1139
Planck Collaboration XIII, 2016, A&A, 594, A13
Popping G. et al., 2019, ApJ, 882, 137
Prada F., Klypin A. A., Cuesta A. J., Betancort-Rijo J. E., Primack J., 2012, MNRAS, 423, 3018
Putman M. E. et al., 2002, AJ, 123, 873
Ramesh R., Nelson D., Pillepich A., 2023, MNRAS, 518, 5754
Reid M. J. et al., 2014, ApJ, 783, 130
Richter P., Sembach K. R., Wakker B. P., Savage B. D., Tripp T. M., Murphy E. M., Kalberla P. M. W., Jenkins E. B., 2001, ApJ, 559, 318
Richter P. et al., 2017, A&A, 607, A48
Sander B., Hensler G., 2021, MNRAS, 501, 5330
Savage B. D. et al., 2000, ApJS, 129, 563
Scannapieco E., Brüggen M., 2015, ApJ, 805, 158
Schneider E. E., Robertson B. E., 2017, ApJ, 834, 144
Sembach K. R. et al., 2000, ApJ, 538, L31
Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, MNRAS, 452, 575
Smith M. C., Sijacki D., Shen S., 2018, MNRAS, 478, 302
Sparre M., Pfrommer C., Ehlert K., 2020, MNRAS, 499, 4261
Springel V., 2010, MNRAS, 401, 791
Springel V., Hernquist L., 2003, MNRAS, 339, 289
Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726
Springel V. et al., 2018, MNRAS, 475, 676
Springel V., Pakmor R., Zier O., Reinecke M., 2021, MNRAS, 506, 2871
Thom C., Peek J. E. G., Putman M. E., Heiles C., Peek K. M. G., Wilhelm R., 2008, ApJ, 684, 364
Tripp T. M., Song L., 2012, ApJ, 746, 173
Tripp T. M. et al., 2003, AJ, 125, 3122
Tufte S. L., Reynolds R. J., Haffner L. M., 1998, ApJ, 504, 773
Vogelsberger M. et al., 2014a, MNRAS, 444, 1518
Vogelsberger M. et al., 2014b, Nature, 509, 177
Wakker B. P., 1991, A&A, 250, 499
Wakker B. P., 2001, ApJS, 136, 463
Wakker B. P., van Woerden H., 1997, ARA&A, 35, 217
Wakker B. P. et al., 2007, ApJ, 670, L113
Wakker B. P., York D. G., Wilhelm R., Barentine J. C., Richter P., Beers T. C., Ivezić Ž., Howk J. C., 2008, ApJ, 672, 298
Weinberger R. et al., 2017, MNRAS, 465, 3291
Westmeier T., 2018, MNRAS, 474, 289
Williams M. E., Shelton R. L., 2022, ApJ, 926, 36
Yao Y., Shull J. M., Danforth C. W., 2011, ApJ, 728, L16
Yirak K., Frank A., Cunningham A. J., 2010, ApJ, 722, 412
Zech W. F., Lehner N., Howk J. C., Dixon W. V. D., Brown T. M., 2008, ApJ, 679, 460
Zhang D., Thompson T. A., Quataert E., Murray N., 2017, MNRAS, 468, 4801
Zinger E. et al., 2020, MNRAS, 499, 768

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 4

Azimuthal Anisotropy of Magnetic Fields in the Circumgalactic Medium Driven by Galactic Feedback Processes

4.1 Statement of contribution

- Scientific Analysis: My contribution was central. In addition, the analysis benefited greatly from feedback provided by collaborators.
- Figures: I independently produced the figures based on the planned scientific analysis. The process also greatly benefited from iterative input by collaborators.
- Writing: I primarily wrote most of the manuscript text. The drafts were further refined based on feedback from collaborators.
- Code and Simulation Development: This paper made use of the publicly available TNG50 simulation. The code to analyze the simulation data and generate plots was written by me, but it also benefited from input provided by colleagues.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **526,** 5483–5493 (2023) https://doi.org/10.1093/mnras/stad3104
Advance Access publication 2023 October 12

Azimuthal anisotropy of magnetic fields in the circumgalactic medium driven by galactic feedback processes

Rahul Ramesh ,[1★] Dylan Nelson ,[1] Volker Heesen [2] and Marcus Brüggen [2]

[1] *Universität Heidelberg, Zentrum für Astronomie, Institut für theoretische Astrophysik, Albert-Ueberle-Str. 2, D-69120 Heidelberg, Germany*
[2] *Hamburg University, Hamburger Sternwarte, Gojenbergsweg 112, D-21029 Hamburg, Germany*

Accepted 2023 October 9. Received 2023 September 15; in original form 2023 May 17

ABSTRACT

We use the TNG50 cosmological magnetohydrodynamical simulation of the IllustrisTNG project to show that magnetic fields in the circumgalactic medium (CGM) have significant angular structure. This azimuthal anisotropy at fixed distance is driven by galactic feedback processes that launch strong outflows into the halo, preferentially along the minor axes of galaxies. These feedback-driven outflows entrain strong magnetic fields from the interstellar medium, dragging fields originally amplified by small-scale dynamos into the CGM. At the virial radius, $z = 0$ galaxies with $M_\star \sim 10^{10}\,M_\odot$ show the strongest anisotropy (∼0.35 dex). This signal weakens with decreasing impact parameter, and is also present but weaker for lower mass as well as higher mass galaxies. Creating mock Faraday rotation measure (RM) sightlines through the simulated volume, we find that the angular RM trend is qualitatively consistent with recent observational measurements. We show that rich structure is present in the circumgalactic magnetic fields of galaxies. However, TNG50 predicts small RM amplitudes in the CGM that make detection difficult as a result of other contributions along the line of sight.

Key words: galaxies: haloes – galaxies: magnetic fields.

1 INTRODUCTION

Observational and theoretical studies suggest that galaxies are surrounded by a halo of gas that typically extends out to roughly the virial radius of their parent dark matter halos. Termed the circumgalactic medium (CGM), this multiscale multiphase reservoir of gas is believed to play a critical role in the evolution of galaxies [see Donahue & Voit (2022) for a recent review of the CGM].

Simultaneously, the CGM is affected by physical processes that take place within the galaxy. For instance, outflows driven by the central supermassive black hole (SMBH) can launch gas out of high-mass galaxies with $M_\star \gtrsim 10^{10.5}\,M_\odot$ and into the CGM at high velocities (Oppenheimer et al. 2020; Ramesh, Nelson & Pillepich 2023a), increasing cooling times and preventing future accretion, and star formation (Davies et al. 2020; Zinger et al. 2020). They can create bubbles of hot rarified gas (Pillepich et al. 2021) similar to the Fermi/eROSITA bubbles emerging from the galactic centre of the Milky Way (Su, Slatyer & Finkbeiner 2010; Predehl et al. 2020). For less massive galaxies, outflows driven by supernovae and stellar winds dominate (Fielding et al. 2017; Li & Tonnesen 2020), reshaping the CGM with significant inputs of mass, momentum, energy, and metals (Nelson et al. 2019b; Mitchell et al. 2020; Pandya et al. 2021).

Cosmological galaxy formation simulations tend to find that outflows are anisotropic, preferentially propagating perpendicular to the galactic disc, i.e along directions where the density of gas is lower than in the disc (e.g Nelson et al. 2019b). As a result, the CGM is predicted to possess angular anisotropies in key physical quantities including metallicity (Péroux et al. 2020), density, and temperature (Truong et al. 2021). The star formation activity of satellite galaxies is expected to respond to this gas anisotropy with an 'angular conformity' signal (Martín-Navarro et al. 2021). This picture of an azimuthally anisotropic CGM has also received observational support with inferences of angular dependencies in metallicity (Cameron et al. 2021) and density (Zhang & Zaritsky 2022) profiles.

In addition to hydrodynamical and thermal gas processes, non-thermal components including magnetic fields may play a key role in the CGM. Theoretical studies suggest that they can lengthen the lifetimes of small, cold gas clouds by suppressing fluid instabilities (Ji, Oh & McCourt 2018; Berlok & Pfrommer 2019; Sparre, Pfrommer & Ehlert 2020), or by contributing support in the form of magnetic pressure (Nelson et al. 2020; Ramesh, Nelson & Pillepich 2023b). It has recently become possible to probe the impact of magnetic fields in large-volume simulations of cosmic structure formation that self-consistently include magnetohydrodynamics within the context of a realistic galaxy population – IllustrisTNG is a notable example. Simulations have shown how magnetic fields may influence the propreties of halo gas (Pakmor et al. 2020; van de Voort et al. 2021), outflows (Steinwandel et al. 2020), and the hot intracluster medium (Vazza et al. 2014).

Observationally, extragalactic magnetic fields are difficult to measure, particularly owing to their relatively low values (Pomakov et al. 2022). However, recent observational studies have inferred $\sim \mu$G magnetic field strengths in the CGM of a number of galaxies

★ E-mail: oorc06@gmail.com

(Mao et al. 2017; Prochaska et al. 2019; Lan & Prochaska 2020; Mannings et al. 2022; Heesen et al. 2023; O'Sullivan et al. 2023), which is of order of the field strengths found in the gaseous halos of galaxies at $z \sim 0$ in the IllustrisTNG simulations (Marinacci et al. 2018; Nelson et al. 2018a; Ramesh, Nelson & Pillepich 2023a). Such inferences are made possible through measurements of Faraday rotation measure (RM), the phenomenon by which the polarization vector changes as light propagates through a region with non-zero magnetic field (see e.g. Kim et al. 2016; Rudnick & Cotton 2023; Takahashi 2023).

Recently, Heesen et al. (2023) used the LOFAR Two-metre Sky Survey (LoTSS; O'Sullivan et al. 2023) to identify a trend of RM strength with galaxy azimuthal angle. Stacking a sample of 21 unique galaxies (29 sightlines) at $z \sim 0$, with impact parameters $\lesssim$ 100 kpc, they find that RM values are ~2–3 times higher close to the minor axis of galaxies in comparison to the major axis. Similar work was performed by Böckmann et al. (2023, in press) who used RM values from the MIGHTEE-POL survey by MeerKAT in the XMM-LSS and COSMOS fields.However, they did not study the azimuthal dependence of the RM. Still, they find a central RM excess of 5.6 ± 2.3 rad m^{-2} with 2.5σ significance around star-forming galaxies with a median redshift of $z = 0.46$ for impact parameters below 130 kpc. They find no evidence for a correlation between RM and redshift. They conclude that mostly luminous, star-forming galaxies with impact parameters of < 130 kpc contribute to the RM of the distant radio sources.

In this paper, we use the cosmological magnetohydrodynamical simulation TNG50 to explore the angular structure of magnetic fields in the CGM, in order to provide theoretical interpretation for these RM observations, while also making predictions for future observational surveys. The paper is organized as follows: in Section 2, we describe the TNG50 simulation and the methods we utilize throughout the paper. We present and discuss our results in Section 3, and summarize our findings in Section 5.

2 METHODS

2.1 The TNG50 simulation

For our analysis we use the TNG50-1 simulation (hereafter, TNG50; Nelson et al. 2019b; Pillepich et al. 2019) of the IllustrisTNG project (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018a; Pillepich et al. 2018b; Springel et al. 2018), a series of cosmological magnetohydrodynamical simulations of galaxy formation. This is the highest resolution TNG simulation, with an average baryonic mass resolution of $\sim 8 \times 10^4$ M$_\odot$ within a volume of $\sim (50\ \mathrm{cMpc})^3$. It was run with the moving mesh code AREPO (Springel 2010), and the fiducial 'TNG model' for galaxy formation physics (Weinberger et al. 2017; Pillepich et al. 2018a): this includes key processes such as primordial and metal-line cooling, heating from a metagalactic background radiation field, self-shielding of dense gas, star formation, stellar evolution and enrichment, tracking of supernovae type Ia, II, and AGB stars, SMBH formation, accretion, merging, magnetic fields, and other physics that is expected to play an important role in the growth and evolution of galaxies.

Importantly, the TNG model includes both stellar and supermassive blackhole (SMBH) feedback processes. Stellar feedback transfers energy to gas in the interstellar medium through a decoupled wind scheme (Springel & Hernquist 2003) to produce mass-loaded galactic-scale outflows (Pillepich et al. 2018a). On the other hand, SMBH feedback operates in one of two modes depending on the accretion state of the SMBH. At high accretion rates, i.e. when the ratio of Bondi-to-Eddingtion rates is large, thermal energy is continuously deposited into neighbouring gas cells. At low accretion rates, feedback energy is instead imparted in the form of time-discrete high-velocity randomly oriented kinetic kicks (Weinberger et al. 2017). Low-mass SMBHs are preferentially in the thermal mode, while more massive SMBHs are typically in the kinetic mode, with the transition occurring roughly at $M_\star \sim 10^{10.5}$M$_\odot$ (corresponding to $M_{\rm BH} \sim 10^8$M$_\odot$; Nelson et al. 2018a; Weinberger et al. 2018).

Importantly, for the structure and directionality of galactic-scale outflows, the TNG feedback processes always inject energy isotropically. For stellar winds and thermal input from SMBHs in the high-accretion state this is true by construction for each energy injection, while the kinetic energy inputs from SMBHs in the low-accretion state are randomly oriented, such that when time-averaged energy injection in this mode is also isotropic. None the less, for both low-mass and high-mass galaxies, outflows in TNG naturally collimate and propagate along preferred directions (Nelson et al. 2019b).

The TNG simulations adopt a cosmology consistent with the Planck 2015 analysis (Planck Collaboration 2016), with: $\Omega_\Lambda = 0.6911$, $\Omega_{\rm m} = 0.3089$, $\Omega_{\rm b} = 0.0486$, and $h = 0.6774$.

2.2 Magnetic fields in TNG

A significant difference between TNG and other large-volume cosmological simulations is the inclusion of (ideal) magnetohydrodynamics (MHD; Pakmor, Bauer & Springel 2011; Pakmor, Marinacci & Springel 2014). At the start of the simulation ($z = 127$), a uniform primordial field of 10^{-14} comoving Gauss is seeded. This field then becomes amplified as a natural result of structure formation and feedback processes, including exponential early-time growth due to small-scale dynamos within halos, followed by a slower linear amplification phase due to differential rotation within disc galaxies (Pakmor et al. 2017, 2020).

In the TNG simulations, magnetic field strengths within galaxies reach $1 - 10$ μG levels by $z \sim 0$, consistent with observational inferences (Marinacci et al. 2018). Magnetic fields have complex topology and connections with galaxy and feedback processes. For instance, at the Milky Way-mass scale, $|B|$ strengths are systematically higher for blue versus red galaxies (Nelson et al. 2018a). Magnetic fields permeate the CGM of Milky Way-like galaxies (Ramesh, Nelson & Pillepich 2023a). Amplified magnetic fields are also ejected to large scales by feedback, producing over-magnetized bubbles around massive halos (Arámburo-García et al. 2021). Finally, the presence of magnetic fields alters galaxy properties, scaling relations, and global statistics, including galaxy sizes, SMBH and stellar masses, halo-scale gas fractions, the $z = 0$ stellar mass function, and the cosmic star formation rate density (Pillepich et al. 2018a).

Throughout this work, in order to derive magnetic field strength as a function of projected distance away from galaxies, we compute two-dimensional maps of $|B|$ using mass-weighted projections and the standard SPH kernel, with the smoothing length scaled in accordance with the gas cell radius. Unless otherwise stated, the projection depth (i.e. along the line of sight direction) scales with the impact factor: for an impact factor b, all gas within $\pm b$ along the perpendicular direction is included.

2.3 Synthetic Faraday RMs

We calculate synthetic Faraday RM values by integrating sightlines through the TNG50 simulation volume. At $z = 0$, we generate 10 000 primary sightlines around a galaxy sample spanning a range of stellar masses (Section 3.1), randomly located in angle and distance from

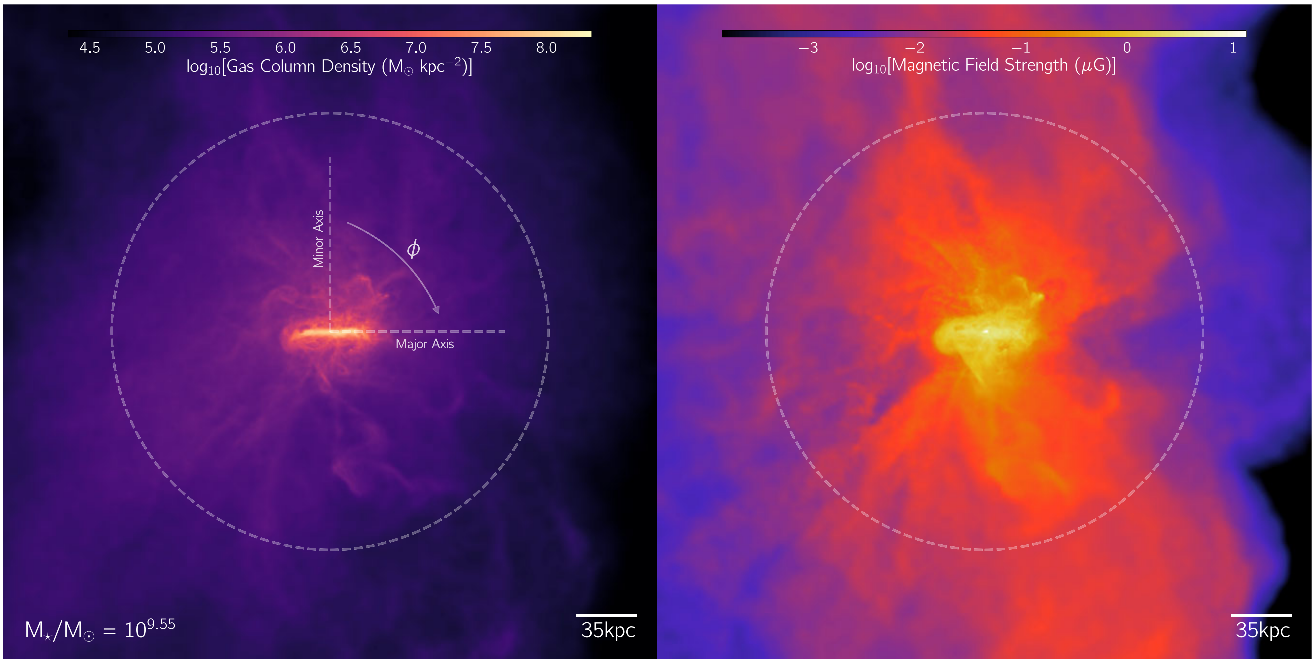

Figure 1. Visualization of a random TNG50 $M_\star \sim 10^{9.55} M_\odot$ galaxy at $z = 0$ rotated edge-on. The left panel shows the gas column density along a projection that extends $\pm 1.5\, R_{200c}$ in the line-of-sight direction, while the right panel shows a mass-weighted projection of magnetic field strength. In both panels, the dashed circles show the virial radius of the halo. At this mass range, except in the innermost regions of the halo, magnetic field strengths are preferentially stronger closer to the minor axis, i.e. at low values of ϕ, defined as the angle with respect to the minor axis, as shown in the left panel.

the galaxy centre out to twice the virial radius. All sightlines are oriented along the $\hat{z}$-axis of the simulation domain, and thus random with respect to the orientations of all galaxies. In addition, we use eight discrete snapshots at redshifts $z = \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 1.0\}$ to generate $N = 2000 \times 2000 = 4 \times 10^6$ random background sightlines and propagate each for a total distance equal to the simulation box length of 35 cMpc h^{-1}. Background sightlines can then be randomly stacked together with primary sightlines such that the total line-of-sight distance (i.e. projection depth) is equal to the light travel distance to a given redshift $z_{\rm bk}$. In all cases we calculate the line-of-sight integral,

$$\mathrm{RM} = 0.812 \int n_{\rm e}\, B_{||}\, (1+z)^{-2} \mathrm{d}l \quad \left[\mathrm{rad\,m^{-2}}\right] \tag{1}$$

by ray tracing through the gas distribution. We do so directly on its native representation of an unstructured Voronoi tessellation, computing the cell-by-cell intersections with each ray and the path lengths dl (in pc) through each cell (Nelson, in preparation). The values of the line-of-sight component of the magnetic field, $B_{||}$ (in μG), as well as the electron number density, $n_{\rm e}$ (in cm^{-3}), are constant within each cell, and taken directly from the simulation output. The only exception is for star-forming gas, which is subject to an additional pressure term from our effective two-phase interstellar medium model (Springel & Hernquist 2003). For these star-forming gas cells, we decompose the mass into its cold and hot components, and take $n_{\rm e}$ of the hot (ionized) phase only (following Nelson et al., in preparation). This approach is most consistent with the gas state in the TNG model and was also used in previous RM studies from simulations with similar ISM models (Pakmor et al. 2020; Mannings et al. 2022).

For the primary RM rays, we focus on galaxies that are oriented roughly edge-on in projection along the $\hat{z}$-axis of the simulation box, allowing us to define an azimuthal angle as in observations. We do so by restricting our analysis to galaxies with an ellipticity parameter $e \gtrsim 0.75$, where $e = 1 - b/a$ and a (b) is the semimajor (semiminor) axis length when viewed along this projection, as derived by STATMORPH morphological measurements in SDSS r-band synthetic stellar light images (Rodriguez-Gomez et al. 2019).

2.4 Galaxy sample and definitions

In this work we only select central galaxies, i.e. those that lie at the potential minimum of their respective dark matter halos, as determined by the SUBFIND substructure identification algorithm (Springel et al. 2001). Throughout this work, $M_\star$ refers to the stellar mass within twice the stellar half mass radius.

When considering true magnetic field strengths (and not RM sightlines), in order to connect to galaxy orientation and derive a well-defined azimuthal angle, we rotate galaxies to be edge-on using a diagonalization of the moment of inertia tensor.

3 RESULTS

In Fig. 1, we begin with a visualization of a gaseous halo around a ($z = 0$) TNG50 $M_\star \sim 10^{9.55} M_\odot$ galaxy. The image extends $\pm 1.5\, R_{200c}$ from edge-to-edge and in the projection direction, and shows the central galaxy oriented edge-on. In both panels, the dashed circle shows the virial radius, R_{200c}, of the parent halo.

In the left panel, we show the gas column density. Values are largest along the disc major axis direction, and decrease with increasing distance from the disc plane. The panel also illustrates the definition we follow for the azimuthal angle (ϕ), which we define with respect to the minor axis, i.e. the direction perpendicular to the disc when viewed in this edge-on projection. A value of $\phi = 0°$ thus corresponds to sightlines along the minor axis, while sightlines along the major axis are situated at $\phi = 90°$.

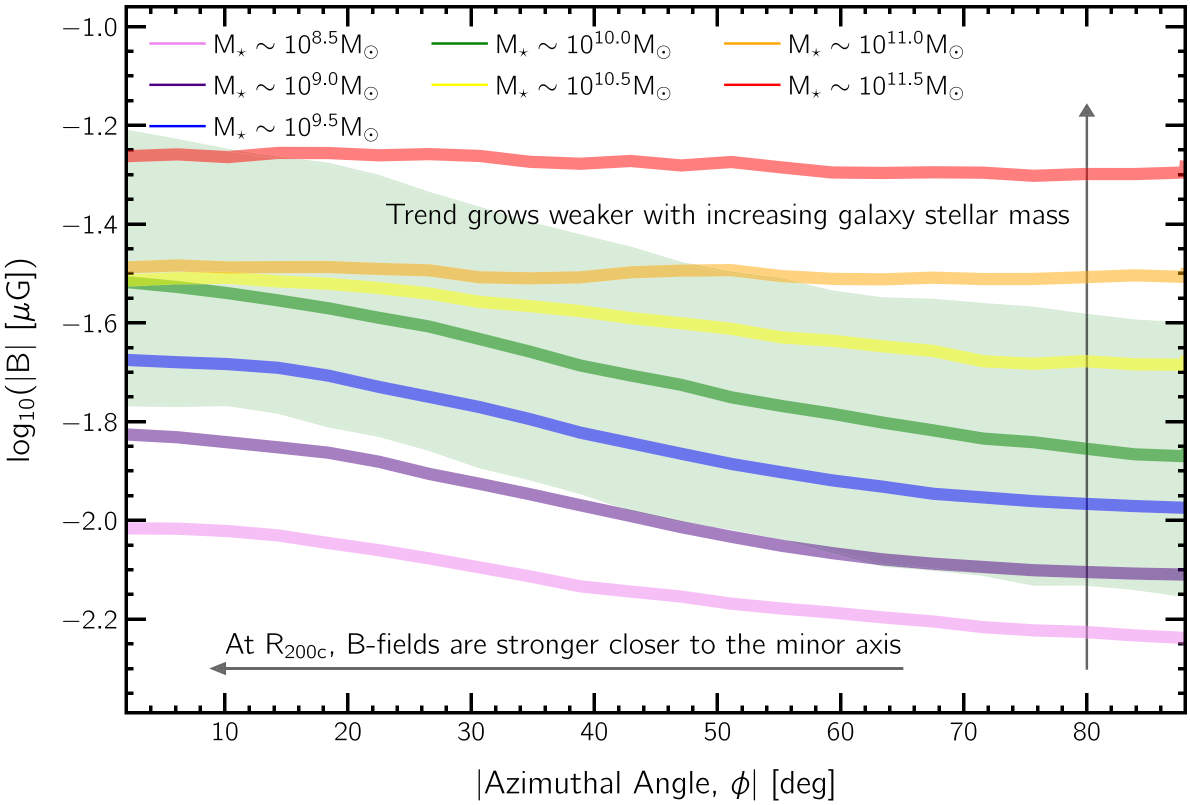

Figure 2. Magnetic field strength as a function of azimuthal angle (ϕ) at the virial radius, for galaxies in seven different mass bins from $10^{8.5} \lesssim M_\star/\mathrm{M}_\odot \lesssim 10^{11.5}$, from TNG50 at $z = 0$. Individual curves show medians. The shaded region, corresponding to 16th–84th percentile values, is shown only for one bin $M_\star \sim 10^{10}\ \mathrm{M}_\odot$ for visual clarity. The anisotropy is strongest for $M_\star \sim 10^{10}\ \mathrm{M}_\odot$ galaxies, and grows weaker for more massive galaxies.

The panel on the right shows a mass-weighted projection of magnetic field strength for the same galaxy, in the same projection. Magnetic fields are strongest along the disc, owing to the relatively large densities (Dolag et al. 2005). As a result, at small impact parameters ($\lesssim 0.25\ R_{200c}$), sightlines closer to the major axis ($\phi \sim 90^\circ$) are more magnetized than those at lower azimuthal angles, where the densities are typically lower than that of the disc. However, at larger galactocentric distances, sightlines at lower azimuthal angles are more magnetized versus higher ϕ counterparts. Such a trend is also visible in gas density (left panel), although the level of anisotropy is weaker.

In Fig. 2, we quantify this angular anisotropy in magnetic field strengths seen visually above. In seven different colours, we show the median magnetic field strength as a function of azimuthal angle for galaxies stacked in seven different mass bins (of width ±0.2 dex) spanning $10^{8.5} \lesssim M_\star/\mathrm{M}_\odot \lesssim 10^{11.5}$. The number of galaxies in each bin varies from a few ten to a few hundred, with the lower stellar mass bins comprising of more objects. In each case, we stack sightlines with impact parameters in the range [0.95, 1.05] R_{200c}, i.e. around the virial radius of the halo. For low-mass galaxies ($M_\star \lesssim 10^{10.5} \mathrm{M}_\odot$), a clear angular signal is visible, wherein B-field strengths are higher by 0.25–0.35 dex along the minor axis ($\phi = 0$) with respect to the major axis ($\phi = 90$). A strong stellar mass trend exists. Above $M_\star \sim 10^{10.0}\ \mathrm{M}_\odot$, the anisotropy weakens with increasing stellar mass, with the difference between the minor and major axes eventually decreasing to ~0 for the highest stellar mass bin.

In the shaded region, we show the 16th–84th percentile values, albeit only for one bin ($M_\star \sim 10^{10.0}\ \mathrm{M}_\odot$) for visual clarity. The percentile regions are relatively broad, stretching ~±0.25 dex across the median at all azimuthal angles. As discussed in Péroux et al. (2020), this poses a challenge to observational studies: if only a handful of sightlines are available for galaxies of a given stellar mass bin, randomly distributed in azimuthal angle, it is possible that such a trend may be missed due to the large scatter.

In Fig. 3, we show the trend with galactocentric distance. We consider galaxies with $M_\star \sim 10^{10.0}\ \mathrm{M}_\odot$ and stack sightlines at four different impact parameters: $0.25R_{200c}$ (dashed curves), $0.50R_{200c}$ (dash–dotted), R_{200c} (solid; same as Fig. 2), and $2R_{200c}$ (dotted).

At the smallest impact parameters, the median trend differs qualitatively from Fig. 2. As discussed above, sightlines at such small impact parameters intersect the extended gaseous disc, where the density (and hence magnetic field strength) is higher at larger ϕ, i.e. closer to the major axis. At larger impact parameters, the trend inverts, and the anisotropy is then stronger at larger galactocentric distances: a median difference between the minor and major axes of ~ 0.2 dex at $0.50R_{200c}$ increases to ~ 0.35 dex at R_{200c}, and further to ~ 0.45 dex at $2R_{200c}$. At this mass range, the anisotropy persists up to impact parameters as large as $\sim 10R_{200c}$, although the field strengths at such distances are much smaller ($\lesssim 10^{-3.5}\ \mu$G) in comparison to the halo. While not shown explicitly, we mention that all trends of Fig. 3 are qualitatively similar for galaxies in other stellar mass bins.

We note that these trends of magnetic field strength versus azimuthal angle in the CGM are qualitatively similar to angular anisotropies of other gas properties predicted by recent cosmological hydrodynamical simulations including TNG (Péroux et al. 2020; Truong et al. 2021; Yang et al. 2023). In particular, the direction of the effect is consistent with the angular modulation of CGM gas density (Truong et al. 2021), as expected.

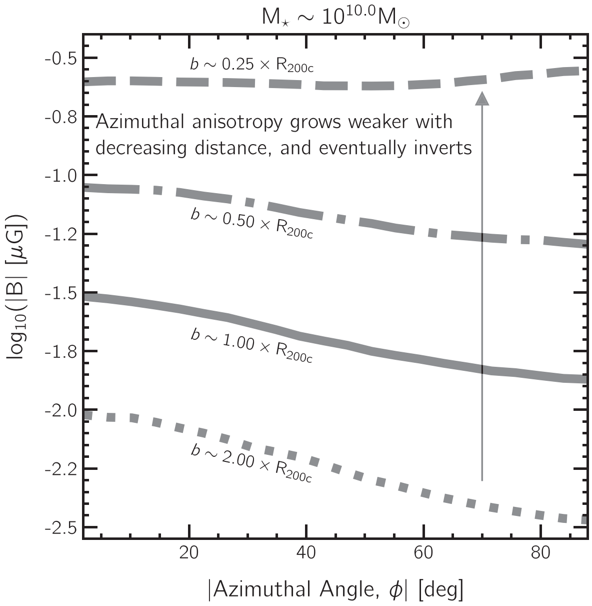

Figure 3. Magnetic field strength as a function of azimuthal angle (ϕ) at different impact parameters, for galaxies with stellar mass $\sim 10^{10} M_\odot$. The angular anisotropy grows weaker with decreasing distance, and eventually inverts close to the disc ($\lesssim 0.25 \times R_{200c}$).

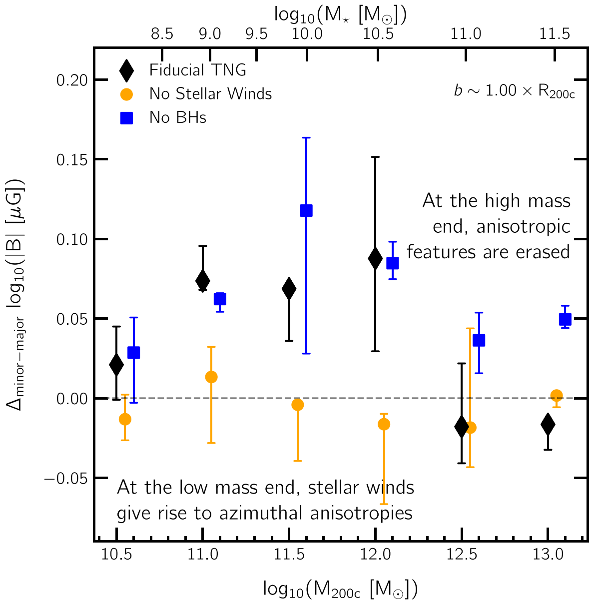

Figure 4. Difference in B-field strengths between minor/major axes at R_{200c}, as a function of halo mass, for different TNG model variation runs, with the fiducial model in diamonds (see text). For reference we include the mean stellar mass as a function of halo mass on the top x-axis, from TNG50. Stellar feedback processes launch dense magnetized gas perpendicular to the disc, giving rise to angular anisotropies at the low stellar mass end ($M_\star \lesssim 10^{10.5} M_\odot$). Without stellar feedback (orange), the anistropy is not present. For more massive galaxies, the kinetic mode of SMBH feedback dominates and results in galaxy quenching. The resulting outflows act to erase the anistropy of CGM magnetic fields at this mass scale, which is otherwise still present (blue).

In Fig. 4, we identify the physical processes responsible for producing the anisotropy in B-fields discussed so far. To do so, we use the TNG variations, a large set of simulations run with perturbations to the fiducial model [as discussed and first presented in Pillepich et al. (2018a)]. These use a smaller box of $L_{\rm box} \sim 37$ Mpc and have a resolution equal to that of TNG100-1 ($m_{\rm b} \sim 2 \times 10^6\ M_\odot$). Specifically, we consider three variants: (i) Fiducial TNG: the baseline model with all aspects unchanged, (ii) No Stellar Winds: excluding resolve stellar feedback processes i.e. the stellar feedback-driven galactic winds, and (iii) No BHs: the TNG model, but with no black holes seeded, and hence no black hole feedback.

To quantify the anisotropy between the minor and major axes, we stack all sightlines into two bins of azimuthal angle: $[0, 45]^\circ$ and $[45, 90]^\circ$. All sightlines in the former bin are associated with the minor axis, and the latter with the major axis. We focus on the difference between these two bins, labelled as $\Delta_{\rm minor-major}$ on the y-axis, as a function of halo virial mass M_{200c}. For reference, the median stellar mass from TNG50 is shown on the top x-axis. Symbols correspond to median differences, while error bars show the 16th–84th percentile values. Results from the fiducial, no stellar winds, and no BHs variation boxes are shown in black, orange, and blue, respectively. In this plot, we only select sightlines with impact parameters close to the virial radius. Although not shown explicitly, the trends discussed here are qualitatively similar at other impact parameters as well.

For lower halo masses $M_{200c} \lesssim 10^{12}\ M_\odot$, corresponding to $M_\star \lesssim 10^{10.5}\ M_\odot$ in TNG50, values of $\Delta_{\rm minor-major}$ are more or less consistent between the Fiducial and No BHs runs, and are positive, i.e. in this mass regime, angular anisotropies are present, and are largely unaffected by the activity of the central SMBH. However, when stellar winds are turned off (orange points), the anisotropy disappears, with $\Delta_{\rm minor-major}$ oscillating randomly around a value of zero. This suggests that outflows driven by stellar winds launch over-dense magnetized gas preferentially perpendicular to the disc, and are largely responsible for the angular structure of magnetic fields. As noted in earlier studies (e.g. Nelson et al. 2019b), outflows prefer the direction perpendicular to the disc since densities along these paths are lower in comparison to that of the disc, and a path of least resistance is established along the pressure gradient.

For more massive halos, i.e. $M_{200c} \gtrsim 10^{12.5}\ M_\odot$, corresponding to $M_\star \gtrsim 10^{11.0}\ M_\odot$ in TNG50, the picture begins to change: at this mass scale, differences arise between the Fiducial and No BHs cases (black versus blue points), while the Fiducial and No Stellar winds cases (black versus orange points) are consistent with each other. Specifically, strong anisotropies between the minor and major axes cease to exist ($\Delta_{\rm minor-major} \sim 0$) when black holes are present. As discussed in Truong et al. (2021), this is likely due to the quenching of central galaxies by the kinetic mode of SMBH feedback: as galaxies become quenched, a well-defined disc-like structure ceases to exist, due to which outflows no longer have a preferred direction. As a result, the kinetic kicks distribute energy uniformly in all directions, suppressing any angular structure previously present.

While no other large volume cosmological simulation beyond TNG currently include the effects of magnetic fields, a few zoom-in projects do. This enables the study of the impact of MHD on the evolution of cosmic gas. For instance, the Auriga simulations produce magnetic field strengths and radial profiles which are overall similar to those in TNG, albeit only for Milky Way-like galaxies (Pakmor et al. 2020). The FIRE simulations include variations with MHD (e.g. Ponnada et al. 2022), and previous studies have shown that field strengths at $z = 0$ are generally weaker than in TNG (see Ramesh, Nelson & Pillepich 2023a, and discussion therein). In both cases, although the angular anisotropy of magnetic fields has not yet

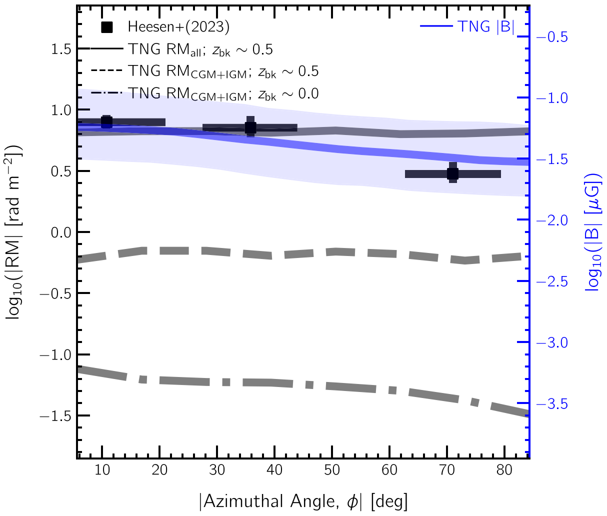

Figure 5. RM as a function of azimuthal angle (ϕ) at ~100 kpc, for galaxies with stellar mass $\sim 10^{10}\,M_\odot$. Black points show results from Heesen et al. (2023), while the three grey lines correspond to three measurements from TNG, as elaborated in the main text. The blue curve shows the true median three-dimensional magnetic field strength, for comparison.

been studied, given their origin in the collimation of galactic-scale outflows, we speculate that such signals may be present as well.[1]

3.1 The observable: RM

In Fig. 5, we compare the results of the TNG50 simulation to observations, where we show the RM as a function of azimuthal angle (left y-axis and grey lines). The black points correspond to those from Heesen et al. (2023), for impact parameters $b < 100$ kpc. The three grey curves correspond to three different measurements using TNG for a selection of 70 galaxies with stellar masses $M_\star \sim 10^{10}\,M_\odot$, the typical stellar mass of objects in the sample of Heesen et al. (2023). The dash–dotted line only includes the $z = 0$ ($d_{LoS} \sim 50$ Mpc) primary sightlines for each galaxy, i.e. the effect of the CGM and the local IGM ($RM_{CGM} + RM_{IGM}$; $z_{bk} \sim 0.0$). The dashed line includes additional background sightlines to yield a total line-of-sight distance out to $z \sim 0.5$, i.e. the effect of the cosmological IGM is also included ($RM_{CGM} + RM_{IGM}$; $z_{bk} \sim 0.5$; see Section 2.3 for details). This is the typical background polarized source redshift of the Heesen et al. (2023) sample (see also Hackstein et al. 2020). In the TNG simulations, these contributions from the cosmological IGM can be non-negligible, especially at $z \lesssim 2$, as a result of star formation and SMBH driven outflows expelling magnetized gas into their local environments (Arámburo-García et al. 2023). In addition, other intervening galaxies may impact the net RM signal, although this is expected to be rare in the TNG simulations: Arámburo-García et al. (2023) find that a fraction $\sim 10^{-4}$ of sightlines are impacted by such contributions. However, some other studies suggest that contributions due to intervening galaxies may be more important, possibly yielding a non-Gaussian distribution of background RMs (e.g. Shah & Seta 2021).

In addition to the intervening CGM and IGM, other contributions to the observed RM signals are present. In particular, from the Milky Way galaxy and halo (foreground), and from the host environment/galaxy of the polarized source itself (background). The former contribution is removed by subtracting a model for the Galactic RM, i.e. the effect of gas in the Milky Way. However, this subtraction unavoidably introduces uncertainty. In addition, instrumental noise and other contaminating contributions to the RM may also be present (see e.g. Hackstein et al. 2019).

We collectively account for all these external contributions, which we label RM_{ext}, by adding random Gaussian noise, with the mean (~ 5.68 rad m^{-2}) and standard deviation (~ 5.50 rad m^{-2}) computed from the Heesen et al. (2023) RMs at impact parameters $b > 500$ kpc (see also Basu et al. 2018). The solid black curve shows the resulting RM values, i.e. $RM_{all} = RM_{ext} + RM_{CGM} + RM_{IGM}$; $z_{bk} \sim 0.5$.

The dash–dotted curve (the intervening CGM and local IGM only) shows a drop of ~0.4 dex from the minor to major axes, consistent with the ~0.4 dex drop seen in the the black points. Note, however, that the RM excess, i.e. the difference between the values at the minor and major axes in linear scale, predicted by TNG (~ 0.05 rad m^{-2}) is about two orders of magnitude smaller than that of Heesen et al. (2023; ~ 4 rad m^{-2}). While the relative anisotropy predicted by TNG is consistent with that of Heesen et al. (2023), the absolute excess is thus smaller (see discussion below).

However, when the effect of the cosmological IGM is included (dashed line), or once external effects are accounted for (solid line), the trend is buried in these non-local signals. This is because the true signal (RM_{CGM}) predicted by TNG is almost always weaker than the other contributions (RM_{ext} and RM_{IGM}; $z_{bk} \sim 0.5$) along the line of sight (see also Basu et al. 2018).

In blue, for the purpose of comparison, we overlay the true (three-dimensional) magnetic field strength-azimuthal angle median relation of these galaxies, i.e. the two y-axes are not related and have an arbitrary relationship for visualization only. The true magnetic field trend shows a clear decrease of ~0.4 dex between the minor versus major axes directions. At the same time, the slope at large azimuthal angles $\phi \gtrsim 45^\circ$ is less steep (in $\log(B)$) than the observations (in log(RM)), although the 16th–84th percentile (shaded) region is broad. TNG is therefore in qualitative agreement with the results of Heesen et al. (2023), while also predicting that such trends are challenging to observe (see below). That is, the azimuthal angle dependence of RM is largely overtaken by external contributions of gas along the line of sight: either gas surrounding the polarized source, or IGM gas, in addition to any other sources of observational noise.

This is similar to challenges faced in the interpretation of dispersion measure (DM) values from observations of fast radio burst (FRBs). Like RMs, the observed DM values result from the combined contribution of the Milky Way, the IGM, the host, and the source. It is not straightforward to disentangle these components and retrieve the contribution of the source alone, owing to its weak amplitude in comparison to other components (e.g. Zhang et al. 2021).

In Fig. 6, we move away from angular structure, and focus on the broad trend of RM as a function of impact parameter. In the main panel we show three different synthetic values of RM for the same selection of 70 galaxies with stellar masses $M_\star \sim 10^{10}\,M_\odot$ as above. These are: $RM_{CGM} + RM_{IGM}$; $z_{bk} \sim 0.0$ (dot–dashed), $RM_{CGM} + RM_{IGM}$; $z_{bk} \sim 0.5$ (dashed) and RM_{all} (solid). Further, we split the galaxies into two sub-samples based on their specific star formation rate (sSFR)[2] Galaxies with values less than the <10th percentile (i.e.

[1]Note that the majority of the FIRE simulations run so far do not include black holes, and hence only the impact of stellar feedback would be visible.

[2]We define sSFR as the ratio of star formation rate (within twice the stellar half mass radius) and $M_\star$.

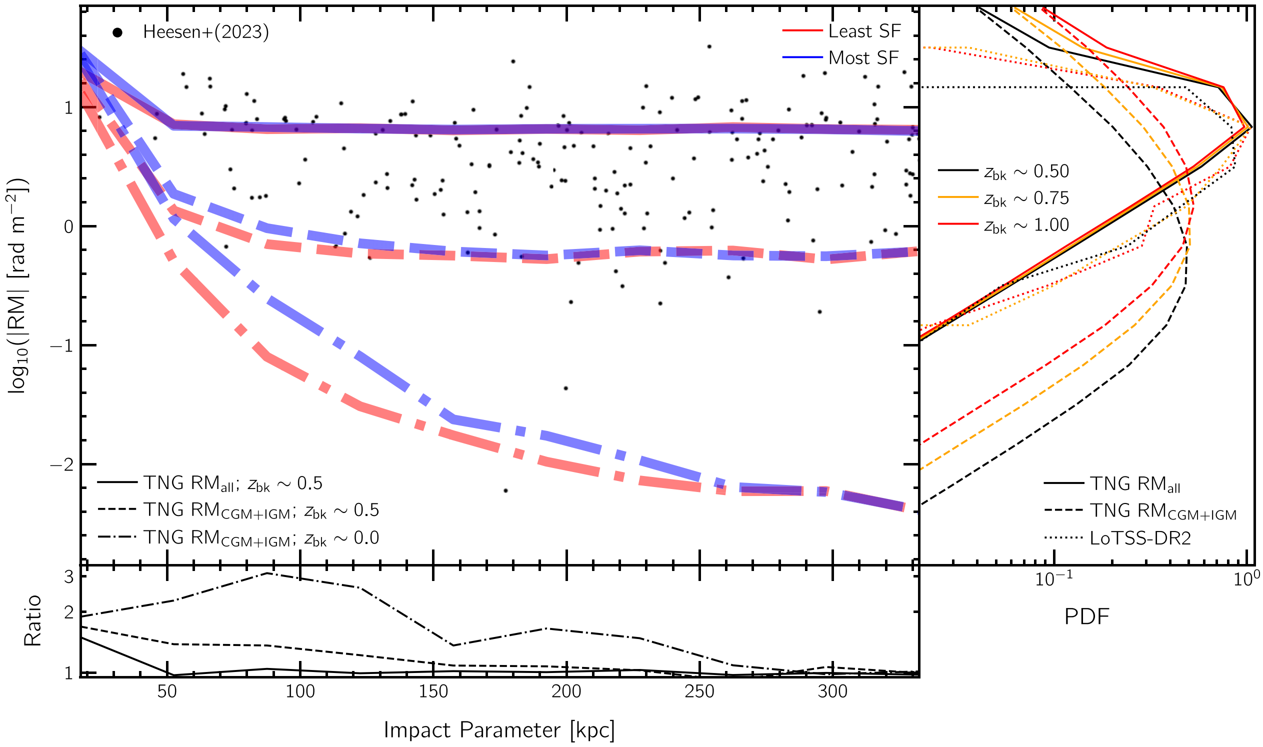

Figure 6. In the main panel, we show the trend of RM as a function of impact parameter for galaxies with $M_\star \sim 10^{10}\ M_\odot$. The sample is split into two bins based on sSFR values: <10th percentile (least SF; red) and >90th percentile (most SF; blue). Solid, dashed, and dot–dashed curves correspond to three different synthetic TNG RM measurements, as elaborated in the main text. The ratio of the blue-to-red curves is shown in the bottom panel. Black dots show data points from Heesen et al. (2023). In the right panel, we show PDFs of RM for random sightlines from TNG (solid and dashed curves), and those from the LoTSS-DR2 catalogue (dotted curves; O'Sullivan et al. 2023). Different colours correspond to different projection depths (refer to main text for details).

the seven least star forming in this sample; median sSFR $\sim 10^{-10.7}$ $\mathrm{yr^{-1}}$; red) are contrasted against the >90th percentile (i.e. the seven most star forming; median sSFR $\sim 10^{-9.8}$ $\mathrm{yr^{-1}}$; blue). We therefore study whether star-forming versus quiescent galaxies have different RM amplitudes in their CGM, motivated by similar dichotomies present in the TNG simulations e.g. in halo OVI abundance (Nelson et al. 2018b) and in X-ray luminosity (Truong et al. 2020), the latter potentially detected in eROSITA stacking (Chadayammuri et al. 2022; Comparat et al. 2022).

In the CGM plus local IGM only RMs (dot–dashed curve), a difference between the blue versus red curves is apparent out to impact parameters of ~250 kpc, wherein galaxies with higher star formation rates have larger values of RM. We interpret this as higher star forming rate (SFR) galaxies launching stronger galactic-scale outflows, driving more magnetized gas into the CGM. When the cosmological IGM is included, i.e. out to $z_{bk} \sim 0.5$ (dashed curves), the differences between the two curves disappear beyond ~125 kpc. With the addition of $\mathrm{RM_{ext}}$ component, differences are only clear close to the galaxy ($\lesssim 50$ kpc). To quantify these differences between the two sub-samples, the bottom panel shows the ratio of the blue to red curves. The peak amplitude of this effect is a factor of ~3 larger RMs in high versus low SFR galaxies, which occurs at impact parameters of ~100 kpc. The quantitative strength of this dichotomy undoubtedly varies with galaxy mass as well as redshift, but its qualitative existence shows that observables of the magnetic fields in the CGM encode information on galactic feedback processes.

In black dots, we show the full sample from Heesen et al. (2023). These observed galaxies have not been split by star formation rate. All but a few measurements are for impact parameters $\gtrsim 50$ kpc, and the vast majority are at much larger distances, as statistically expected for random background-foreground associations. Observational detection of the differential RM signal between the blue and red sub-samples would clearly be a challenge given the level of scatter, once again primarily as a result of external contributions by gas along the line of sight. Only at small impact parameters does the local CGM signal dominate, making small separation pairs extremely valuable. Future large surveys such as POSSUM on ASKAP (Gaensler et al. 2010), with LOFAR (O'Sullivan et al. 2020), and SKA-era instruments (Heald et al. 2020), will enable this science. Note that the linearity of the solid curves at small impact parameters is a result of using large bins.

In the right panel, we show PDFs of RM values for several cases of interest which are *not* related to intervening galaxies. For a selection of 10^6 *random* TNG sightlines, we show $\mathrm{RM_{all}}$ in solid curves, and $\mathrm{RM_{CGM}} + \mathrm{RM_{IGM}}$ in dashed lines. The dotted lines are observational data from the LoTSS-DR2 catalogue (O'Sullivan et al. 2023). In each of these cases, we split the samples (i.e. select) three different values of z_{bk}: 0.5 (black), 0.75 (orange), and 1.0 (red). Each has different path-lengths of the cosmological IGM contribution included. For the LoTSS catalogue, we bin the spectroscopic redshifts of polarized sources around the three different redshift values.

In the $\mathrm{RM_{CGM}} + \mathrm{RM_{IGM}}$ case (dashed), increasing the traversal depth through the IGM leads to a systematic increase in RM values. The black curve, corresponding to $z_{bk} \sim 0.5$, peaks at a value of ~0.4 rad $\mathrm{m^{-2}}$. Increasing the projection depth to $z_{bk} \sim 0.75$ shifts the PDF by ~0.4 dex towards higher RM values, and further by ~0.4 dex for $z_{bk} \sim 1.0$. Integrating to greater distances thus yields a larger value of RM, indicating that large-scale contributions as sightlines traverse the magnetized cosmic web continue to accrue. This suggests that, even though the integral in equation (1) is sensitive to the sign of $B_{||}$, positive and negative values i.e. field reversals do not average out (see also Vacca et al. 2016).

The solid curves show the PDFs with the RM_{ext} included. In this case, the three PDFs peak at roughly the same value, with differences between them apparent only in the high-RM tails ($\gtrsim 15$ rad m^{-2}). That is, in the presence of the RM_{ext} component, the effect of the cosmological IGM is sub-dominant and largely disappears, except at values where the IGM component is comparable to RM_{ext}. The three LoTSS (dotted) curves are in broad agreement with TNG, i.e. independent of z_{bk} except at the high-RM end. In the presence of a RM_{ext} component obtained from Heesen et al. (2023), TNG thus predicts that extracting RM_{IGM} is a challenge, as for RM_{CGM}.

While recent studies using Auriga (Pakmor et al. 2020) and FIRE (Ponnada et al. 2022) have quantified the radial profiles of RM using RM mocks (see also Liu, Kretschmer & Teyssier 2022), only the RM_{CGM} component is included. Direct comparison with these studies is hence not possible, but we mention here that the RM values of Pakmor et al. (2020) are comparable to our RM_{IGM}; $z_{bk} \sim 0.0$ case, i.e. when only the local IGM is included, while FIRE predicts lower values of RM, much like the weaker magnetic field strength values.

4 DISCUSSION

Overall, we confirm that RM_{CGM} measurements contain rich information on halo magnetic fields. TNG50 predicts that the angular anisotropy signal of magnetic fields observed by Heesen et al. (2023) is present over a large range of halo and stellar masses, and not just for the $M_\star \sim 10^{10}\,M_\odot$ bin that Heesen et al. (2023) mainly probed. Moreover, TNG50 also predicts that the observed RM values are sensitive to feedback processes within the galaxy. Observing such trends would require denser sampling and better sensitivity, likely to be made possible by future surveys including POSSUM on ASKAP, as well as the Square Kilometer Array (SKA). Such measurements will be instrumental in improving the feedback physics implemented in future cosmological hydrodynamical simulations. However, if the TNG50 simulation is realistic, then this suggests that CGM RMs are challenging to measure observationally, especially at large impact parameters, due to their low values.

While the RM predicted by TNG is about two orders of magnitude smaller than the LoTSS catalogue (O'Sullivan et al. 2023), the discrepancy is even larger in comparison to some other results (e.g. Prochaska et al. 2019; Sobey et al. 2019; Lan & Prochaska 2020; Seta & Federrath 2021). One possibility is that the RM values predicted by TNG50 are indeed too low, in which case the most likely culprit is insufficiently large late time magnetic field strengths. Although under-predicted RM values can be a result of simulated n_e values being too low, this is unlikely to dominate our results (Fig. A1). Further, while magnetic fields could be missing some degree of large-scaling ordering and hence have smaller coherence lengths than reality, it is unlikely that this can improve the RM signal by over an order of magnitude (Fig. B1).

If simulated field strengths are too low, this could point to missing small-scale physical processes which boost amplification, or simply insufficient numerical resolution to capture small-scale dynamos on sufficiently short timescales (Pakmor et al. 2017). While distributions of field strengths are believed to be converged at TNG50-1 resolution (e.g. Marinacci et al. 2015; Ramesh & Nelson 2023), some studies suggest that the Powell et al. (1999) divergence-cleaning scheme utilized by AREPO may not be sufficiently accurate (e.g. Hopkins & Raives 2016), although Pakmor & Springel (2013) show that the results produced by AREPO do not differ significantly in comparison to constrained transport schemes. Carretti et al. (2023) recently showed that a uniform, homogeneous primordial magnetic field, like the one assumed in TNG, may not be an accurate representation of the universe, although various studies suggest that the exact nature of the seeding does not have a significant impact (e.g. Pakmor & Springel 2013; Marinacci et al. 2015; Vazza et al. 2017; Martin-Alvarez et al. 2021). Nevertheless, the TNG model can certainly be improved by incorporating more sophisticated seed fielding processes (e.g. Garaldi, Pakmor & Springel 2021).

Finally, the apparent tension of RM could be due to complexities of the data analysis and comparison itself. Most clearly, observed values include the combined line-of-sight contributions from many sources other than the extragalactic CGM of the target galaxy. These encompass the local hot interstellar medium, the Milky Way halo itself, the cosmological IGM, and the intrinsic source. Many of these terms are highly uncertain. For example, the exceptionally high time variability of some repeating FRBs suggests that the source contribution, its subtraction to obtain RM_{CGM} for a given line of sight, and thus any comparison with cosmological simulations inherits this non-negligible uncertainty (Hilmarsson et al. 2021).

The simulation outcomes will also change with the inclusion of more complex physics such as non-ideal MHD, or with the inclusion of cosmic rays which is believed to alter the flow structure of CGM gas (Butsky et al. 2022). This could potentially yield larger RM and magnetic field strengths in the CGM, enhancing angular anisotropies in extragalactic halos to a level similar to that observed by Heesen et al. (2023). Moreover, differences in magnetic field properties between cosmological simulations run with different codes suggest that the numerical techniques employed to solve the equations of magnetohydrodynamics may still play a significant role (Ramesh, Nelson & Pillepich 2023a). Constraints from observational data will help constrain magnetic fields in future cosmological magnetohydrodynamical simulations.

5 SUMMARY AND CONCLUSIONS

In this paper, we have shown that magnetic fields in the CGM of TNG50 galaxies have significant angular structure, such that gas along the minor axes of galaxies is more strongly magnetized with respect to the major axes. Our main findings can be summarized as follows:

(i) At the virial radius, these angular anisotropies are strongest for galaxies with stellar masses $M_\star \sim 10^{10}\,M_\odot$, decreasing in strength by $\sim$0.35 dex between the two galactic axes. The anisotropy grows weaker for both lower and higher mass galaxies (Fig. 2).

(ii) A trend with distance is also present: the strength of the anisotropy weakens with decreasing distance, and for very small impact parameters ($\lesssim 0.25 \times R_{200c}$), the direction of anisotropy eventually inverts, i.e gas along the major axis is more strongly magnetized, as a result of higher densities along the disc (Fig. 3).

(iii) This anisotropy arises as a result of two different feedback processes. At lower galaxy masses, stellar feedback launches dense, magnetized gas perpendicular to the disc plane of galaxies, yielding higher B-field strengths along the direction of the minor axis. This mode of feedback dominates for galaxies with $M_\star \lesssim 10^{10.5} M_\odot$. For more massive galaxies, SMBH (active galactic nuclei) feedback quenches galaxies, washing out this angular structure as strong outflows produce underdense bubbles, while outflows also lose their directional collimation in high-mass galaxies lacking gaseous discs (Fig. 4).

(iv) We make predictions for Faraday RM studies. TNG shows a qualitatively similar dependence of RM on galactic azimuthal angle as seen observationally (Heesen et al. 2023). However, TNG also predicts small RM amplitudes that are difficult to observe above external contributions from gas present along the line of sight between the observer and the polarized source (Fig. 5). In order to reproduce observed RM values, the magnetic field would have to be stronger.

(v) The profile of RM as a function of distance depends on galaxy SFR, with greater RM values around galaxies with higher SFR. This dichotomy reflects overall stronger magnetic fields in the CGM of high SFR galaxies in comparison to their lower SFR and quenched counterparts. However, TNG50 predicts that these differences will be challenging to observe owing to their relatively low magnitudes, except at small impact parameters (Fig. 6).

With the inclusion of magnetic fields, the TNG simulations provide a unique starting point to study the amplification, late-time structure, and impact of magnetic fields in/on galaxies and their gaseous halos. In the future, hydrodynamical simulations with substantially improved resolution in the CGM, together with more sophisticated models for stellar and SMBH feedback, as well as currently absent physics including cosmic rays, are needed to fully capture the role of magnetic fields in galaxies and their halos.

ACKNOWLEDGEMENTS

RR and DN acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Group (grant number NE 2441/1-1). RR is a Fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). MB acknowledges support from the DFG under Germany's Excellence Strategy – EXC 2121 "Quantum Universe' – 390833306. The authors thank the anonymous referee for constructive feedback that has greatly helped improve the quality of this work. The TNG50 simulation was run with compute time granted by the Gauss Centre for Supercomputing (GCS) under Large-Scale Projects GCS-DWAR on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS). This analysis has been carried out on the VERA supercomputer of the Max Planck Institute for Astronomy (MPIA), operated by the Max Planck Computational Data Facility (MPCDF).

DATA AVAILABILITY

The IllustrisTNG simulations, including TNG50, are publicly available at www.tng-project.org/data (Nelson et al. 2019a). Other data related to this publication is available upon request.

REFERENCES

Arámburo-García A., Bondarenko K., Boyarsky A., Nelson D., Pillepich A., Sokolenko A., 2021, MNRAS, 505, 5038
Arámburo-García A., Bondarenko K., Boyarsky A., Neronov A., Scaife A., Sokolenko A., 2023, MNRAS, 519, 4030
Basu A., Mao S. A., Fletcher A., Kanekar N., Shukurov A., Schnitzeler D., Vacca V., Junklewitz H., 2018, MNRAS, 477, 2528
Berlok T., Pfrommer C., 2019, MNRAS, 489, 3368
Butsky I. S. et al., 2022, ApJ, 935, 69
Cameron A. J. et al., 2021, ApJ, 918, L16
Carretti E., O'Sullivan S. P., Vacca V., Vazza F., Gheller C., Vernstrom T., Bonafede A., 2023, MNRAS, 518, 2273
Chadayammuri U., Bogdán Á., Oppenheimer B. D., Kraft R. P., Forman W. R., Jones C., 2022, ApJ, 936, L15
Comparat J. et al., 2022, A&A, 666, A156
Davies J. J., Crain R. A., Oppenheimer B. D., Schaye J., 2020, MNRAS, 491, 4462
Dolag K., Grasso D., Springel V., Tkachev I., 2005, J. Cosmol. Astropart. Phys., 2005, 9
Donahue M., Voit G. M., 2022, Phys. Rep., 973, 1
Fielding D., Quataert E., McCourt M., Thompson T. A., 2017, MNRAS, 466, 3810
Gaensler B. M., Landecker T. L., Taylor A. R., POSSUM Collaboration,, 2010, in American Astronomical Society Meeting Abstracts #215. p. 470.13
Garaldi E., Pakmor R., Springel V., 2021, MNRAS, 502, 5726
Hackstein S., Brüggen M., Vazza F., Gaensler B. M., Heesen V., 2019, MNRAS, 488, 4220
Hackstein S., Brüggen M., Vazza F., Rodrigues L. F. S., 2020, MNRAS, 498, 4811
Heald G. et al., 2020, Galaxies, 8, 53
Heesen V. et al., 2023, A&A, 670, L23
Hilmarsson G. H. et al., 2021, ApJ, 908, L10
Hopkins P. F., Raives M. J., 2016, MNRAS, 455, 51
Ideguchi S., Tashiro Y., Akahori T., Takahashi K., Ryu D., 2017, ApJ, 843, 146
Ji S., Oh S. P., McCourt M., 2018, MNRAS, 476, 852
Kim K. S., Lilly S. J., Miniati F., Bernet M. L., Beck R., O'Sullivan S. P., Gaensler B. M., 2016, ApJ, 829, 133
Lan T.-W., Prochaska J. X., 2020, MNRAS, 496, 3142
Lerche I., 1970, Ap&SS, 6, 481
Li M., Tonnesen S., 2020, ApJ, 898, 148
Liu Y., Kretschmer M., Teyssier R., 2022, MNRAS, 513, 6028
Mannings A. G. et al., 2022, preprint(arXiv:2209.15113)
Mao S. A. et al., 2017, Nature Astron., 1, 621
Marinacci F., Vogelsberger M., Mocz P., Pakmor R., 2015, MNRAS, 453, 3999
Marinacci F. et al., 2018, MNRAS, 480, 5113
Martin-Alvarez S., Katz H., Sijacki D., Devriendt J., Slyz A., 2021, MNRAS, 504, 2517
Martín-Navarro I., Pillepich A., Nelson D., Rodriguez-Gomez V., Donnari M., Hernquist L., Springel V., 2021, Nature, 594, 187
Mitchell P. D., Schaye J., Bower R. G., Crain R. A., 2020, MNRAS, 494, 3971
Naiman J. P. et al., 2018, MNRAS, 477, 1206
Nelson D. et al., 2018a, MNRAS, 475, 624
Nelson D. et al., 2018b, MNRAS, 477, 450
Nelson D. et al., 2019a, Comput. Astrophys. Cosmol., 6, 2
Nelson D. et al., 2019b, MNRAS, 490, 3234
Nelson D. et al., 2020, MNRAS, 498, 2391
O'Sullivan S. P. et al., 2020, MNRAS, 495, 2607
O'Sullivan S. P. et al., 2023, MNRAS, 519, 5723
Oppenheimer B. D. et al., 2020, MNRAS, 491, 2939
Pakmor R., Springel V., 2013, MNRAS, 432, 176
Pakmor R., Bauer A., Springel V., 2011, MNRAS, 418, 1392
Pakmor R., Marinacci F., Springel V., 2014, ApJ, 783, L20
Pakmor R. et al., 2017, MNRAS, 469, 3185
Pakmor R. et al., 2020, MNRAS, 498, 3125
Pandya V. et al., 2021, MNRAS, 508, 2979
Péroux C., Nelson D., van de Voort F., Pillepich A., Marinacci F., Vogelsberger M., Hernquist L., 2020, MNRAS, 499, 2462
Pillepich A. et al., 2018a, MNRAS, 473, 4077
Pillepich A. et al., 2018b, MNRAS, 475, 648
Pillepich A. et al., 2019, MNRAS, 490, 3196
Pillepich A., Nelson D., Truong N., Weinberger R., Martin-Navarro I., Springel V., Faber S. M., Hernquist L., 2021, MNRAS, 508, 4667
Planck Collaboration, 2016, A&A, 594, A13
Pomakov V. P. et al., 2022, MNRAS, 515, 256
Ponnada S. B. et al., 2022, MNRAS, 516, 4417
Powell K. G., Roe P. L., Linde T. J., Gombosi T. I., De Zeeuw D. L., 1999, J. Comput. Phys., 154, 284
Predehl P. et al., 2020, Nature, 588, 227

Prochaska J. X. et al., 2019, Science, 366, 231
Ramesh R., Nelson D., 2023, preprint(arXiv:2307.11143)
Ramesh R., Nelson D., Pillepich A., 2023a, MNRAS, 518, 5754
Ramesh R., Nelson D., Pillepich A., 2023b, MNRAS, 522, 1535
Rodriguez-Gomez V. et al., 2019, MNRAS, 483, 4140
Rudnick L., Cotton W. D., 2023, preprint (arXiv:2304.02728)
Seta A., Federrath C., 2021, MNRAS, 502, 2220
Shah H., Seta A., 2021, MNRAS, 508, 1371
Sobey C. et al., 2019, MNRAS, 484, 3646
Sparre M., Pfrommer C., Ehlert K., 2020, MNRAS, 499, 4261
Springel V., 2010, MNRAS, 401, 791
Springel V., Hernquist L., 2003, MNRAS, 339, 289
Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, MNRAS, 328, 726
Springel V. et al., 2018, MNRAS, 475, 676
Stasyszyn F. A., de los Rios M., 2019, MNRAS, 487, 4768
Steinwandel U. P., Dolag K., Lesch H., Moster B. P., Burkert A., Prieto A., 2020, MNRAS, 494, 4393
Su M., Slatyer T. R., Finkbeiner D. P., 2010, ApJ, 724, 1044
Takahashi K., 2023, PASJ, 75, S50
Truong N. et al., 2020, MNRAS, 494, 549
Truong N., Pillepich A., Nelson D., Werner N., Hernquist L., 2021, MNRAS, 508, 1563
Vacca V. et al., 2016, A&A, 591, A13
van de Voort F., Bieri R., Pakmor R., Gómez F. A., Grand R. J. J., Marinacci F., 2021, MNRAS, 501, 4888
Vazza F., Brüggen M., Gheller C., Wang P., 2014, MNRAS, 445, 3706
Vazza F., Brüggen M., Gheller C., Hackstein S., Wittor D., Hinz P. M., 2017, Class. Quantum Gravity, 34, 234001
Weinberger R. et al., 2017, MNRAS, 465, 3291
Weinberger R. et al., 2018, MNRAS, 479, 4056
Yang T., Davé R., Cui W., Cai Y.-C., Peacock J. A., Sorini D., 2023, preprint(arXiv:2305.00602)
Zhang H., Zaritsky D., 2022, ApJ, 941, 18
Zhang Z. J., Yan K., Li C. M., Zhang G. Q., Wang F. Y., 2021, ApJ, 906, 49
Zinger E. et al., 2020, MNRAS, 499, 768

APPENDIX A: SIMULATED ELECTRON NUMBER DENSITIES

In Fig. A1, we show probability distribution functions (PDFs) of the electron number density of gas cells from the different TNG50 boxes. TNG50-1, which has been used throughout the bulk of this work, is shown in black. The lower resolution boxes of TNG50-2 (mass resolution lower by a factor of 8 with respect to TNG50-1), TNG50-3 (factor of 64), and TNG50-4 (factor of 512) are shown in blue, orange, and red, respectively. The peaks of these PDFs shift slightly towards lower values of n_e with improving resolution, although convergence of these distributions between runs is decent barring TNG50-4.

The dashed vertical grey line shows the best observational estimate of n_e currently available from the Milky Way CGM (Donahue & Voit 2022). This is a value of 10^{-4} cm^{-3}, estimated at a galactocentric distance of ~ 50–100 kpc. For a direct comparison with this value, in the black dashed line, we show the PDF constructed from a set of Milky Way-like galaxies from TNG50-1 (for details on sample selection, see Ramesh, Nelson & Pillepich 2023a). Similar to the above mentioned observational result, this PDF is restricted to gas cells within a distance of 50–100 kpc. This PDF peaks close to $\sim 10^{-4}$ cm^{-3}, and is in good agreement with the observed value. Although this comparison is limited to a single observational value, and hence halo mass, this suggests that the under-prediction of RM by TNG is unlikely a result of simulated electron number densities being too low.

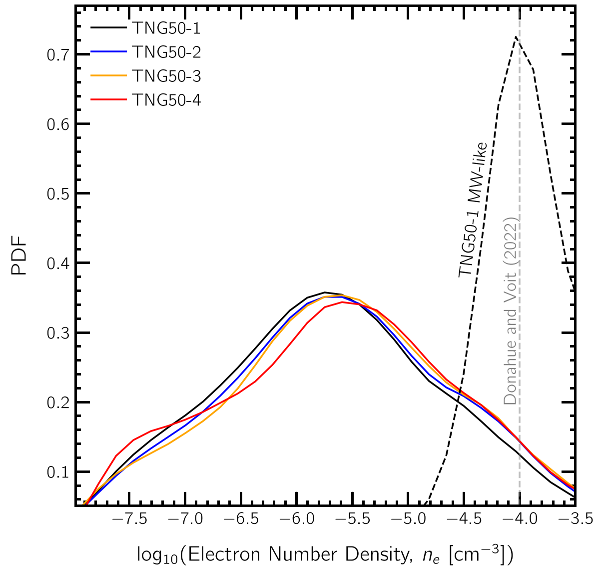

Figure A1. PDFs of electron number density of gas cells from the different TNG50 boxes. Except for the low-resolution TNG50-4 box, PDFs show decent convergence. In addition, the dashed vertical grey line shows the best estimate of n_e available from the Milky Way CGM currently (Donahue & Voit 2022). For a direct comparison, in the dashed black line, we also include the PDF for those gas cells from a sample of TNG50-1 Milky Way-likes, which is in good agreement with the observed value.

APPENDIX B: IMPACT OF TURBULENCE AND COHERENCE LENGTHS

The measured RM signal is believed to be impacted by factors like the magnetic coherence length (e.g. Stasyszyn & de los Rios 2019) and turbulence (e.g. Lerche 1970; Ideguchi et al. 2017) of gas along the line of sight leading to depolarization. In Fig. B1, we explore the impact of such effects on the derived TNG RM values by studying several test cases.

We pick 100 halos at random from ± 0.2 dex bins of $\log_{10}(M_\star) \in [8.5, 9.0, 9.5, 10.0, 10.5, 11.0, 11.5]$, i.e. a subset of halos from Fig. 2. For each of these halos, we calculate three different sets of RM values by varying the magnetic field component along the line of sight ($B_\parallel$):

(i) Real: using the actual values of $B_\parallel$ from the TNG simulation.
(ii) Perfect: assuming perfect alignment of all magnetic field vectors along the line of sight direction, with $B_\parallel$ set to the norm of the magnetic field of the corresponding gas cell.
(iii) Random: for each gas cell, we set $B_\parallel$ to the product of the norm of its magnetic field and a random number drawn uniformly from the interval $[-1.0, 1.0]$. This yields a random alignment of magnetic field vectors along the line of sight.

The perfect alignment case is a proxy for a scenario where the coherence length of B-fields is large, as large as the extent of the halo in this case. The random alignment case probes the other extreme where coherence lengths are small, possibly caused by highly turbulent gas throughout the line of sight interval.[3] To assess the impact of these tests specifically on CGM gas, we restrict to those

[3] In our case, small corresponds roughly to the gas spatial resolution in the CGM, i.e. on $\sim$ kpc scales. We cannot assess decoherence on even smaller scales as our method requires a constant magnetic field across a single resolution element.

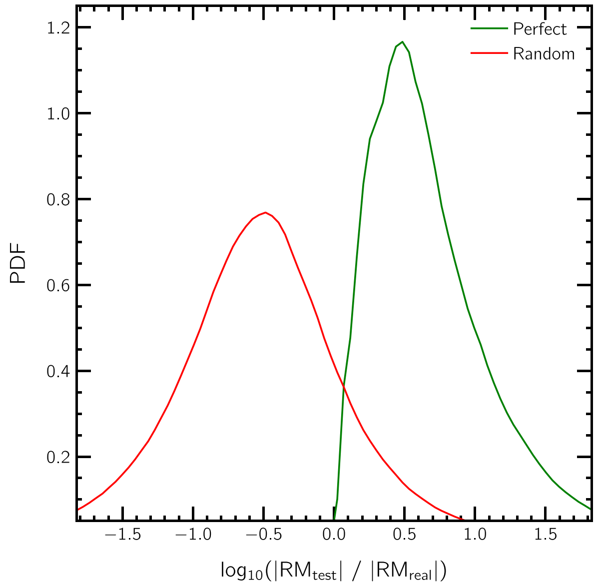

Figure B1. Several test cases studying the effect of magnetic coherence length on the obtained RM signal. As explained in the main text, we construct cases where the field lines are perfectly coherent (green) and randomly oriented (red). PDFs of the ratios of these cases with respect to the real case are shown. These tests suggest that the RM signal typically varies by less than an order of magnitude when the coherence length is varied between these two extremes.

sightlines with impact parameters in the range [0.15, 1.0] × $R_{200,c}$, and a projection column of $\pm R_{200,c}$ centred on the galaxy. For each selected sightline, we compute ratios of RM of the test case with respect to the real case. The PDF of these ratios is shown in Fig. B1: the perfect alignment case in green, and the random alignment in red.

As expected, the green curve shows that a larger coherence length results in a larger RM signal. The typical 'strengthening' of the signal is a factor of ∼ 3, although the values of RM along some sightlines are boosted by factors of $\gtrsim 10$, i.e. by an order of magnitude or more. On the other extreme, randomizing $B_{\parallel}$ leads to the RM signal growing weaker, typically by a factor of ∼ 0.3 with respect to the real case. Surprisingly, there are cases where the RM signal grows despite this randomization, suggesting that these sightlines are already dominated by highly turbulent fields.

Overall, while the magnetic coherence length can have an impact on the obtained RM value, the above test cases suggest that the effect is typically less than an order of magnitude. It is thus unlikely that the under-prediction of RM values by TNG50-1 is dominated by this effect.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 5

Zooming in on the Circumgalactic Medium with GIBLE: Resolving small-scale gas structure in cosmological simulations

5.1 Statement of contribution

- Scientific Analysis: My contribution was central. In addition, the analysis benefited greatly from feedback provided by collaborators.
- Figures: I independently produced the figures based on the planned scientific analysis. The process also greatly benefited from iterative input by collaborators.
- Writing: I primarily wrote most of the manuscript text. The drafts were further refined based on feedback from collaborators.
- Code and Simulation Development: This paper made use of the new set of GIBLE simulations, which I set up and ran during the first half of my PhD. This also involved developing a 'CGM refinement' module into the AREPO code, which required extensive coding in C. The cumulative effort of Project GIBLE greatly benefited from input from my thesis supervisor, including, but not limited to, feedback on optimizing the code architecture and generating initial conditions for these new simulations. The code to analyze the simulation data and generate plots was written by me, but it also benefited from input provided by colleagues.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY
MNRAS **528,** 3320–3339 (2024) https://doi.org/10.1093/mnras/stae237
Advance Access publication 2024 January 22

Zooming in on the circumgalactic medium with GIBLE: Resolving small-scale gas structure in cosmological simulations

Rahul Ramesh ★ and Dylan Nelson
Universität Heidelberg, Zentrum für Astronomie, Institut für theoretische Astrophysik, Albert-Ueberle-Straße 2, D-69120 Heidelberg, Germany

Accepted 2024 January 17. Received 2024 January 11; in original form 2023 July 19

ABSTRACT

We introduce Project GIBLE (Gas Is Better resoLved around galaxiEs), a suite of cosmological zoom-in simulations where gas in the circumgalactic medium (CGM) is preferentially simulated at ultra-high numerical resolution. Our initial sample consists of eight galaxies, all selected as Milky Way-like galaxies at $z = 0$ from the TNG50 simulation. Using the same galaxy formation model as IllustrisTNG, and the moving-mesh code AREPO, we re-simulate each of these eight galaxies maintaining a resolution equivalent to TNG50-2 ($m_{\rm gas} \sim 8 \times 10^5$ M$_\odot$). However, we use our super-Lagrangian refinement scheme to more finely resolve gas in the CGM around these galaxies. Our highest resolution runs achieve 512 times better mass resolution ($\sim 10^3$ M$_\odot$). This corresponds to a median spatial resolution of $\sim$75 pc at 0.15 $R_{200,\,\rm c}$, which coarsens with increasing distance to $\sim$700 pc at the virial radius. We make predictions for the covering fractions of several observational tracers of multiphase CGM gas: H I, Mg II, C IV, and O VII. We then study the impact of improved resolution on small scale structure. While the abundance of the smallest cold, dense gas clouds continues to increase with improving resolution, the number of massive clouds is well converged. We conclude by quantifying small scale structure with the velocity structure function and the autocorrelation function of the density field, assessing their resolution dependence. The GIBLE cosmological hydrodynamical simulations enable us to improve resolution in a computationally efficient manner, thereby achieving numerical convergence of a subset of key CGM gas properties and observables.

Key words: galaxies: evolution – galaxies: haloes– methods: numerical.

1 INTRODUCTION

Galaxy formation and evolution is a complex process. Over the years, various cosmological simulation projects have attempted to improve our understanding of the physics governing cosmic growth. While the first set of cosmological simulations were dark matter only, more recent simulations over the past decade or so have begun including important baryonic processes, thereby making it possible to statistically reproduce baryonic structures of the observed universe (for a recent overview of cosmological simulations, see Vogelsberger et al. 2020).

These cosmological hydrodynamical simulations have often fared well when reproducing galaxy- and halo-scale global properties. For instance, the stellar size–mass relation obtained from simulations (e.g. Genel et al. 2018; Henden, Puchwein & Sijacki 2020; Feldmann et al. 2023) is increasingly consistent with those of observations (e.g. Shen et al. 2003; Bernardi et al. 2014). Simulated galaxies can separate into two populations based on their star formation activity, thereby giving rise to a bimodal galaxy colour distribution (e.g. Nelson et al. 2018a; Cui et al. 2021), akin to observational studies (e.g. Strateva et al. 2001; Balogh et al. 2004). Quenched fractions of simulated satellite galaxies are found to correlate with their environment (e.g. Wright et al. 2019; Donnari et al. 2021), as shown by observations (e.g. Geha et al. 2012; Medling et al. 2018).

However, the consistency between cosmological hydrodynamical simulations and observations is less straightforward with respect to the distribution of gas around galaxies – the circumgalactic medium (CGM). This reservoir of gas is believed to be multiscale and multiphase, with cold, dense clouds embedded in a volume-filling warm–hot phase, making it a challenge to numerically simulate (see Donahue & Voit 2022 for a recent review of the CGM).

Cosmological hydrodynamical simulations typically make robust predictions for the properties of CGM gas on relatively larger scales and lower densities. For instance, at typical resolutions realized by modern cosmological simulations, predictions of the column density distribution functions of the warm–hot phase, traced by O VI, O VII, and O VIII (e.g. Nelson et al. 2018b; Wijers et al. 2019), are in decent agreement with those of observations (e.g. Thom & Chen 2008; Danforth et al. 2016). At the Milky Way mass, the IllustrisTNG simulations produce AGN feedback-driven bubbles of hot, rarified gas on either side of galactic discs (Pillepich et al. 2021), much like the *Fermi*/eROSITA bubbles emerging from the galactic centre of the Milky Way (Su, Slatyer & Finkbeiner 2010; Predehl et al. 2020). Observed angular anisotropies of metallicity, density, and magnetic field strengths in CGM gas (Cameron et al. 2021; Zhang & Zaritsky 2022; Heesen et al. 2023) are also found in some simulations (Péroux et al. 2020; Truong et al. 2021; van de Voort et al. 2021; Ramesh et al. 2023c).

★ E-mail: rahul.ramesh@stud.uni-heidelberg.de

However, the picture differs for the more clumpy, colder, and dense phase. Some studies, although not all, find higher H I covering fractions with improving resolution, without any clear sign of convergence (Peeples et al. 2019; van de Voort et al. 2019). Recent work also suggests that improving numerical resolution results in the existence of a larger number of cold, dense discrete CGM gas clouds (Nelson et al. 2020; Ramesh, Nelson & Pillepich 2023b). In reality, observations of clouds in the Milky Way (e.g. Hsu et al. 2011) and extragalactic haloes (e.g. Zahedy et al. 2019) reveal the existence of cloud sizes at least as small as $\sim$100–200 pc, thereby requiring simulations to adequately sample the underlying gas distribution on these scales to capture such objects.

While recent cosmological large-volume and zoom-in simulations are significantly better resolved than those of the past (see e.g. introduction of Pillepich et al. 2023), it is still a challenge to simulate CGM gas at the level of resolution discussed earlier. As a result, various idealized simulations are commonly used to study phenomena at these scales, either using wind-tunnel (e.g. Scannapieco & Brüggen 2015; Goldsmith & Pittard 2016; Gronke & Oh 2020) or halo-scale (Fielding et al. 2017; Kopenhafer, O'Shea & Voit 2023; Tan & Fielding 2023) setups. Such numerical experiments can explore small-scale physics, albeit without the realism of a cosmological CGM.

Recently, an alternative approach has emerged, using cosmological hydrodynamical simulations that preferentially increase the resolution of CGM gas while maintaining a relatively coarse resolution in other regions, most importantly the galaxy itself (Hummels et al. 2019; Peeples et al. 2019; Suresh et al. 2019; van de Voort et al. 2019). These simulations provide the ideal bridge between highly resolved idealized setups, and cosmological setups that naturally capture the complexity of the CGM.

In this paper, we describe and introduce a new set oschmidf simulations run with our 'CGM refinement' technique, an improvement of past work (Suresh et al. 2019). These are the GIBLE (Gas Is Better resoLved around galaxiEs) simulations, currently comprised of eight Milky Way-like galaxies at $z = 0$, drawn from the TNG50 simulation. The rest of this paper is organized as follows: in Section 2, we describe the galaxy formation model and the CGM refinement technique used. Results are presented in Section 3, and we summarize the paper in Section 4.

2 METHODS

2.1 Galaxy formation model and initial conditions

Project GIBLE uses the well-studied and validated IllustrisTNG model (Weinberger et al. 2017; Pillepich et al. 2018a) to simulate the basic elements of galaxy formation and evolution, within the moving mesh code AREPO (Springel 2010). The model accounts for star formation, stellar and AGN feedback, stellar enrichment, metal-line cooling of gas, and other important processes, along with the inclusion of magnetic fields (Pakmor, Marinacci & Springel 2014). The star formation and stellar feedback model is the subgrid two-phase model of Springel & Hernquist (2003).

The IllustrisTNG project is made of three suites of simulations: the relatively large-volume boxes of TNG100 and TNG300 (Marinacci et al. 2018; Naiman et al. 2018; Springel et al. 2018; Nelson et al. 2018a; Pillepich et al. 2018b), and the high-resolution, smaller box TNG50 (Nelson et al. 2019; Pillepich et al. 2019). Project GIBLE (re-)simulates eight galaxies from the TNG50 simulation, all of which are a part of the Milky Way-like (MW-like) sample from Ramesh et al. (2023a).[1] To recap, these are a set of galaxies that (i) have a stellar mass in the range $10^{10.5}\,\mathrm{M}_\odot$ to $10^{10.9}\,\mathrm{M}_\odot$, measured within a 3D aperture of 30 kpc, (ii) have no other galaxy having $\mathrm{M}_\star > 10^{10.5}\,\mathrm{M}_\odot$ within a distance of 500 kpc, (iii) reside in haloes that satisfy virial mass, $\mathrm{M}_{200,c} < 10^{13}\,\mathrm{M}_\odot$, (iv) are discy, either based on a constraint on the minor-to-major axis ratio of the stellar mass distribution ($s < 0.45$) or through visual inspection of stellar light maps, and (v) lie at the potential minimum of their host haloes, as identified by the friends-of-friends algorithm (Davis et al. 1985), that is, they are central galaxies.

While the MW-like sample is assembled using TNG50-1, the highest resolution version of the TNG50 box, we derive initial conditions for Project GIBLE from TNG50-2 (lower resolution by a factor eight in mass), following a cross-match of galaxies across these two runs. The eight galaxies in our sample have been carefully selected to ensure a good amount of diversity: two have been selected based on the presence of an SMC- and LMC-like satellite within a distance of 300 kpc (Engler et al. 2021), two that host *Fermi*-/eROSITA-like bubbles on either side of the central galaxy (Pillepich et al. 2021), two that possess a bar and/or stellar disc length similar to the 'real' Milky Way galaxy, one with a M31-like companion (i.e. a galaxy with a similar mass as M31 that is moving towards the central MW-like galaxy; Pillepich et al. 2023), and the eighth is chosen at random from the aforementioned set of 132 TNG50 MW-like galaxies. The galaxy with the M31-like companion does not have a significant cold gas disc in TNG50-2 at $z = 0$, due to prior disruption by AGN feedback during the quenching process. This galaxy is an outlier in terms of multiple properties, adding to the diversity of the sample.

With the sample selected, we construct zoom initial conditions using the multimass version of the N-GENIC code (Angulo et al. 2012; Springel et al. 2021). At $z = 0$, we identify particles that are within 4 times the halo virial radius R_{200c}. The high-resolution region is then constructed using the convex hull of the $z = 127$ positions of these particles. The rest of the box is populated with resolution elements of higher mass, with a gradual worsening of resolution with increasing distance from the central zoom region.

2.2 Super-Lagrangian refinement scheme

Project GIBLE uses an improvement of the 'CGM refinement' previously explored in Suresh et al. (2019), with several technical updates to the underlying code. The most major change concerns the functionality of black hole particles: earlier, it was possible to seed, that is, create, at most a single-black hole particle. Moreover, this black hole particle was only a psuedo-particle, with accretion and subsequent feedback processes disabled. Our new CGM refinement technique removes these limitations, and black holes now function as normal TNG black holes (Weinberger et al. 2017).

Black holes are important in our refinement scheme as they are used to trace the centre of the target halo, and thus the refinement region where CGM gas is super-refined (Suresh et al. 2019). Since only one refinement region is generally desired per run, the earlier CGM Zoom refinement scheme limited the number of seeded black holes to one. However, to be consistent with the TNG runs and the black hole feedback models, we now allow for the possibility of multiple black holes to be seeded, but 'flag' the first one to be formed. By default, this takes place in the first halo to cross a threshold halo

[1]This sample is a subset of the MW/M31-like galaxies of TNG50 presented in Pillepich et al. 2023.

mass of $\sim 7 \times 10^{10}$ M$_\odot$. When the flagged black holes merges, the flag is passed on to the descendent, thus ensuring that the flagged black hole is never lost through the course of the simulation.

Our refinement scheme builds on the already existing (de)refinement operations that AREPO performs to maintain (high-resolution) gas cells within a factor of two[2] of a given (constant) target mass, $m_{\rm target}$ (Springel 2010). Gas cells whose mass falls below the defined threshold, that is, $<0.5 \times m_{\rm target}$, are dissolved and their properties are conservatively distributed to their Voronoi neighbours. Those whose mass rises above the upper limit, that is, $>2.0 \times m_{\rm target}$, are split into two. This procedure ensures that all (high-resolution) gas cells contain roughly the same mass, at least up to a desired threshold. In our scheme, the constant target mass is replaced by a distance dependent target mass, $m_{\rm dist,\,target}(r)$, which is defined for every high-resolution gas cell, based on its distance from the flagged black hole, r

$$m_{\rm dist,target}(r) = \begin{cases} m_{\rm target} & \text{if } r < r_{\rm galaxy} \\ m_{\rm target}/{\rm RF} & \text{if } r_{\rm CGM,min} < r < r_{\rm CGM,max}, \\ m_{\rm target} & \text{if } r > r_{\rm IGM} \end{cases} \quad (1)$$

where $r_{\rm galaxy}$, $r_{\rm CGM,\,min}$, $r_{\rm CGM,\,max}$, and $r_{\rm IGM}$ are the radii of the galaxy, inner boundary of the CGM, outer boundary of the CGM, and the IGM, respectively. The variable RF is the 'refinement factor', by which CGM gas is better refined with respect to $m_{\rm target}$, and hence better refined with respect to gas in the galaxy. These five values remain constant throughout the run. To avoid resolution jumps, at intermediate buffer regions between (a) the galaxy and the inner boundary of the CGM and (b) the outer boundary of the CGM and the IGM, the target mass is (linearly) interpolated.

In all our runs, we define $[r_{\rm galaxy}, r_{\rm CGM,\,min}, r_{\rm CGM,\,max}, r_{\rm IGM}] = [0.10, 0.15, 1.00, 1.50] \times R_{200,\,c}$, where $R_{200,\,c}$ is the $z = 0$ virial radius of the corresponding halo from TNG50-2. As a result, a constant comoving volume bounded by $r_{\rm CGM,\,min}$ and $r_{\rm CGM,\,max}$ is maintained at super-high resolution (i.e. a factor of RF better than $m_{\rm target}$) between the point the first black hole is seeded ($z \sim 3$–5 for the haloes in our sample) and $z = 0$. At all redshifts, we thus always define the CGM as the region between these two (pre-specified) radii.

We provide a visual illustration of our technique in Fig. 1. Each quadrant corresponds to a different resolution level of a given galaxy. The top-left corresponds to the halo from TNG50-2 (CGM gas mass resolution, $m_{\rm gas,\,CGM} \sim 8.5 \times 10^5M_\odot$[3]), the top right to the RF8 run of the corresponding halo of GIBLE ($m_{\rm gas,\,CGM} \sim 1.2 \times 10^5M_\odot$), the bottom right to RF64 ($m_{\rm gas,\,CGM} \sim 1.4 \times 10^4M_\odot$), and the bottom-left to RF512 ($m_{\rm gas,\,CGM} \sim 1.8 \times 10^3M_\odot$). Background colours show the average gas cell radius of gas in a ± 20 kpc depth projection. The two circles are drawn at radii $[0.15, 1.00] \times R_{200,\,c}$, the two boundaries of the CGM, and the circular arrow signifies the direction of improving numerical resolution, that is, clockwise from the top-left to bottom-left. While the background CGM, on average, shifts towards brighter colours along this arrow, signifying better numerical resolution throughout the halo, a greater amount of small scale structure is also apparent. The RF512 run resolves the denser portions of the CGM with cell radii of 10s of parsecs.

A more quantitative illustration of the scheme is provided in Fig. 2. Both panels show trends of numerical resolution as a function of galactocentric distance normalized by the virial radius. Curves from RF8, RF64, and RF512 are shown through orange, red, and black curves, respectively. The two grey curves correspond to TNG50-1 (solid) and TNG50-2 (dashed). The vertical lines portray the definitions we use for the radii of the galaxy, the two boundaries of the CGM and the IGM.

The left panel shows the mass resolution as a function of distance, which traces equation (1). By definition, at distances $<0.1\, R_{200,\,c}$, that is, within the galaxy, the mass resolution of all three resolution runs of GIBLE are roughly the same, and roughly equal to the resolution of TNG50-2.[4] The same holds for distances $>1.5\, R_{200,\,c}$, that is, beyond the IGM. Between $0.15\, R_{200,\,c}$ and $R_{200,\,c}$, that is, in the CGM, resolution of the RF8 run is similar to TNG50-1, although not exactly the same, since we only (de)refine gas cells if they are off by a factor of two with respect to the (distance dependent) target gas mass. The RF64 curve (dashed) is vertically offset by a factor of eight, and the RF512 curve (solid) further by another factor of eight. The three GIBLE curves converge in the two buffer regions, eventually overlapping each other in the galaxy and beyond the IGM. The two TNG50 curves show constant mass resolution at all radii, which is also what would be seen in a 'normal' zoom-in simulation, at least within the high-resolution region. Table 1 succinctly summarizes the various resolution levels realized by our three RF levels.

The right panel shows how spatial resolution changes with distance, for one of our eight simulated haloes. Similar to the left panel, the three GIBLE curves converge with TNG50-2 in the galaxy, and beyond the IGM. This corresponds to a median gas cell size of $\sim$450 pc in the galaxy, and $\gtrsim$10 kpc in the far reaches of the IGM. In the CGM, the spatial resolution progressively increases for the higher RF-runs: the RF8 curve (orange) is in the ballpark of TNG50-1 (solid grey), with a median cell size of $\sim$300 pc at $0.15\, R_{200,\,c}$, increasing gradually to $\sim$3 kpc at the virial radius. The red (RF64) and black (RF512) curves are better resolved by factors of $\sim$2 and 4, respectively, at all distances, resulting in a substantially improved median spatial resolution of $\sim$75 pc at $0.15\, R_{200,\,c}$ for the RF512 run.

The three shaded regions show three different percentile regions for the RF512 run: 5th–95th, 0.1th–95th, and 0.0001th–95th for the darkest to lightest bands, that is, the latter two bands portray the extreme outliers of the distribution. The 5th–95th region is relatively broad, as a result of multiphase gas at all distances: the halo averaged spatial resolution varies from $\sim$100 pc for the cold phase ($T < 10^{4.5}$ K) to $\sim$500 pc for the warm phase ($10^{4.5} < T < 10^{5.5}$ K), to $\sim$700 pc for the hot component ($T > 10^{5.5}$ K). The dips in the percentile regions seen at, for example $\sim 0.8\, R_{200,\,c}$, are a result of ISM gas of satellites being super-refined when they pass through the refinement region. The best resolved cell of this RF512 run is $\sim$3 pc ($n_H \sim 2000$ cm^{-3}), located in the dense ISM of an infalling satellite, and is shown by the dashed horizontal line at the bottom end of the panel. This is roughly 4 times smaller than the best resolved gas cell ($\sim$10 pc) of the TNG50 box (Pillepich et al. 2019).

2.3 A comparison with other projects

Before proceeding further, we compare our simulation(s) versus several other similar projects, including analogous CGM refinement techniques. Fig. 3 shows the (median) spatial resolution of simulated

[2]Note that, in principle, this factor can be set arbitrarily. A value of two is employed in the TNG runs.

[3]In TNG50-2, this resolution is maintained throughout the box.

[4]This resolution, that is, TNG100-1, is where the TNG model is calibrated versus observational constraints (Weinberger et al. 2017; Pillepich et al. 2018a), and hence the regime where the models for galactic feedback return the most realistic and robust results. That is, the energetics across this 'inner boundary condition' from the galaxy itself into the CGM are best when we maintain the galaxy at TNG50-2 resolution, as in our simulations.

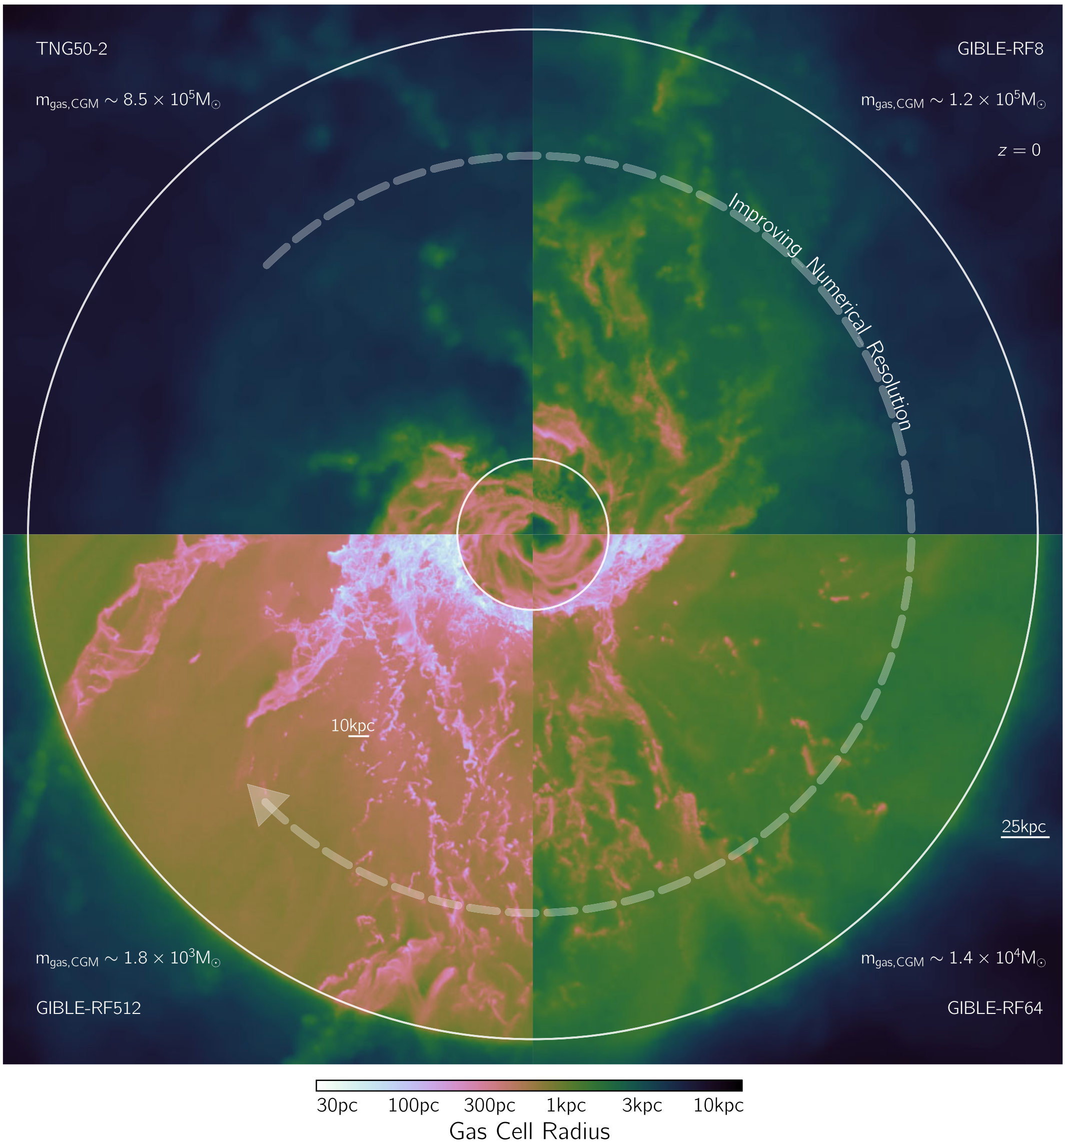

Figure 1. A qualitative illustration of our refinement scheme. Each quadrant shows a part of the same halo from four different simulations. The top left corresponds to TNG50-2, while the others are taken from GIBLE. Starting from the top-left, the (CGM gas) resolution increases in the clockwise direction. The image extends $\pm$ 1.05 $R_{200,\,c}$ along the plane, and $\pm$ 20 kpc perpendicular to the plane. Colours show the gas cell radius, that is, the spatial resolution, with the lighter pixels corresponding to better numerical resolution. At higher resolutions, finer structure is apparent on smaller scales, embedded in an altogether better resolved halo.

haloes as a function of galactocentric distance normalized by the virial radius. The left panel only includes cold ($T < 10^{4.5}$ K) gas, while the right panel shows results for all gas, with satellite gas excised in both cases. Note that the right panel shows only a subset of curves from the left.

The three green curves correspond to $z = 0$ Milky Way-like haloes selected from the three TNG boxes. These are all large-volume uniform-resolution boxes, run at different resolutions: TNG50-1 (dashed; $m_b \sim 8 \times 10^4$ M$_\odot$), TNG100-1 (dot-dashed; $m_b \sim 10^6$ M$_\odot$), and TNG300-1 (dotted; $m_b \sim 10^7$ M$_\odot$). The blue curves show examples of projects that aim to substantially improve numerical resolution: the dashed-blue curve corresponds to Auriga Level-2 (Grand et al. 2021), the (currently) best-resolved zoom-in simulation of a Milky Way-like galaxy, in which the mass resolution of gas is constant throughout the halo ($m_b \sim 800$ M$_\odot$). The dot-dashed curve shows CGOLS (Schneider & Robertson 2018), an idealized $10 \times 10 \times 20$ kpc^3 simulation of a $M_\star \sim 10^{10}$ M$_\odot$ disc galaxy simulated at a constant $\sim$5 pc spatial resolution throughout.

As opposed to the earlier simulations that maintain a fixed resolution, either mass- or spatial-, throughout the region of interest, the black curves show examples of CGM-refinement simulations, wherein resolution is preferentially increased within a refinement region that traces the CGM. The solid line shows one of the GIBLE-RF512 haloes. The ($z \sim 2$) 10^{12} M$_\odot$ halo from Suresh et al. (2019), where a similar refinement technique was used as GIBLE to attain a CGM gas mass resolution of $\sim$2200 M$_\odot$, is shown through the dashed curve. The dot-dashed curve shows TEMPEST (Hummels et al. 2019), a simulation of a Milky Way-like galaxy at $z = 0$ run with the grid code ENZO (Bryan et al. 2014), with enhanced spatial resolution of $\sim$500 pc within a 200 kpc box centred on the galaxy, and

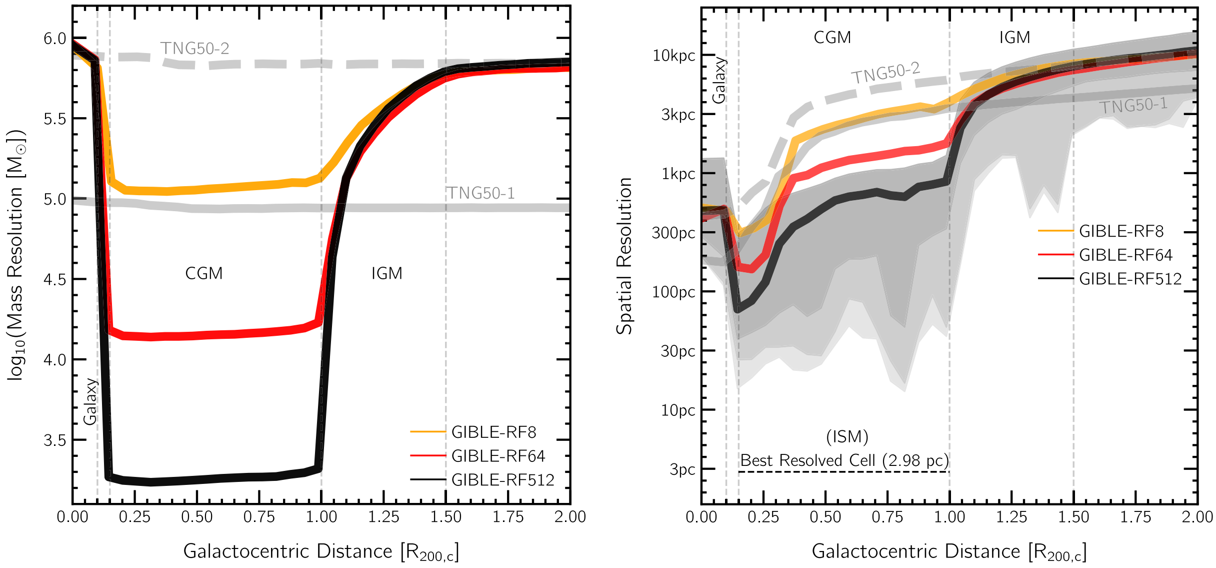

Figure 2. A schematic illustration of our refinement scheme. Orange, red, and black curves correspond to RF8, RF64, and RF512, respectively, while the two grey lines correspond to TNG50. In the right panel, we show three different percentile regions for the RF512 run: 5th–95th, 0.1th–95th, and 0.0001th–95th for the darkest to lightest bands, that is, the latter two bands portray the extreme outliers of the distribution. The dips seen in the shaded bands, as well as the best resolved cell in this run (~3 pc; horizontal dashed line), correspond to the super-refined ISM gas of satellites as they pass through the host CGM. We show the radial profile of mass resolution on the left, and spatial resolution on the right. The different dashed vertical lines signify the different boundaries used to define the refinement region. The distance dependent target mass of our CGM refinement scheme is used to achieve a boosted mass/spatial resolution in the CGM region.

Table 1. Characteristic numbers associated with our simulations: the first column lists the refinement factor (RF), that is, the factor by which CGM gas is better resolved compared to the galaxy. The second column specifies the average gas mass resolution in the galaxy, which is roughly the same for all runs. The average CGM gas mass resolution is shown in the third column, which progressively improves from top to bottom. The average computational time for our sample of MW-like haloes run to $z = 0$ is shown in the last column.

Refinement Factor	$m_{\mathrm{gas,\,galaxy}}$ [$M_\odot$]	$m_{\mathrm{gas,\,CGM}}$ [$M_\odot$]	Compute time [10^6 CPU h]
8	$\sim 8.5 \times 10^5$	$\sim 1.2 \times 10^5$	~0.05
64	$\sim 8.5 \times 10^5$	$\sim 1.4 \times 10^4$	~0.1
512	$\sim 8.5 \times 10^5$	$\sim 1.8 \times 10^3$	~0.5

coarser resolution at farther distances. While the first set of FOGGIE simulations used a similar technique as TEMPEST (Peeples et al. 2019), the latest runs additionally refine based on a cooling length criterion, reaching a spatial resolution of ~275 cpc in the cold phase (e.g. Lochhaas et al. 2023); this FOGGIE run is shown using dashed lines with alternate long and short dashes. The halo from van de Voort et al. (2019) is shown through the dotted curve. This is a zoom-in simulation of a Milky Way-like galaxy run with AREPO at a base resolution of $m_b \sim 5 \times 10^4 M_\odot$, that is, similar to TNG50-1, but with added spatial refinement that sets a maximum spatial resolution for gas of 1 kpc.

As expected, the zoom-in simulations achieve much better resolutions in comparison to the large-volume boxes, although at the expense of a smaller sample size. With the CGM-refinement technique, it is possible to simulate a relatively large sample of galaxies at super-high (CGM gas) resolution. For instance, the current sample of GIBLE is composed of eight haloes that are comparable in (CGM) resolution to Auriga Level-2, of which there is only one galaxy.

In the case of CGM refinement simulations, the choice of the refinement technique is important, since it has a preferential impact on the numerical resolution of a particular phase of gas. For instance, Peeples et al. (2019) show that at a constant spatial resolution of ~500 cpc, the mass resolution can vary by many orders of magnitude, between $\lesssim 100\,M_\odot$ for the hot phase and $\gtrsim 10^5\,M_\odot$ for the colder phase. As discussed in, for example Lochhaas et al. (2023); Smith et al. (2023), good resolution in the hot phase can help resolve the effects of turbulence in and around hot outflows. Conversely, at fixed mass resolution, as is the case with GIBLE, the spatial resolution of the cold phase is better with respect to the warm-hot component: the median spatial resolution of cold gas at RF512 is $\lesssim 500$ pc throughout the halo (left panel of Fig. 3), while the median is ~1 kpc at $R_{200,\,c}$ once the other phases of gas are included (right). Subsequently, the gas distribution in dense regions is sampled by a larger number of resolution elements, which is important to improve the accuracy of interactions with the ambient medium (e.g. Goldsmith & Pittard 2016). In addition, resolving all phases is important since the evolution of multiphase gas is influenced by the interactions between these different phases (Fielding & Bryan 2022).

While the dependence of mass resolution on temperature at fixed spatial resolution is discussed earlier, we here briefly consider the other case, that is, the impact of added spatial resolution on a simulation in which refinement is also performed based on mass. This is seen when the van de Voort et al. (2019) curves are contrasted against TNG50, both of which have similar mass resolutions, but the former includes added spatial refinement. The TNG50 and van de Voort et al. (2019) curves are comparable out to large distances in the left panel, since the size of cold gas cells at these distances is already smaller than or comparable to the 1 kpc enforced maximum. However, the difference between the two is much more pronounced once the warm- and hot-phases of gas are included (right), wherein all gas beyond ~0.3 $R_{200,\,c}$ is significantly better resolved.

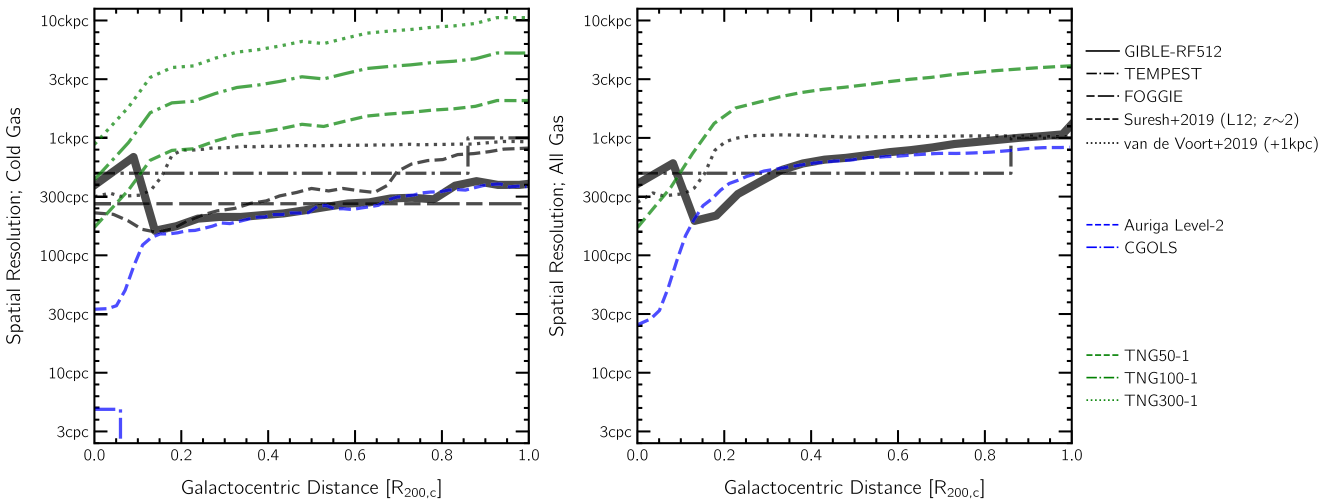

Figure 3. A comparison between our simulation(s) and other projects through radial profiles of spatial resolution. The three green curves correspond to the three TNG boxes at $z = 0$, that is, large volume uniform-resolution simulations. The blue curves are examples of projects that simulate a smaller volume at high resolution, either in a cosmological (Auriga; dashed) or an idealized (CGOLS; dot-dashed) context, but with uniform resolution, either mass or spatial, within the region of interest. The black curves show those that employ a CGM-refinement technique, all run as cosmological zoom-in simulations, with a further improvement in resolution within a refinement region that traces the CGM. The left panel shows results for cold gas alone, while the right panel includes all gas. The axes ranges of both panels are identical.

3 RESULTS

3.1 An overview of global properties

We begin with a tabulation of several important global properties of our sample (at $z = 0$) in Table 2. The first column shows the different halo IDs (in TNG50-2) from which initial conditions were derived. Each halo is split into three sub-rows, corresponding to the different RF choices, shown in the second column. The rest of the table is split into two parts. The left half shows quantities not directly associated with CGM gas: the first four columns give key properties used to describe galaxies and haloes: from left to right, these are the halo virial mass ($M_{200,\,c}$), halo virial radius ($R_{200,\,c}$), galaxy stellar mass[5] ($M_\star$), and the mass of the central supermassive black hole ($M_{\rm BH}$). Note that the three masses are shown in $\log_{10}$ units. The next column briefly summarizes the merger history of these galaxies, with the three values corresponding to the number of major, minor and all mergers since $z = 2$. Major mergers are defined as those with merger mass ratios $\mu > 0.25$, while minor mergers are those with $0.1 < \mu < 0.25$ (for more details, see Rodriguez-Gomez et al. 2016).

As galaxies in our sample are selected based on their similarity to the real Milky Way, halo-, and galaxy stellar-masses are confined to relatively narrow bins, that is, by construction (see Section 2). As a result, black hole masses and the merger histories of these galaxies are not too dissimilar as well (see Pillepich et al. 2023 for details on the level of variance expected within such a constrained sample).

The next part of the table shows values related to CGM gas: the total CGM gas mass ($M_{\rm g,CGM}$), mass fraction of cold gas in the CGM ($f_{\rm m,\,CGM,\,cold}$), number of discrete, cold CGM gas clouds ($N_{\rm cld,\,CGM}$), average temperature of CGM gas ($\overline{T_{\rm CGM}}$),[6] and volume fractions of CGM gas split into three different phases ($f_{\rm vol,\,CGM}$), respectively, from left to right. Gas is split into three phases based on temperatures cutoffs: cold ($T < 10^{4.5}$ K), warm ($10^{4.5} < T < 10^{5.5}$ K), and hot ($T > 10^{5.5}$ K). Following Ramesh et al. (2023b), we define cold gas clouds as contiguous sets of cold gas cells, restricting to those cells that are not bound to any satellites (see also Nelson et al. 2020). We only consider clouds composed of at least 10 cells.

Despite a relatively narrow bin in halo masses, properties of CGM gas show a large variation across the sample of eight galaxies (see also Ramesh et al. 2023a). For instance, the coldest halo in the sample has an average temperature of $\sim 10^{4.8}$ K, more than an order-of-magnitude cooler than the hottest ($\sim 10^{5.9}$ K). This results in a large variation of cold gas mass fractions, from as low as $\lesssim 1$ per cent to as high as ~ 50 per cent.

To visually demonstrate the effect of numerical resolution, we plot a subset of these quantities in Fig. 4. The left half of the figure shows the first four columns of the first part of Table 2, while the first four columns of the second part of the table are shown on the right panel. Each row corresponds to a single-physical quantity. Both panels are split into eight shaded regions, with each colour corresponding to a given halo. Runs at RF8, RF64, and RF512 resolutions are shown through circles, squares, and pentagons, respectively.

Overall, halo mass shows the least variation with resolution, differing by at most ~ 0.03 dex between runs of different resolutions. Galaxy stellar masses and BH masses show larger variations ($\lesssim \pm 0.1$ dex), as do some of the gas quantities, like CGM gas mass ($\lesssim \pm 0.2$ dex), cold gas mass fractions ($\lesssim$20–30 per cent), and average halo temperatures ($\lesssim \pm 0.3$ dex). However, no global monotonic trend is apparent for any of these quantities. This suggests that such variations are likely a result of the randomness of the butterfly effect (Genel et al. 2019) and/or the zoom-timing effect (Springel et al. 2008), and not directly sensitive to changing numerical resolution.

The one quantity for which a clear resolution trend is visible is the number of discrete, cold CGM gas clouds. For all but the first halo (S91), a greater number of clouds are found at higher resolution. As discussed in Ramesh et al. (2023b), this is a direct result of the existence of a resolution limit for every simulation. Structures

[5]Computed through the sum of masses of stars within an aperture of 30 kpc with respect to the centre of the galaxy.

[6]The temperature of star-forming gas in the TNG model represents the effective state of the two phases. To overcome this limitation, we always set the temperature of star-forming gas to 10^3 K, its cold phase value.

Table 2. A summary of selected global properties of our sample at $z = 0$. The first two columns show the halo number and RF. The shapes next to the halo numbers reflect the selection criteria used to choose these galaxies, as shown by the legend (see also Section 2) The next five columns show quantities not directly related to CGM gas: halo mass, virial radius, galaxy stellar mass, central supermassive black hole mass, and number of mergers since $z = 2$, split based on merger mass ratio μ ($\mu > 0.25$: major, $0.1 < \mu < 0.25$: minor). The last five columns show quantities related to CGM gas: total CGM gas mass, cold gas mass fraction, number of cold CGM gas clouds, average CGM gas temperature, and volume fractions of the three phases of CGM gas. Refer to main text for details (Section 3.1).

Halo number	RF	$M_{200,c}$ [log $M_\odot$]	$R_{200,c}$ [kpc]	$M_\star$ [log $M_\odot$]	M_{BH} [log $M_\odot$]	$N_{merger,\, z<2}$ ma., mi., A	$M_{g,CGM}$ [log $M_\odot$]	$f_{m,\, CGM}$ [C; per cent]	$N_{cld,\, CGM}$ ($\geq$ 10 cells)	$\overline{T_{CGM}}$ [log K]	$f_{vol,\, CGM}$ [per cent] C, W, H
S91*	8	12.18	242.74	10.57	8.40	0, 0, 16	10.12	0.75	4	5.95	00.19, 18.37, 81.44
	64	12.18	243.23	10.61	8.33	0, 0, 18	10.31	0.07	7	5.90	00.10, 95.60, 04.30
	512	12.19	243.28	10.53	8.32	0, 0, 21	10.05	0.59	36	5.87	00.24, 29.19, 70.56
S98$^\triangle$	8	12.31	266.84	10.72	8.33	2, 0, 15	11.13	47.51	345	4.80	12.44, 82.29, 05.26
	64	12.32	270.62	10.82	8.30	1, 0, 15	11.20	51.24	2831	4.79	13.10, 69.59, 17.31
	512	12.33	271.19	10.84	8.28	1, 0, 17	11.20	47.60	18152	4.86	14.11, 76.80, 09.09
S105$^\diamond$	8	12.28	261.91	10.63	8.13	1, 2, 09	11.08	16.79	510	5.29	06.50, 85.02, 08.48
	64	12.29	263.46	10.65	8.07	1, 0, 14	11.12	16.79	2793	5.34	04.47, 95.16, 00.37
	512	12.29	263.34	10.65	8.10	2, 3, 07	11.12	22.30	14998	5.22	05.32, 91.58, 03.10
S146$^\spadesuit$	8	12.19	245.03	10.74	8.11	0, 0, 07	11.08	32.03	310	5.13	05.82, 85.62, 08.57
	64	12.16	239.48	10.74	8.17	0, 0, 08	10.81	37.17	600	5.21	01.93, 23.76, 74.31
	512	12.21	247.74	10.78	8.09	0, 0, 15	11.11	40.13	11551	4.97	07.84, 91.21, 00.95
S167$^\spadesuit$	8	12.17	242.12	10.85	8.17	0, 1, 20	10.89	8.39	238	5.53	03.11, 96.12, 00.77
	64	12.16	239.23	10.80	8.22	0, 1, 16	10.95	14.44	2076	5.59	06.75, 83.40, 09.86
	512	12.16	239.19	10.92	8.23	0, 0, 17	10.90	5.04	6218	5.65	00.53, 99.16, 00.31
S201$^\diamond$	8	12.07	223.39	10.61	8.14	0, 0, 08	10.84	2.41	126	5.68	00.11, 99.86, 00.03
	64	12.07	222.11	10.72	8.08	0, 0, 08	10.75	13.46	1534	5.73	03.97, 55.03, 41.00
	512	12.04	218.37	10.65	8.09	0, 0, 09	10.49	18.47	3491	5.48	00.53, 99.16, 00.31
S221$^\clubsuit$	8	11.97	206.19	10.58	7.79	0, 0, 12	10.72	18.52	47	5.38	00.72, 99.20, 00.07
	64	11.98	208.47	10.62	8.12	0, 0, 08	10.77	9.70	803	5.51	00.93, 98.92, 00.15
	512	11.98	207.87	10.70	7.83	0, 0, 10	10.70	18.27	7895	5.45	04.78, 89.42, 05.80
S264$^\triangle$	8	11.95	203.88	10.67	7.96	0, 0, 11	10.59	10.74	70	5.49	02.09, 97.89, 00.02
	64	11.95	203.36	10.69	8.00	0, 0, 09	10.44	46.97	306	4.95	10.00, 78.34, 11.66
	512	11.94	200.93	10.74	8.10	0, 0, 10	10.25	36.68	1784	5.14	12.58, 66.40, 21.02

Note. * : M31-like companion; $^\triangle$: *Fermi*-like bubble; $^\diamond$: Bar and/or stellar disc length; $^\spadesuit$: SMC/LMC-like satellite; $^\clubsuit$: Random selection.

smaller than the resolution limit are simply absent. We return to this discussion in Fig. 9.

In Fig. 5, we further assess numerical convergence of overall-halo gas properties by studying PDFs. For each halo, we compute averages over ten snapshots (~300 Myr) starting at $z = 0$. Time-averaging over such a relatively short interval helps suppress the impact of the butterfly effect and the timing problem discussed earlier, while also ensuring that the dynamical evolution of the halo does not play a major role. In each panel, the solid curves show the median of each time averaged halo-curve, while the shaded regions shown the variation across the sample. Results from RF8, RF64, and RF512 are shown in orange, red, and black, respectively.

From left to right, panels show the PDFs of gas temperature, metallicity, magnetic field strength, and radial velocity, respectively. In each case, the medians appear to be converged with resolution, at least at the level of resolution considered here. However, as we will explore throughout the remainder of the paper, this does not imply that all quantities related to CGM gas are necessarily converged, especially those of small-scale properties, as has already been seen with the increase in number of cold clouds with improving resolution, despite the PDFs of gas temperature being well converged.

3.2 H I covering fractions of halo gas

The covering fraction of a given species is a commonly used observational metric to quantify the distribution of halo gas when viewed in 2D projection. Here, we make predictions for the covering fractions (at $z = 0$) of commonly studied ions, and explore their sensitivity to numerical resolution.

We do not track the abundances of specific metal species in GIBLE, and so assume that hydrogen contributes 76 per cent to the mass of each gas cell, and adopt solar abundance for all other metal species. When estimating the fraction of H I that comprises the (neutral) hydrogen-mass of the cell, we use the model of Gnedin & Kravtsov (2011) to excise the contribution of H_2 (following Popping et al. 2019). We use CLOUDY (Ferland et al. 2013, v13.03) to model the abundances of various ionization states of the different metal species (following Nelson et al. 2018b; Ramesh et al. 2023a).

In Fig. 6, we consider the covering fractions of H I (y-axis) as a function of (2D) impact parameter (normalized by $R_{200,c}$; x-axis), a ubiquitous tracer of cold gas. The main panel considers three commonly used column density thresholds for H I: $> 10^{16}$ cm^{-2} (solid; pLLSs), $> 10^{19.6}$ cm^{-2} (dashed; sub-DLAs), and $> 10^{20.3}$ cm^{-2} (dotted; DLAs). Results from the RF8, RF64, and RF512 runs are shown in orange, red, and black, respectively. Each galaxy is considered in a random orientation, and curves show median trends of the sample. For the $> 10^{16}$ cm^{-2} case alone, we show the halo-to-halo variation through the 16th–84th percentile shaded regions.

For a $N_{H\,I}$ threshold of 10^{16} cm^{-2}, median covering fractions are close to unity at the inner boundary of the CGM (0.15 $R_{200,c}$), for all three resolution runs. For the best resolved run (RF512; black), there is a steady decrease to ~0.4 at 0.5 $R_{200,c}$, followed by a steeper drop to ~0.005 at the virial radius. At this $N_{H\,I}$ threshold, results from the RF64 run (red) are more or less consistent with the black curve, except in the innermost region of the halo. For the RF8 run, however, covering fractions are systematically lower by ~0.1–0.3 dex with respect to the black curve out to an impact parameter of ~0.9 $R_{200,c}$. Beyond this distance, in the outermost regions of

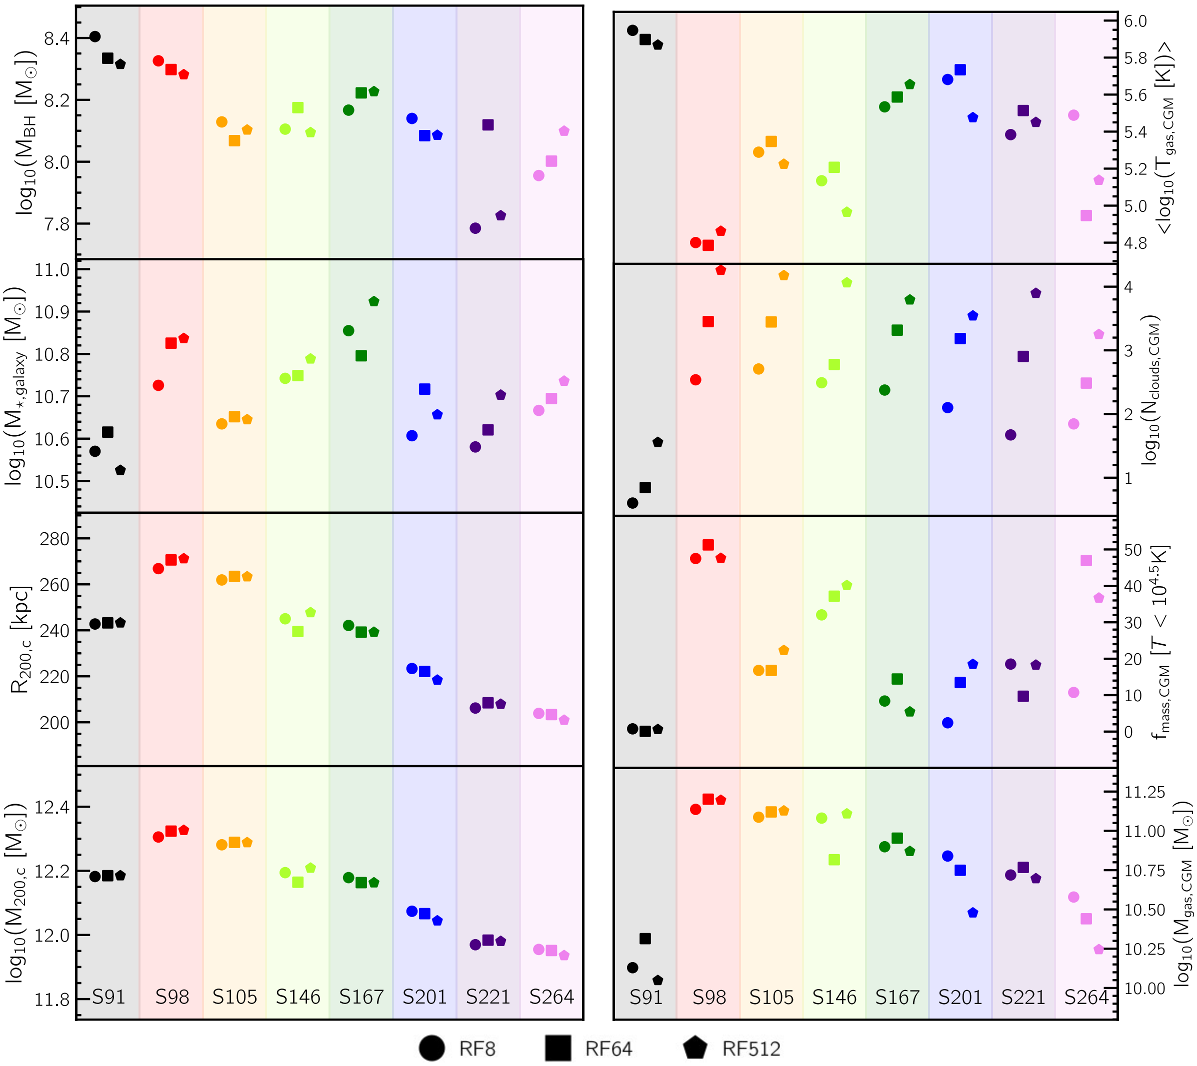

Figure 4. Resolution trends for physical properties of our simulated galaxies and their gaseous haloes, taken as a subset from Table 2. Each colour corresponds to a given halo. Runs at RF8, RF64, and RF512 are represented using circles, squares, and pentagons, respectively. While all quantities fluctuate with changing resolution, due to inherent stochasticity and timing offset effects, no global resolution trend is present, except for the case of CGM gas clouds.

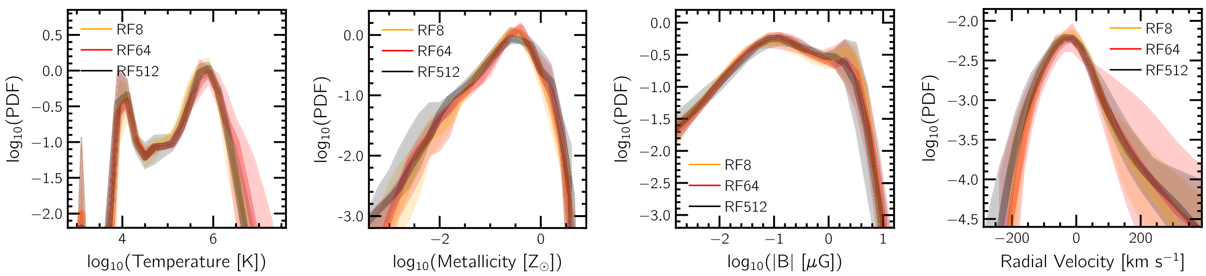

Figure 5. Probability distribution functions (PDFs) of four different properties of CGM gas. For each halo in our sample, we compute (time) averages over ten snapshots (∼300 Myr) starting at $z = 0$. The solid curves show the median of these time-averaged curves, while the shaded regions show the 16th–84th percentile variation of the sample. Results of the RF8, RF64, and RF512 runs are shown in orange, red, and black, respectively. Although these time- and sample-averaged PDFs are converged with resolution, this does not necessarily imply that all results are also converged, as explored in detail throughout the rest of the paper.

the halo, all three curves are converged. The overlap of the black and red curves demonstrates that the H I covering fraction, at this column density threshold, is numerically converged already at our RF64 resolution. We note that the resolution of RF8 is comparable to TNG50-1 (in the CGM), and results derived from TNG50-1 are largely comparable to the ones from RF8 here (not explicitly shown). This suggests that the H I covering fraction for Milky Way-like haloes at $z = 0$ above a low threshold of $N_{\rm HI} > 10^{16} \rm cm^{-2}$ is not

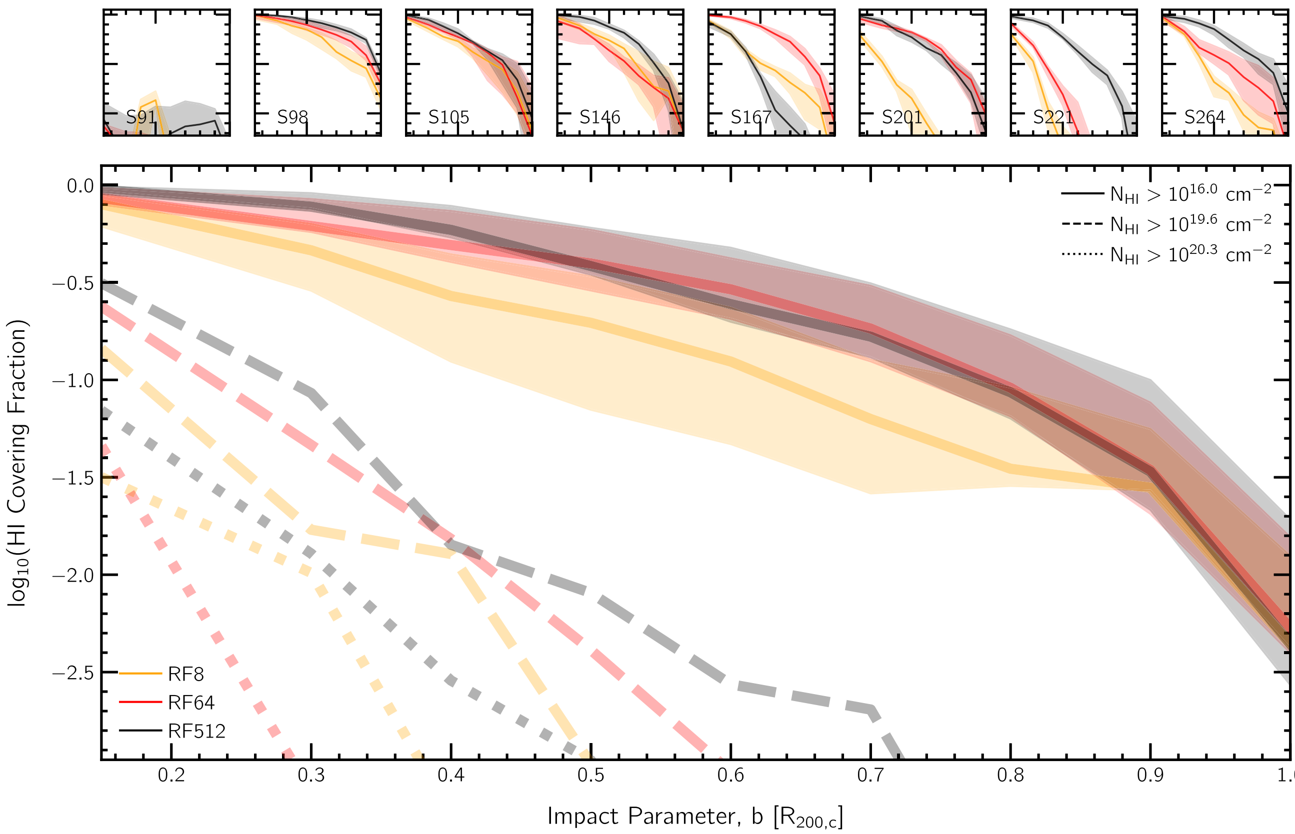

Figure 6. Trends of H I covering fractions as a function of impact parameter. In the main panel, we show median curves for the sample, split based on numerical resolution (different line colours), and $N_{\rm H\,I}$ cutoffs (different line types). Shaded regions, shown only for the 10^{16} cm^{-2} case, represents the halo-to-halo variation across the sample. Median covering fractions decrease with distance, and are lower for higher column density thresholds. Increased numerical resolution generally leads to higher covering fractions, except at the high column density end, where the impact of resolution is superseded by other stochastic effects. The small panels in the top row show the radial profiles of H I covering fraction for the eight individual haloes in our sample. Both the axes labels and the axes ranges of these eight stamps are identical to those of the main panel. Refer to main text for details.

yet fully converged at TNG50-1 resolution, depending on impact parameter.

As one transitions towards higher $N_{\rm H\,I}$ thresholds, the dependence on resolution becomes stronger. At a cutoff of $10^{19.6}$ cm^{-2} (dashed lines), the median covering fraction of the RF512 run is ∼0.3 at 0.15 $R_{200,\rm c}$, and is over two orders of magnitude smaller (∼0.002) already at 0.7 $R_{200,\rm c}$. Columns of gas at such density thresholds are thus rare to find at large impact parameters. For the RF64 run, the median is lower by ∼0.1 dex in the inner halo ($\lesssim$0.3 $R_{200,\rm c}$) with respect to RF512, and the offset grows beyond ∼0.5 $R_{200,\rm c}$. The RF8 run has even lower converting fractions, demonstrating that dense gas (at this $N_{\rm H\,I}$ threshold) is effected by numerical resolution, and results are not well converged, at least up to a resolution of $\sim10^4$ M$_\odot$.

For $N_{\rm H\,I} \gtrsim 10^{20.3}$ cm^{-2} (dotted curves), median covering fractions drop further, for all resolution levels. For instance, in the RF512 runs, the median covering fraction is lower by ∼0.6 dex with respect to the black-dashed curve. The red and orange curves are again vertically offset with respect to the black curve, although no monotonic trend with resolution is present here. This suggests that, at column density thresholds as high as $10^{20.3}$ cm^{-2}, the impact of resolution is superseded by other effects, possibly by the choice of a random orientation for each galaxy, or due to the randomness of stochastic processes, as discussed earlier.

In the top panels, we show a collection of eight stamps, each corresponding to a given halo. Each stamp contains three curves, corresponding to the three different resolution runs, with the same colour scheme as the main panel. For visual clarity, we show only the case with a $N_{\rm H\,I}$ threshold of 10^{16} cm^{-2}. For each galaxy, we pick 100 random orientations, and show the median with solid lines. Shaded regions represent the 16th–84th percentile variation of the different orientations. As seen, at least at this column density threshold, the choice of orientation does not play a dominant role, with percentile regions being generally narrow ($\lesssim$0.1 dex).

While most individual haloes are consistent with the findings of the main panel, there are exceptions. For instance, in the fourth case from the left, the RF64 curve lies below the RF8, while the fifth halo gas the RF512 curve below the other two. This variability demonstrates that it is impossible to accurately assess the impact of resolution based on a single halo, and a sample of multiple haloes is needed to overcome uncertainties related to effects including numerical stochasticity and/or timing offsets.

Adding further complication, the H I content of gas in simulations is sensitive to the (possibly self-shielded) photo-ionization rate. Using a semi-empirical model, Faerman & Werk (2023) make predictions for the radial profile of H I column densities. In addition, they study the effect of varying the amount of non-thermal pressure support of cold gas (η). Greater values of η correspond to higher rates of photo-ionization. Out to impact parameters of ∼0.5 $R_{200,\rm c}$, they find that their results are sensitive to the chosen value of η. For instance, at an impact parameter of 0.2 $R_{200,\rm c}$, their low- (high-)η

case predicts a median H I column density of $10^{17.5}$ ($10^{16.5}$) cm^{-2}, as a result of higher photo-ionization rates leading to a removal of H I. The choice of the (sub-grid) self-shielding model employed for dense gas in cosmological simulations is thus important, and predictions of quantities related to H I could possibly be sensitive to this choice. However, we have explicitly checked that trends of H I covering fraction are qualitatively similar to those of H covering fractions (not shown), suggesting that the self-shielding recipe employed by TNG (Rahmati et al. 2013) does not bias results.

To provide an idea of the typical H I-densities and trends of covering fractions recovered from both observations and simulations, we quote some numbers from current literature. Note that most of these are not apple-to-apple comparisons, that is, the galaxy sample selections, redshift ranges, and other details are not replicated, and are intended to be taken at face value only. Observationally, for a set of 32 low-z COS-Haloes galaxies, Prochaska et al. (2017) find that the median $N_{\rm H\,I}$ is $\gtrsim 10^{17.5}$ cm^{-2} at impact parameters <60 kpc, which is consistent with our RF64 and RF512 curves that predicts a covering fraction close to unity for pLLS systems at impact parameters $\lesssim 0.3\ R_{200,\,\rm c}$. For a sample of 32 strong Ly α absorbers in the redshift range $0.2 \lesssim z \lesssim 1.4$, Weng et al. (2023) too infer a trend of decreasing $N_{\rm H\,I}$ with increasing impact parameter, with values dropping from $\sim 10^{19.5}$ cm^{-2} at 0.15 $R_{200,\,\rm c}$ to $\sim 10^{18}$ cm^{-2} at the virial radius. Using a sample of 32 $z = 2$ $N_{\rm H\,I} \simeq 10^{17.2}$ cm^{-2} absorbers, Rubin et al. (2015) find that the covering fraction is $\sim 0.6 - 0.8$ at impact parameters $\lesssim$150 kpc, and drops to $\sim$0.2 at larger distances. Rudie et al. (2012) find that, for a sample of 886 $z \sim 2$–3 galaxies, the covering fraction for $N_{\rm H\,I} > 10^{16}$ cm^{-2} absorbers varies between 0.6 in the inner most part of the halo and $\lesssim$0.1 at larger impact parameters. These trends of decreasing covering fractions of H I with increasing impact parameter is qualitatively consistent with our GIBLE curves, suggesting that the CGM refinement technique does not introduce any noticeable spurious artefact, at least in the case of H I covering fractions.

Using standard cosmological zoom simulations, Hani et al. (2019) also study the distribution and covering fractions of H I and various ionic species. They adopt a sample of 40 $z = 0$ MW-like galaxies from the Auriga project simulated with an average gas resolution of $\sim 5 \times 10^4$ M$_\odot$, that is, at a resolution roughly between our RF8 and RF64 runs. For $N_{\rm H\,I} \gtrsim 10^{14.15}$ cm^{-2} absorbers, they find that the distribution of covering fractions is large, for sightlines stacked between 0.2 $R_{200,\,\rm c}$ and $R_{200,\,\rm c}$: while the distribution peaks at $\sim$0.65, values range between $\sim$0 and $\sim$0.8. They find that this large variation is a result of the $\sim$1 dex scatter in stellar mass, and a wide distribution of disc fractions and AGN luminosity across their sample.

The impact of resolution on H I covering fractions has also been explored by the other 'CGM refinement' projects mentioned earlier. For a set of $z \sim 2$ MW-like progenitors from the FOGGIE simulation project, Peeples et al. (2019) find that their method leads to higher covering fractions for column densities $10^{15} \lesssim N_{\rm H\,I} \lesssim 10^{17}$ cm^{-2}. Similarly, van de Voort et al. (2019) show that increased spatial refinement leads to larger H I covering fractions for $10^{14} \lesssim N_{\rm H\,I} \lesssim 10^{19}$ cm^{-2}, for a single MW-like analogue at $z = 0$, simulated at the same resolution as the Auriga sample discussed earlier, that is, $m_b \sim 5 \times 10^4$ M$_\odot$ (but with added spatial refinement for a subset of runs). Alternatively, our results show decent convergence (for RF64 and above) in the low $N_{\rm H\,I}$ end ($\lesssim 10^{19}$ cm^{-2}). This difference is likely a consequence of the difference in numerical methods employed to run these simulations (see also Section 2.3). Moreover, as previously discussed in Suresh et al. (2019), the 'base' resolution plays an important role in the comparison to van de Voort et al. (2019), that is, the mass resolution at which they run the zoom-in simulation on top of which the enhanced CGM spatial refinement is incorporated. In our RF64 and RF512 runs, the median spatial resolution of cold gas is already below 1 kpc, which is the point at which van de Voort et al. (2019) find a difference with their added spatial resolution. This suggests that covering fractions at the low $N_{\rm H\,I}$ end are already converged at a mass resolution of RF64, and the added 1 kpc spatial refinement (as used by van de Voort et al. 2019) only makes a difference if the base resolution is below this.

3.3 Metal covering fractions

In Fig. 7, we explore trends of covering fractions of commonly used metal ion tracers of multiphase CGM gas (see e.g. Tumlinson, Peeples & Werk 2017): Mg II (cold gas; left panel), C IV (warm gas; centre panel), and O VII (hot X-ray emitting gas; right panel). Each panel shows the covering fraction of the corresponding ion as a function of impact parameter (normalized by $R_{200,\rm c}$). The medians of the three different resolution runs are shown through different colours: orange (RF8), red (RF64), and black (RF512). The different line styles correspond to different column density thresholds, as elaborated further.

For a column density threshold of $N_{\rm Mg\,II} \gtrsim 10^{14}$ cm^{-2} (solid curves, left panel), the median covering fraction at $\sim$0.15 $R_{200,\rm c}$ is close to unity for the RF512 run, which reduces steadily to $<10^{-3}$ at the virial radius. Values of RF64 are consistent with RF512 at all distances, but the RF8 curve is vertically offset to lower covering fractions, almost at all distances. Covering fractions reduce with increasing thresholds on column density. At $N_{\rm Mg\,II} \gtrsim 10^{15}$ cm^{-2}, (dashed) curves are vertically offset by $\sim$1 dex in comparison to the solid curves. Surprisingly, decent convergence is seen between resolution levels at this threshold, although curves are clearly noisy. For a higher threshold of $N_{\rm Mg\,II} \gtrsim 10^{16}$ cm^{-2}, shown through dotted curves, covering fractions are over two orders of magnitudes smaller in comparison to the solid curves. No consistent trend with resolution is seen here, much like the H I case, which is expected since both these species largely trace the same gas.

These trends and results are largely in qualitative agreement with past observational and theoretical work. For instance, using a large sample of $\sim$160 000 Mg II absorbers in the redshift range $0.3 \lesssim z \lesssim 2.3$, Anand, Nelson & Kauffmann (2021) study the covering fractions in the CGM around luminous red galaxies (LRGs) and emission line galaxies (ELGs). At an equivalent width (EW) threshold of 0.4 Å, they find that the Mg II covering fractions of LRGs (ELGs) is $\sim$0.15 (0.6) at an impact parameter of 20 kpc, and drops to $\sim$0.03 (0.04) at 400 kpc. At a higher EW threshold of 2.0 Å, these numbers reduce to $\sim$0.05 (0.1) at 20 kpc, and $\sim$0.001 (0.003) at 400 kpc. With a set of 16 LRGs at intermediate redshift ($0.21 \lesssim z \lesssim 0.55$), Zahedy et al. (2019) estimate the Mg II covering fraction to be 0.4–0.7 (0.0–0.2) for impact parameters $\lesssim$100 kpc (100–160 kpc), for a column density threshold of $N_{\rm Mg\,II} \gtrsim 10^{13}$ cm^{-2}.

Using the EAGLE simulations, Ho, Martin & Schaye (2020) studied the covering fractions around a set of $z \sim 0.3$ galaxies. For a threshold of $N_{\rm Mg\,II} \gtrsim 10^{11.5}$ cm^{-2}, they find that the Mg II covering fraction at $10^{10.5} \lesssim M_\star/{\rm M}_\odot \lesssim 10^{11.0}$ is $\sim$0.55 (0.2) at an impact parameter of 50 (150) pkpc. Moreover, they find that when galaxies are oriented edge-on, sightlines closer to the minor axis have larger Mg II covering fractions, as a result of the directional nature of outflows (e.g. Péroux et al. 2020; Ramesh et al. 2023c; Weng et al. 2024). In addition, they study the effect of the star-formation status of the galaxy. At this mass range, Mg II covering fractions do not correlate with SFR, but in galaxies with smaller stellar masses,

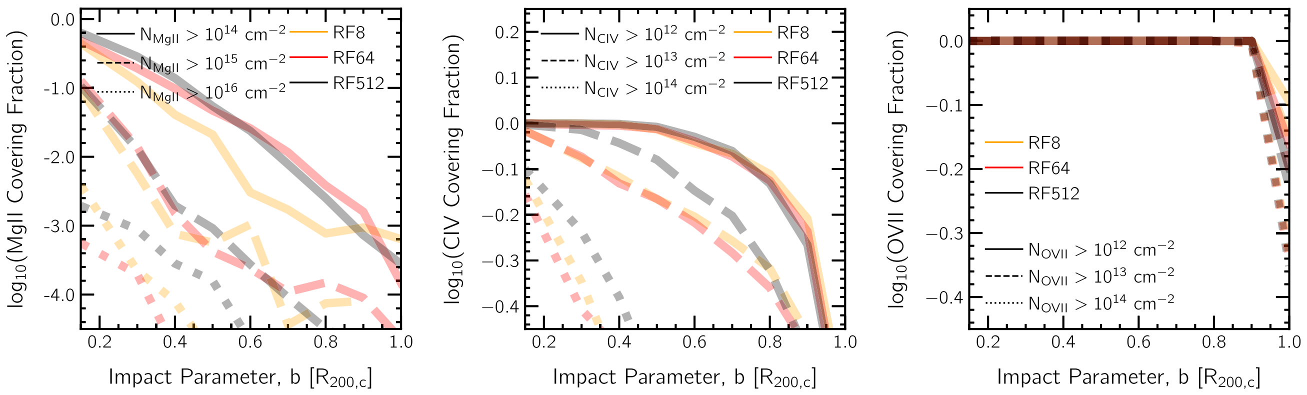

Figure 7. Trends of covering fractions of Mg II (left), C IV (centre), and O VII (right) as a function of impact parameter. In each panel, the colours represent the different resolution runs, while line styles correspond to different column density thresholds for each ion. Overall, the hot phase, traced by O VII, is the best converged (numerically) in our simulations, while the convergence of the other two ions depend on the column density thresholds employed.

covering fractions are typically smaller when the central galaxy is quenched.

Nelson et al. (2020) study the Mg II covering fractions for a sample of TNG50 haloes in the range $10^{13.2} \lesssim M_{200,c}/M_{\odot} \lesssim 10^{13.8}$, stacked at three snapshots ($z = [0.4, 0.5, 0.6]$). They compare TNG50-1 with –2, –3, and –4 (8, 64, and 512 times lower mass resolution, respectively), and find that the covering fraction monotonically decreases with coarsening resolution, for a $N_{Mg\,II}$ cutoff of 10^{16} cm^{-2}. At this threshold, we find no monotonic trend with resolution with GIBLE. This suggests that effects of numerical resolution on Mg II covering fractions are more important at the resolution of TNG50-1 (i.e. RF8) and below, and hence less apparent apparent when comparing between the higher resolution RF64 and RF512 runs.

Covering fractions of C IV, at a column density threshold of $N_{C\,IV} \gtrsim 10^{12}$ cm^{-2} (solid curves, centre panel) are relatively flat near unity between the inner boundary of the CGM (~0.15 $R_{200,c}$) and ~0.5 $R_{200,c}$. At larger impact parameters, the value drops steadily, but is still as large as ~0.2 at the virial radius. Numerical convergence, at this threshold, is good between the different runs. For $N_{C\,IV} \gtrsim 10^{13}$ cm^{-2} (dashed curves), the covering fractions for the RF512 run are offset by ~0.1 dex with respect to the solid curve, except at impact parameters $\lesssim$0.35 $R_{200,\,c}$. The RF8 and RF64 runs show slightly smaller values, possibly as a result of resolution effects kicking in. At even higher thresholds ($N_{C\,IV} \gtrsim 10^{14}$ cm^{-2}; dotted curves), values are further smaller, and no monotonic trend is seen with numerical resolution. This marks the regime where stochastic effects overshadow any clear trend of resolution.

As before, we further provide a zeroth-order comparison with other results from the literature. Using 43 low-mass $z \lesssim 0.1$ galaxies, Bordoloi et al. (2014) showed that the C IV covering fraction for absorption systems with rest-frame EW $W_r > 0.1$ Å (> 0.3 Å) is ~0.8 (0.4) at 0.2 $R_{200,\,c}$, and drops to ~0.4 (0.15) at 0.4 $R_{200,\,c}$. Schroetter et al. (2021) find that the covering fraction of C IV absorbers with $W_r > 0.7$ Å at $z \sim 1.2$ exceeds 50 per cent at impact parameters $\lesssim 23^{+62}_{-16}$ kpc. With a sample of eight $L_\star$ galaxies observed at $z \sim 2$, Rudie et al. (2019) show that, for impact parameters $\lesssim$100 pkpc, eight, seven, and six of those sightlines satisfy $N_{C\,IV}$ thresholds of 10^{12}, 10^{13}, and 10^{14} cm^{-2}, respectively. While direct comparison with our results is not straightforward, we note that these numbers are largely (qualitatively) consistent with the trends of our RF512 (black) curves.

On the theoretical end, the semi-analytic model of Faerman & Werk (2023) predicts that, at fixed distance, the median C IV column density can vary between ~$10^{12.5}$ cm^{-2} to ~$10^{13.5}$ cm^{-2} out to impact parameters of ~0.5 $R_{200,\,c}$, depending on the chosen value of η (see discussion earlier). As with H I, the variation is a result of greater photo-ionization producing larger amounts of C IV. At larger impact parameters, their model predicts that $N_{C\,IV}$ is roughly independent of η, and drops from ~$10^{12.5}$ cm^{-2} at ~0.5 $R_{200,\,c}$ to $\lesssim 10^{12.0}$ cm^{-2} at the virial radius. These numbers are largely consistent with the trend of our solid curves ($N_{C\,IV} > 10^{12}$ cm^{-2}).

With the same set of Auriga galaxies discussed earlier, Hani et al. (2019) studied the covering fractions of $N_{C\,IV} \gtrsim 10^{14.7}$ cm^{-2} absorbers in the CGM. They find that the distribution of covering fractions is broad, varying from ~0 to ~0.8, with the distribution peaking at ~0.6, when all sightlines between 0.2 $R_{200,\,c}$ and $R_{200,\,c}$ are stacked. For these column densities, we find such high covering fractions only in the innermost region of the CGM ($\lesssim$0.3 $R_{200,\,c}$).

Similarly, for a stacked set of sightlines passing through the CGM, the FOGGIE simulations predict that the covering fraction of C IV is ~40 per cent (~5 per cent) for $N_{C\,IV} > 10^{12}$ cm^{-2} (> 10^{14} cm^{-2}) absorbers, albeit at $z = 2$ (Peeples et al. 2019). They find that increased resolution leads to higher covering fractions for $10^{11.5} \lesssim N_{C\,IV} \lesssim 10^{13.5}$ cm^{-2}. However, our results are converged for lower thresholds ($N_{C\,IV} \sim 10^{12}$ cm^{-2}), and the impact of increasing resolution is only seen at higher values ($N_{C\,IV} \gtrsim 10^{13}$ cm^{-2}). As with H I, this is likely a direct consequence of differences in our methods as well as 'base' resolution (see discussion earlier).

For the same column density cuts, we compare the covering fractions of O VII in the right panel, a tracer of the hotter, volume-filling component of the CGM. At all considered cutoffs, the covering fractions are close to unity, except at the very outskirts of the halo ($\gtrsim$0.9 $R_{200,\,c}$). This is consistent with results from the SIMBA and EAGLE simulations, where covering fractions of O VII are close to ~100 per cent for thresholds $N_{O\,VII} \lesssim 10^{14.0}$ cm^{-2} (Wijers et al. 2019; Bradley et al. 2022; Tuominen et al. 2023). Observationally estimating covering fractions of O VII in extragalactic haloes is not yet possible, since this requires X-ray spectroscopy beyond the sensitivity and spectral resolution of currently available technology (e.g. Williams, Mulchaey & Kollmeier 2013). However, observations of O VII absorption from the Milky Way halo and the Local Group suggest that O VII is present ubiquitously out to large distances (Bregman & Lloyd-Davies 2007).

Throughout most of the halo, the convergence of O VII covering fraction with resolution is excellent, and consistent with the results from the TNG50 simulations run with slightly coarser resolutions

(Nelson et al. 2018b). However, it is important to note that the convergence of covering fractions does not imply the same for other properties of the same species of gas. Relatedly, the physical processes that drive the formation of these different species, and their kinematics, may further depend on resolution (e.g. Fielding et al. 2020; Lochhaas et al. 2023), the exploration of which we save for future work.

3.4 Small scale structure and cold clouds

We next explore the impact of numerical resolution on small scale structure in the CGM. We begin with an illustration in Fig. 8. This is similar in layout to Fig. 1, but for a zoomed-in region of the CGM of one halo, at four different resolutions: clockwise, the resolution increases from the top left to lower left. Colour indicates the local gas spatial resolution (cell radius), ranging from $\sim$30 pc in the brightest pixels, to $\sim$3 kpc in the darkest.

In the best resolved simulation (RF512; lower left), structures are clearly apparent at scales of $\sim O(1$ kpc), with the underlying gas distribution sampled by resolution elements that are typically an order-of-magnitude smaller in comparison, with $\sim O(100$ pc). Features become fuzzy as resolution decreases, with small scale structure from the RF512 run visibly absent at TNG50-2 resolution (top left panel). While the small scale structure visible in higher resolution simulations is directly due to the presence of a greater number of resolution elements, it is also possible that condensation of hot-halo gas into cold, and thus denser, gas is also more efficient at increased resolution (e.g. Joung, Bryan & Putman 2012).

To quantify this small scale structure we measure the abundance of clouds that are formed by cold, dense gas. We define clouds as contiguous sets of cold gas cells ($T < 10^{4.5}$ K), restricting to those cells that are not bound to any satellites (Ramesh et al. 2023b), and consider only those clouds that are composed of at least 10 gas cells. Owing to the nature of the Voronoi tessellation which represents our gas distribution, these clouds can trace gas structures of arbitrary sizes, masses, and shapes.

In Fig. 9, we begin by exploring the distribution of cloud masses. The main panel shows the number of clouds as a function of total mass of the cloud. Curves show the median of the sample, while the shaded regions show the 16th–84th percentile values. Results from RF8, RF64, and RF512 are shown in orange, red, and black, respectively. For comparison, we show results from Ramesh et al. (2023b) in the dashed-blue curve, that is, the results for the TNG50 MW-like sample, from which the eight GIBLE haloes were selected.

The RF8 (orange) curve peaks at roughly $10^{6.3}$ M$_\odot$, that is, close to ten times the resolution limit to the of the simulation, with $\sim$100 clouds on average in this bin. Clouds grow less frequent with increasing mass, with the average number dropping to $\sim$30 (4) clouds for clouds with masses 10^7 (10^8) M$_\odot$. The curve drops sharply towards the left of the peak, indicating the resolution limit. Smaller structures are simply absent beyond this point, at this resolution. The TNG50 dashed-blue curve is largely similar to RF8. Although we do not explore cloud properties in this work, it is clear that the results of Ramesh et al. (2023b) are impacted by this improved resolution, particularly with respect to 'small clouds' (i.e. $M_{\rm cl} \lesssim 10^6$ M$_\odot$) that were missing in the previous sample.

The RF64 (red) curve is qualitatively similar to RF8, but with the peak shifted by a factor $\sim$8 to the left, signifying a better mass resolution by a factor of 8. Similarly, the RF512 (black) curve is offset horizontally further by a factor of 8. The number of 'small' clouds thus increases continually with improving resolution, as signified by the arrow towards the top left of the panel. On the high-mass end, very good convergence is seen between the different resolution runs. This is important since it suggests that the smallest, marginally resolved clouds present at a given resolution level are not artefacts of limited resolution, but are rather true, real structures. It is likely that analogues of[7] these small clouds would also be present in higher resolution runs, albeit better resolved.

The eight panels in the top row show the number of clouds as a function of cloud mass for each of the eight haloes separately. In each panel, very good convergence is seen between the different resolution runs at the high-mass end, as in the main panel, although the first case (S91) is visibly noisy as a result of an abnormally small number of clouds.

Recent observations of high-velocity clouds (HVCs) in the Milky Way sky have revealed a large range in the masses of these objects, with values ranging from as high as $\sim 10^7$ M$_\odot$ (Thom & Chen 2008) to $\lesssim 10^5$ M$_\odot$ (Wakker 2001; Adams, Giovanelli & Haynes 2013). With the RF512 runs, we can thus begin probing a similar mass regime as clouds observed in the Milky Way, albeit with the caution of limited resolution at the low-mass end.

In Fig. 10, we explore the distribution of cloud sizes, that is, a measure for the spatial distribution of gas within clouds. Following Ramesh et al. (2023b), we fit the vertices of the Voronoi cells of the outer layer of each cloud to an ellipsoid, and define the size to be half the mean of the lengths of the three axes of the ellipsoid. Similar to Fig. 9, the main panel shows the median results from RF8, RF64, and RF512 in orange, red, and black, respectively. The 16th–84th percentiles of the sample are shown with shaded bands.

The RF8 curve peaks at a cloud size of roughly 4.5 kpc, with $\sim$100 clouds in this bin. The number of more massive clouds is smaller, with only $\sim$10 clouds in the high-size end ($\sim$30 kpc). As with the distribution of cloud masses, the curve drops rapidly to the left of the peak, signifying the resolution limit of the simulation. The RF64 curve peaks at $\sim$2.2 kpc, that is, roughly a factor of two smaller than RF8, corresponding to a factor of eight in mass (or resolution elements). Similarly, the RF512 curve is further offset by a factor of $\sim$2, peaking at $\sim$1 kpc. As before, with increasing resolution, the smallest objects are always more in number.

However, unlike the case of cloud masses, the number of high-size objects does not converge as rapidly, as can be seen by the vertical offsets between different curves at the right side of the main panel. To understand this trend, in the inset on the top-right, we show the relation between cloud masses (y-axis) and cloud size (x-axis). Note that the range of the x-axis of the inset, and the position of ticks, is exactly the same as that of the x-axis of the main panel. For visual clarity, we avoid showing the percentile regions, and only plot the medians in the inset.

At fixed cloud size, on average, cloud masses are lower at better resolution levels. For instance, at a size of 10 kpc, clouds have an average mass of $\sim 10^{6.5}$ M$_\odot$ in the RF8 simulations. At RF64-resolution, these clouds are slightly less massive at $\sim 10^{6.35}$ M$_\odot$. At even higher resolution, analogous clouds in the RF512 simulations have an average mass of $\sim 10^6$ M$_\odot$. We interpret this decrease in mass (at fixed cloud size) to be the result of their spatial distribution being captured to a better extent by the increased number of resolution elements (see also Hummels et al. 2019).

[7] An exact realization of a given cloud is unlikely to be present in another run due to the butterfly and zoom timing effects discussed earlier.

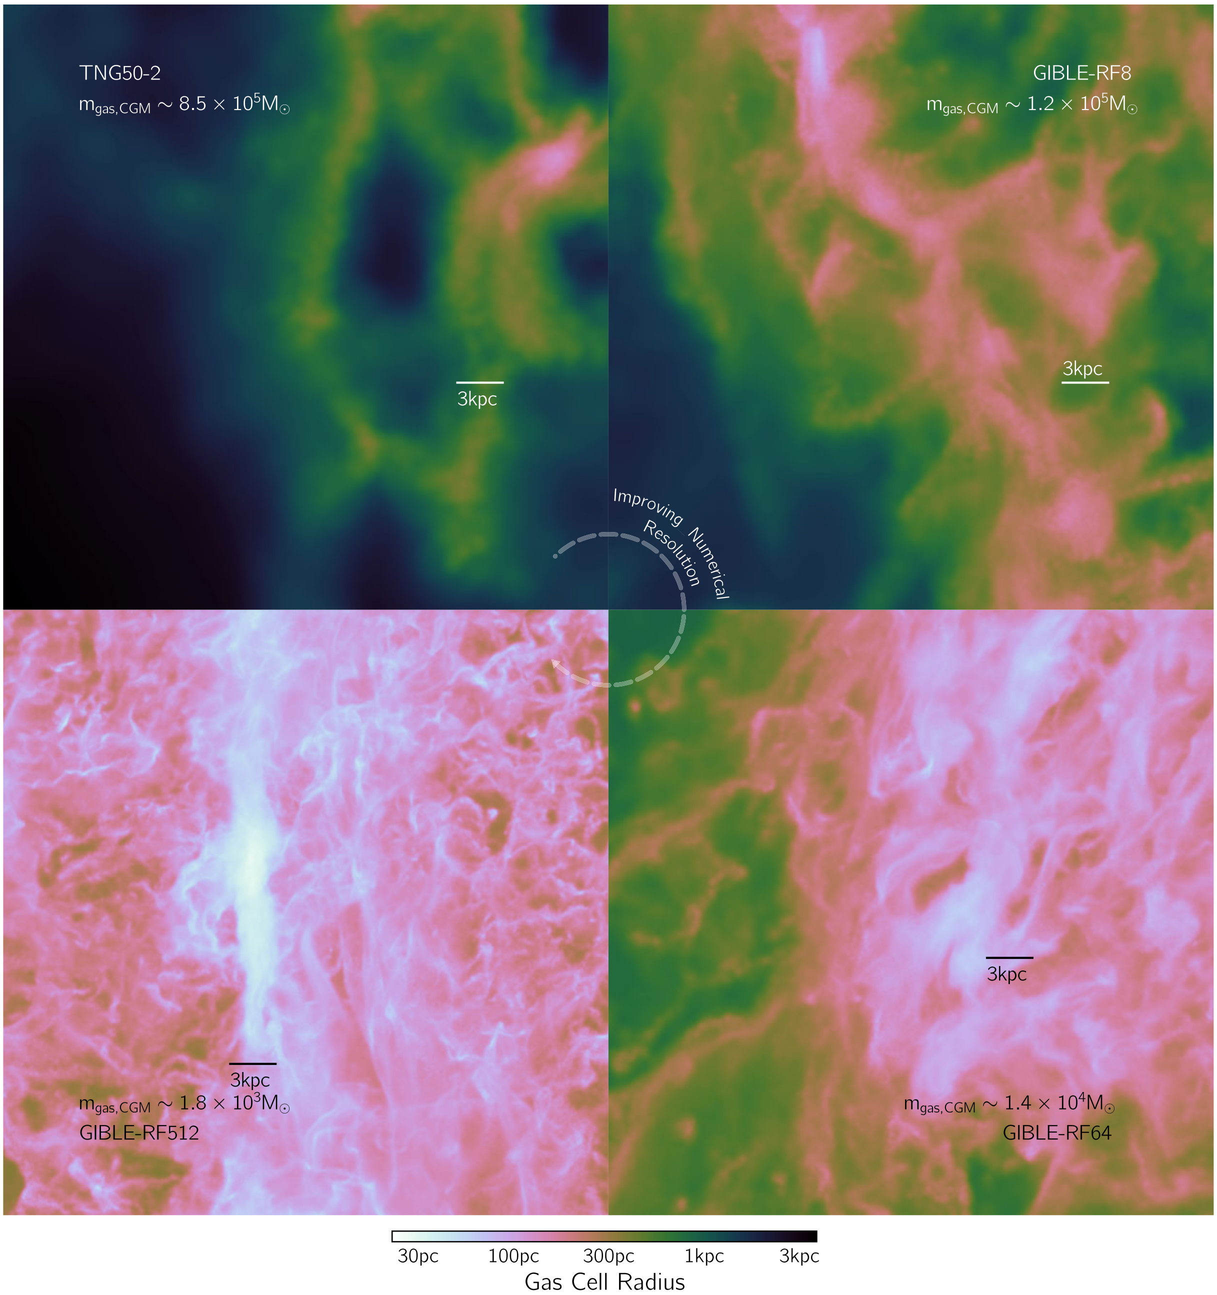

Figure 8. A qualitative illustration of the impact of resolution on small scale structure. The layout is similar to that of Fig. 1, wherein we show a zoomed-in region of one halo from four different simulations, with improving numerical resolution from top left to bottom left in clockwise direction. Colours correspond to the spatial resolution, that is, radii of gas cells. An increased amount of small scale structure is visibly apparent at higher resolutions.

This explains the lack of convergence between the medians of the three resolution levels at the high-size end: at fixed size, cloud masses are smaller at better resolution, and these clouds are typically more in number (Fig. 9). As a result, the higher resolution curves are shifted vertically upwards and horizontally towards the right, leading to a slower apparent convergence. However, for clouds above a certain mass threshold, we suspect sizes to be converged across resolutions, since such massive clouds are always resolved by a large number of resolution elements. This evolution towards convergence can be seen at the right-most edge of the main panel, where the RF64 median asymptotically approaches that of RF512. Unfortunately, the most massive clouds available with the GIBLE sample are not sufficient to demonstrate this for the RF8 runs. Comparison with a future, even higher resolution simulation is required to test this hypothesis.

Similar to Fig. 9, the panels in the top row show the number of clouds as a function of cloud size, for the eight GIBLE haloes separately. While most cases are in agreement with the conclusions of the main panel, many actually show decent convergence between the three resolution levels at the high-size end. This suggests that the lack of convergence clearly seen in the last two haloes, but not at all in the third, fourth, or fifth may be due to the particular state or evolution of the CGM of those haloes.

Observational studies suggest that HVCs typically have sizes of $\lesssim 10$ kpc (e.g. Thom & Chen 2008). Similar to the discussion with

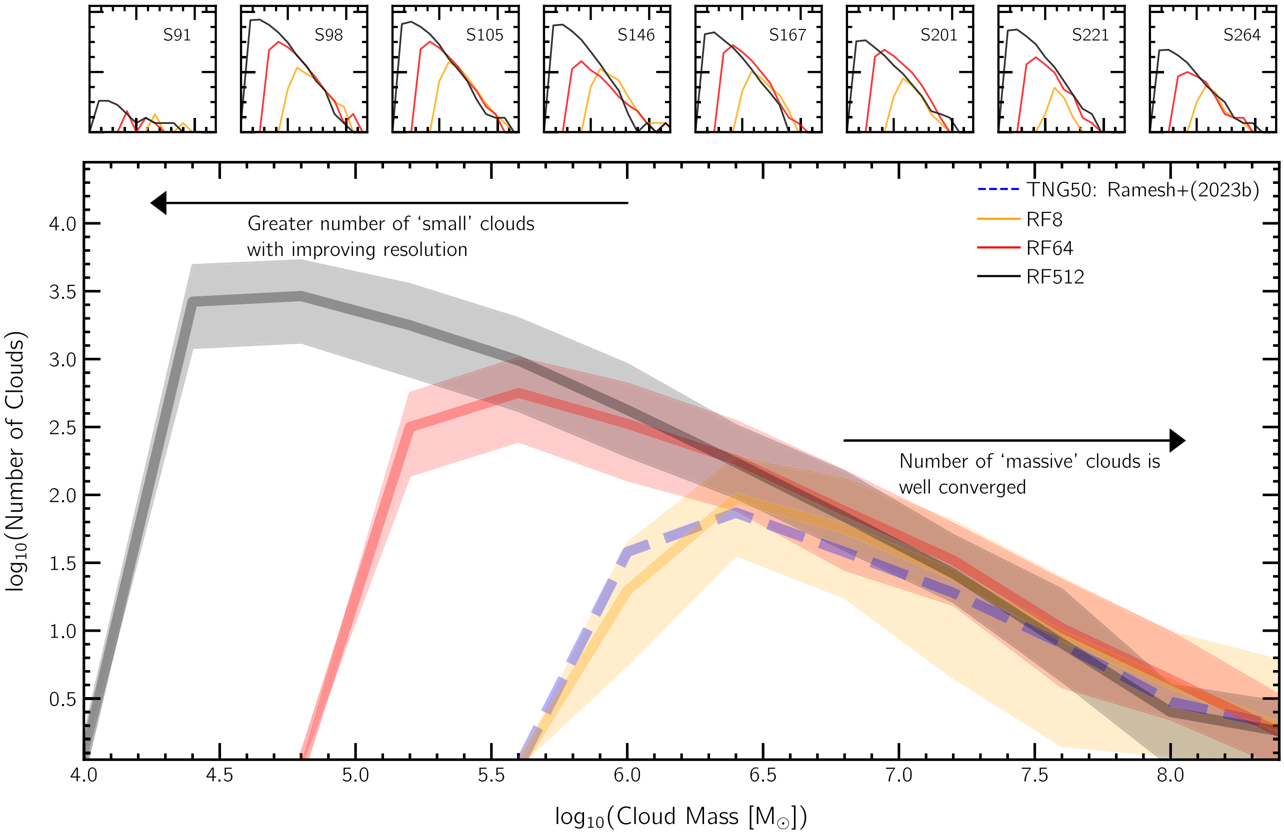

Figure 9. Number of clouds as a function of their mass. The main panel shows the median of the sample with solid lines, while the shaded regions show the 16th–84th percentile values. The solid orange, red, and black curves correspond to RF8, RF64, and RF512, respectively, while the dashed blue curve shows results from Ramesh et al. (2023b), that is, the entire TNG50 MW-like sample. Small panels in the top row show results for each of the eight GIBLE haloes separately. Both the axes labels and the axes ranges of these eight stamps are identical to those of the main panel. Excellent convergence with numerical resolution is seen at the high mass end, while the number of small clouds increases with improving resolution, since these are simply absent in lower resolution runs.

HVC masses above, we note the overlap in sizes of simulated clouds with their observed counterparts, thus making these runs, primarily the RF512s, ideal to begin studying such objects.

3.5 Cloud abundances along absorption sightlines

While the earliest observations of (high-velocity) clouds in the Milky Way sky were made possible through H I emission studies (Muller, Oort & Raimond 1963; Wakker 1991; Wakker & van Woerden 1997), more recent explorations have begun using absorption-line measurements to identify cloud like features in the Milky Way (e.g. Lehner et al. 2012; Richter et al. 2017) and extragalactic haloes (e.g. Zahedy et al. 2019; Weng et al. 2022). In the latter case, the number of clouds along a given line of sight can be estimated through the number of Voigt profile components required to fit the observed spectrum. In addition, gravitational lens absorption studies, that is, those that probe multiple sightlines of a single-foreground CGM using different (lensed) images of the background source, too provide the avenue to better study cloud-like features (e.g. Rubin et al. 2018; Kulkarni et al. 2019).

In Fig. 11, we study the impact that numerical resolution has on making predictions for observations of this kind. For simplicity, we do not make synthetic spectra to exactly mimic observational studies, but rather adopt a simpler approach possible with the simulations. Specifically, we project the mass of each cloud onto a 1000×1000[8] grid of pixels using the standard SPH-kernel, and for every pixel that the cloud contributes a H I column density $> 10^{16}\,\mathrm{cm^{-2}}$ to, the cloud count of that pixel is increased by one. Each contribution of $N_{\mathrm{H\,I}} > 10^{16}\,\mathrm{cm^{-2}}$ by a cloud is thus assumed to create a measure-able dip in an absorption spectrum. Further, it is assumed that every cloud produces a distinct dip in the spectrum, something which may not always be true, for example if two clouds along the line of sight have very similar velocities. Further, the impact of limited spectral resolution and/or presence of noise is ignored. What follows is thus meant as a zeroth-order theoretical study for the number of discrete clouds along a line of sight through the CGM.

The top panel visually demonstrates this technique. We choose the halo with the most clouds for this purpose (S98, RF512), with the central galaxy oriented edge-on at the centre. The image extends $\pm 1.05\ R_{200,\mathrm{c}}$ from edge to edge along the plane of the image, and $\pm 1\ R_{200,\mathrm{c}}$ in the perpendicular direction. We draw two circles at [0.15, 1.00] $R_{200,\mathrm{c}}$, signifying the adopted boundaries of the CGM. Only clouds between these two (3D) boundaries are considered for this analysis, even though clouds are visible within the inner circle as a result of projection effects.

[8]This corresponds to an average pixel size of $\sim$0.2 kpc. For a column density threshold of $10^{16}\,\mathrm{cm^{-2}}$, we find that this choice has a minimal impact on our results for reasonable pixel sizes ($\lesssim$1 kpc)

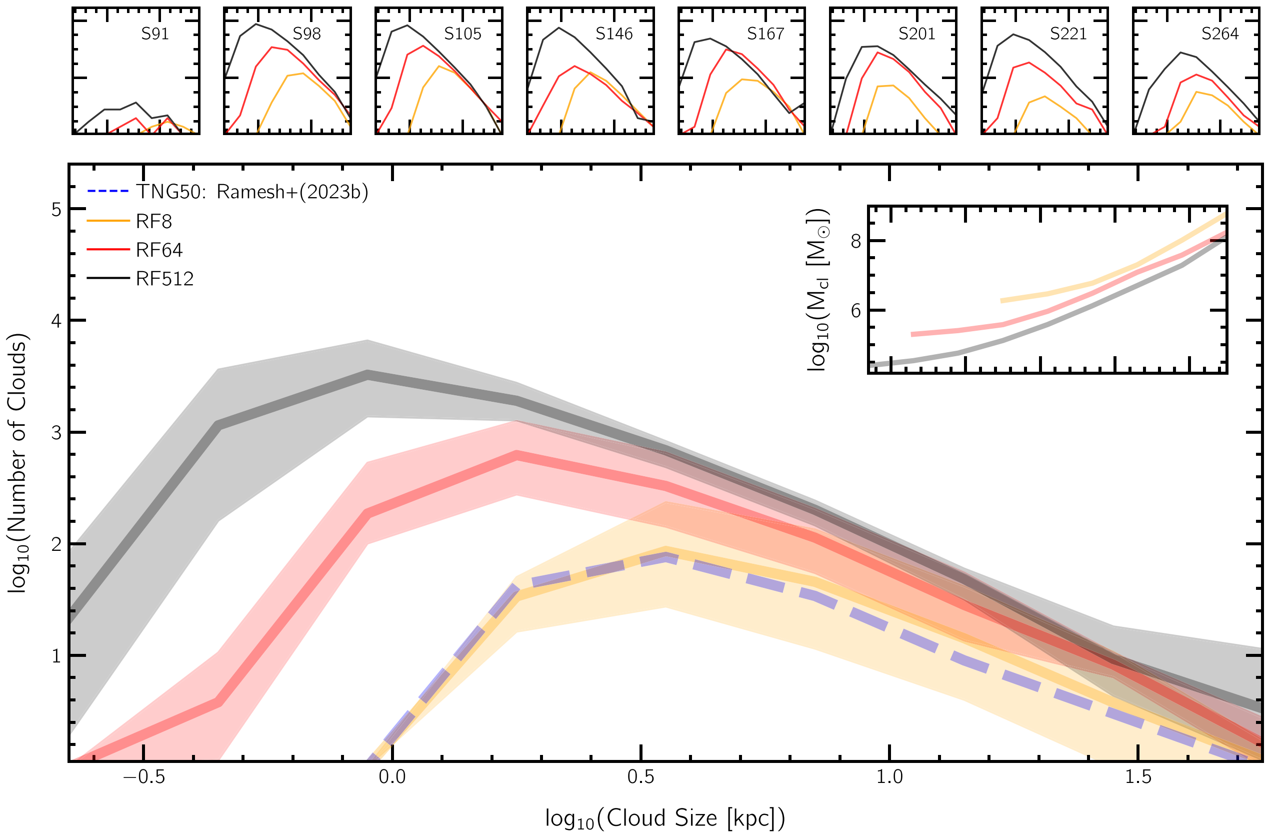

Figure 10. Number of clouds as function of their size. In the main panel, medians of the three resolution runs are shown with different colours, and 16th–84th percentile values with shaded regions. Similar to Fig. 9, smaller clouds are present at higher resolution, which are also greater in number. However, unlike Fig. 9, the number of high-size clouds is not typically converged. We explore a possible reason for this in the inset, which shows cloud mass as a function of cloud size. The panels in the top row show the same results independently for the eight GIBLE haloes. Both the axes labels and the axes ranges of these eight stamps are identical to those of the main panel.

Colours show the number of (H I i.e. cold) clouds along each sightline. Sightlines with no clouds are coloured black, while those with only one cloud are dark red, two and three clouds are orange, four and five clouds are blue, and six or more clouds are white, as shown by the discrete colourbar in the top-left. There is a visible radial trend with more clouds per sightline, on average, at smaller impact parameters, with typical colours shifting from blue/white close to 0.15 $R_{200,\,\mathrm{c}}$ (4 + clouds per sightline) to red close to the virial radius (1 cloud). This is broadly consistent with our earlier findings from the TNG50 simulation, where a larger number of clouds are seen at smaller galactocentric distances, each of which are smaller than their most distant counterparts (Ramesh et al. 2023b). However, even at fixed impact parameter, there is a large variation: for instance, in the innermost regions of the halo, there are clearly sightlines with 1–2 clouds adjacent to other sightlines with 4–6+ clouds.

In the lower panel, we quantify the radial trend discussed earlier. For a given impact parameter, we compute the mean number of clouds along those sightlines that contain at least one cloud. This metric is thus it not a measure of the covering fraction of clouds, a related quantity which has already been explored earlier (Fig. 6), but rather describes the expected number of cold clouds along those sightlines for which an absorption signal is present. For each galaxy, we construct radial profiles for 100 random orientations, and compute their median. The three solid curves show the median of these radial profiles for the sample, while the shaded regions show the 16th–84th percentiles variations. Results corresponding to RF8, RF64, and RF512 are shown in orange, red, and black, respectively.

In our best resolved simulations (RF512), at an impact parameter of ∼0.15 $R_{200,\,\mathrm{c}}$, sightlines that contain clouds have multiple (2–3), on average. The number steadily drops to ∼2 at ∼0.5 $R_{200,\,\mathrm{c}}$, and further to ∼1 at the virial radius. At RF8 and RF64 resolution, clouds are typically less common at impact parameters ≲0.8 $R_{200,\,\mathrm{c}}$, but numbers approach convergence with RF512 at larger distances. We suspect that this is a result of clouds, on average, being more massive at larger distances (Ramesh et al. 2023b), and hence are typically better resolved across different resolution levels (Fig. 9). Overall, our RF512 runs suggest that absorption spectra of cold species at low impact parameters will have complex spectral morphologies, sometimes composed of a half dozen or more discrete absorbing clouds. This is consistent with results from other observational (e.g. Tripp et al. 2008; Werk et al. 2016) and simulation studies (e.g. Peeples et al. 2019; Marra et al. 2022; Hummels et al. 2023), suggesting that the connection between observed sightlines and underlying cloud structures may not be straightforward to interpret. Gravitational lens absorption studies also infer significant variability in component structure between different sightlines, depending on the ionization state probed and the transverse separation between sightlines (e.g. Augustin et al. 2021).

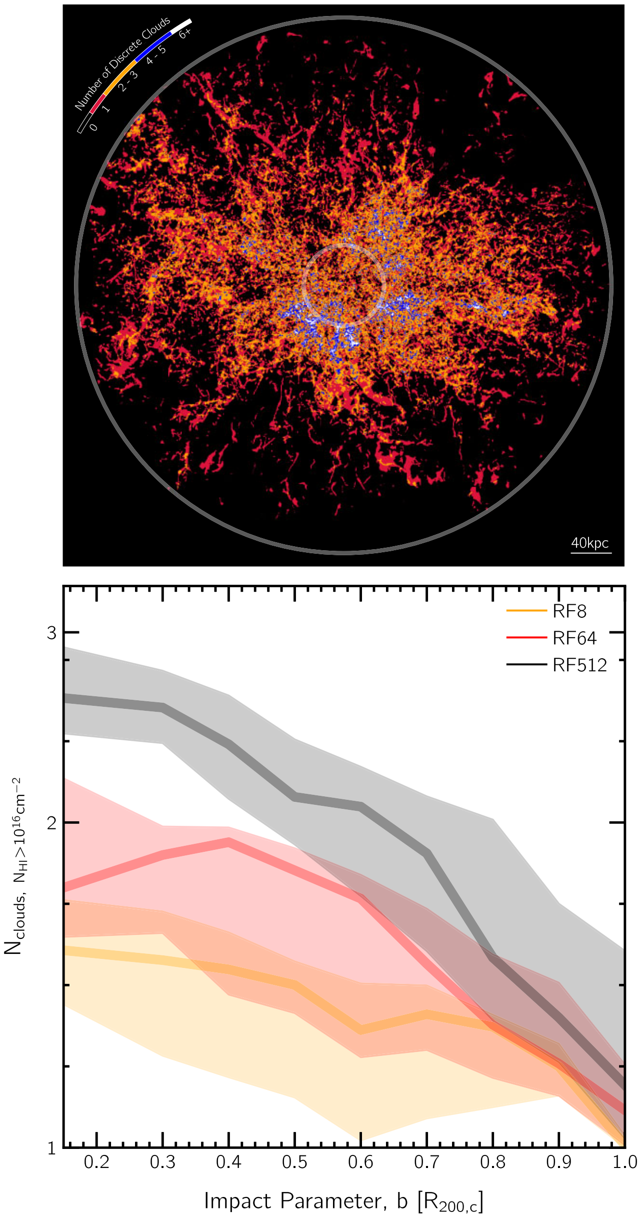

Figure 11. Number of distinct (cold i.e. H I) clouds along sightlines passing through the CGM. As detailed in the main text, a cloud is counted for a given sightline if it contributes $N_{HI} > 10^{16}\,cm^{-2}$. The top panel shows a visualization for the halo that contains the most clouds in our sample, with the galaxy oriented edge-on. The bottom panel shows the mean number of clouds per sightline as a function impact parameter, for those sightlines that contain at least one cloud (i.e. detectable absorption). In the innermost regions of the halo, higher resolution runs predict larger number of distinct clouds per sightline, while results are fairly well converged at larger impact parameters.

3.6 Small-scale density and velocity structure

In the final part of this paper, we explore two additional metrics to quantify the small-scale structure of CGM gas. In Fig. 12, we show the autocorrelation function of the density field (left) and the first-order velocity structure function (VSF; right). Both these metrics describe statistical quantities related to pairs of gas cells at positions $\vec{x}_1$ and $\vec{x}_2$, thus separated by distance $|\vec{x}_1 - \vec{x}_2|$: while the former shows how the densities at these points are related ($< \rho(\vec{x}_1) \times \rho(\vec{x}_2) >$), the latter describes the differences in velocities ($S_1(|\vec{x}_1 - \vec{x}_2|) =< |\vec{v}(\vec{x}_1) - \vec{v}(\vec{x}_2)| >$).[9]

In the main panels, we show results of the eight haloes from the RF512 runs, with each solid curve coloured according to the total CGM cold gas mass of that halo at $z = 0$. Additionally, for one halo (indicated with the thin overlaid black curve), we show a variety of test cases for the density autocorrelation, further described later. The insets in each panel shown the medians of the eight haloes, for the RF8, RF64, and RF512 runs separately, in orange, red, and black, respectively.

In the main panel on the left, curves peak close to 40–60 pc, corresponding to two times the best spatial resolution achieved in the CGM (Fig. 2). This peak, and the sharp drop to the left, is an artefact of finite resolution, and represents the resolution limit of the simulations. To the right of the peak, each curve decays with a finite slope. The amplitude and location of the peak, as well as the slope of the autocorrelation function at larger scales, all depend on the total CGM cold gas mass.

To unravel some of the drivers of these trends, we conduct a number of tests. The dotted-black line shows the autocorrelation function for a case where every gas cell is assigned a new density, given by the one-dimensional radial profile of the halo. This curve thus assesses the impact of the radial dependence of density on the autocorrelation function. Similar to the real curves, this test case drops sharply to the left of ∼40 pc, again signifying the resolution limit. However, at larger separations, the curve is more or less flat, at least out to ∼10–20 kpc. The finite slope seen in the real case is thus not an effect of the radial profile alone.

In the black dot-dashed and dashed curves, we show results when only T ∼ 10^4 and 10^5 K gas is included, respectively, with the latter vertically offset by 2 dex for better visibility. Similar to the dotted curve, these are flat out to ∼10−20 kpc. However, they have different amplitudes, corresponding to different mean densities at those temperatures, and they peak at different values, tracing the density dependence of spatial resolution. We conclude that the slopes of the actual curves, that is, the decrease of the density autocorrelation function with increasing spatial scale, arises due to an averaging across a 'fan' of contributing components, corresponding to different temperatures/densities, each of which have a different amplitude, and each of which dominates at different characteristic separations. This also explains the dependence of the slope on the cold gas content of the CGM: in haloes with less cold gas, the warm–hot phase dominates the autocorrelation at smaller distances to a larger extent, giving rise to a steeper slope.

The inset shows the impact of numerical resolution. With improving resolution, the small-scale peak shifts towards smaller separations, as expected. Additionally, the amplitude of the autocorrelation function increases, signifying the presence of more dense gas at improved resolutions. The slopes of these median curves to the right of the peak are roughly the same irrespective of resolution, indicating that the fraction of gas in different phases is typically not a strong function of resolution (Table 2).

For comparison, analogous metrics are used to quantify the level of clustering of sun-spots on the solar surface (e.g. Zhou, He &

[9]We use the positions of simulated gas cells themselves as tracers, and given that all (CGM) gas cells have roughly the same mass (Section 2), these statistics are therefore essentially mass weighted. An alternate approach would be to sample points randomly (or evenly) throughout the halo, providing a more volume-weighted result, although we do not explore this approach here.

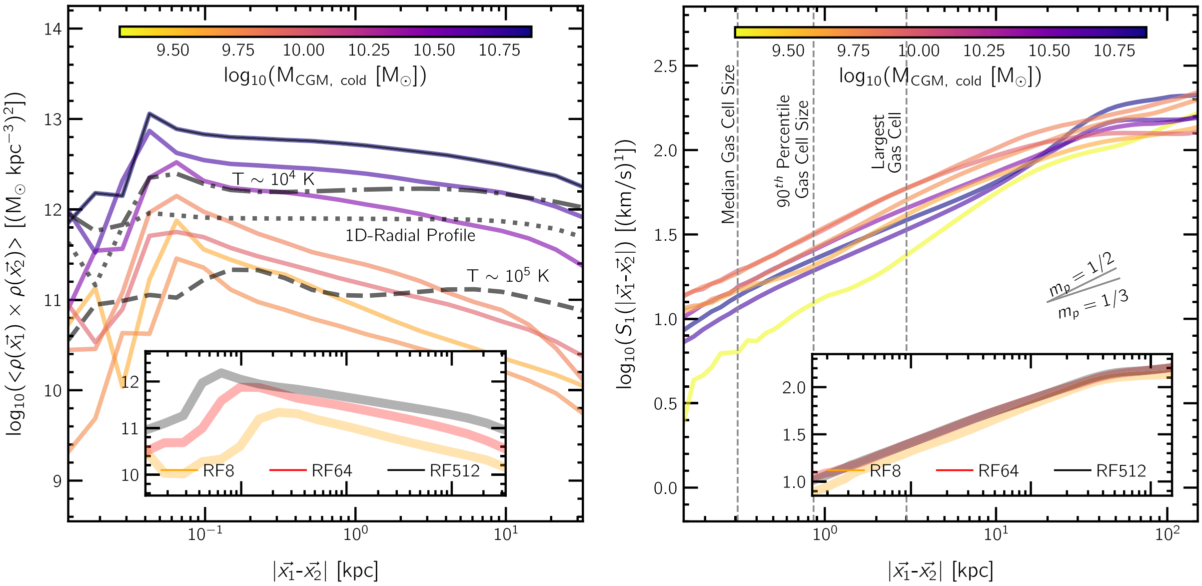

Figure 12. Quantifying small scale structure of CGM gas: the left panel shows the autocorrelation function of the density field, while the right panel shows the first-order VSF. The main panels show results for the eight haloes from the RF512 run, with curves coloured by the cold gas mass in the CGM of the corresponding halo. In the left panel, the dashed, dot-dashed, and dotted curves correspond to a variety of test cases included for comparison, as elaborated in the main text. Insets assess numerical convergence by comparing the median results of the RF8 (orange), RF64 (red), and RF512 (black) runs. Both the axes labels and the axes ranges of the insets are identical to those of the main panel.

Wan 2020), star clusters in galaxies (e.g. Grasha et al. 2017), and galaxies in the universe (e.g. Keihänen et al. 2019). In all cases, a given finite slope corresponds to greater clustering at smaller scales. Our measurement is also closely related to other metrics of cloud–cloud clustering, such as the Δ_{10} measurement used in Ramesh et al. (2023b). This phase-dependent (i.e. temperature binned) density autocorrelation could be a valuable tool to quantify the cold-phase versus warmer-phase of the CGM in simulations, and possibly also a way to, for example compare different simulations with different physics and across different numerical resolutions.

The right panel of Fig. 12 moves from density to velocity, showing the first-order VSF of CGM gas from $\sim$100 pc to $\sim$160 kpc scales. The three dashed vertical lines denote the median ($\sim$0.31 kpc), 90th percentile ($\sim$0.86 kpc), and maximum ($\sim$3 kpc) of CGM gas cell sizes in the RF512 runs. Since the dissipation of turbulent energy through numerical viscosity occurs on scales set by the gas cell sizes, dissipation in our simulations take place on various scales, corresponding to the varied gas cell radii (Fig. 2). The resulting impact on turbulence statistics is complex and depends on the numerical details of the hydrodynamical solver (see e.g. Bauer & Springel 2012; Zier & Springel 2022; Grete, O'Shea & Beckwith 2023).

All curves steadily increase with increasing separation, with a typical slope of $m_p \sim 0.45$ (VSF(r) $\propto r^{m_p}$), before plateauing at larger separations ($\gtrsim$50 kpc). This slope is steeper than that predicted by Kolmogorov turbulence ($m_p = 1/3$; Kolmogorov 1941), and slightly smaller than a value of $m_p = 1/2$ (Burgers 1948), both of which are shown by light grey curves in the main panel. Pairs of gas cells closer to each other are thus more likely to be moving with similar velocities. Similar to the left panel, a trend is present with respect to CGM cold gas mass, albeit not as pronounced. The inset assesses numerical convergence between the RF-runs: RF64 and RF512 yield similar results, with RF8 underpredicting values.

For comparison, various studies have explored the VSF of halo gas, albeit typically with idealized simulations where resolutions can be higher. For instance, Hillel & Soker (2020) analyse simulations of intracluster medium (ICM) gas being heated by jet-inflated hot bubbles. They find that, depending on the separation scale, slopes can either be steeper or shallower than $m_p = 1/3$, that is, Kolmogorov turbulence. Wang et al. (2021) study the impact of AGN jets in a Perseus-like cluster, using both hydrodynamical (HD), and magnetohydrodynamical (MHD) simulations. For the hot phase of gas, they find slightly shallower slopes in the MHD case ($m_p \lesssim 0.5$) than the HD case ($m_p \sim 0.5$), with a similar trend for the cold phase as well. Using an idealized simulation of a Milky Way-like halo, Lochhaas et al. (2020) report a slope of $m_p = 0.3$–0.4 for hot ($T > 10^{4.7}$ K) CGM gas, once again in the ballpark of expectations from Kolmogorov turbulence. These simulations are non-cosmological and do not include mergers, which may drive turbulence in gaseous haloes (Schmidt, Schmidt & Grete 2021). At higher mass scales, Vazza et al. (2017) find a slope slightly steeper than Kolmogorov turbulence in a cosmological zoom-in simulation of a cluster ($M_{halo} \sim 10^{14}$ M$_\odot$ at $z = 0$).

Recent observations have studied VSFs in a number of nearby clusters. For instance, Li et al. (2020) probe cool gas at the centres of clusters through their Hα emission, and find a slope of $m_p \sim 0.5$ for Persesus, and steeper slopes of $m_p \sim 0.6$ for Abell 2597 and for the inner 2.5 kpc of the Virgo cluster. While these values vary significantly, this may not reflect changing properties of turbulence, but could instead be due to measurement and projection effects in observations (see Li et al. 2020). Although most of the studies mentioned here focus on rather different environments than that of haloes like the Milky Way, we present them for context. Moreover, note that the first two studies mentioned earlier, as well as the results from GIBLE, use a three-dimensional VSF, in contrast to observations that are limited to the VSF in 2D projection.

The VSF is clearly a useful tool to gain some understanding of, and quantify, the kinematic structure of CGM gas. Similar to the density autocorrelation, this could be a powerful method to compare simulations run with different models and/or resolutions.

Moreover, one could compute a two-dimensional projected VSF to mock observations studies, thereby improving the theoretical understanding and the interpretation of such results (e.g. Mohapatra et al. 2022). Future X-ray spectroscopic imaging concept missions such as the Line Emission Mapper (LEM; Kraft et al. 2022) will enable us to measure the VSF of the hot CGM for Milky Way like galaxies in the future. Comparisons with such observations would likely prove useful to test and validate different feedback models in simulations, which could have an impact of the VSF of CGM gas.

4 SUMMARY AND CONCLUSIONS

In this paper we introduce the GIBLE simulations, a new suite of cosmological hydrodynamical zoom-in simulations of galaxies with a 'CGM refinement' method, designed to increase numerical resolution in the CGM. While gas in the galaxy is maintained at a lower resolution, CGM gas is super-refined to significantly better mass (and spatial) resolution. Here we introduce a sample of eight, Milky Way-like galaxies, drawn originally from the TNG50 simulation, which reach a mass resolution of $\sim 10^3\,\mathrm{M}_\odot$ in the CGM, competitive with current state-of-the-art zoom-in simulations of Milky Way-like galaxies at $z = 0$. The high computational efficiency of our CGM refinement method enables us to simulate a relatively large and diverse set of haloes in order to explore their diversity at high resolution. Our main findings of this presentation paper, the first in the GIBLE series are

(i) increased CGM gas resolution not only results in the halo being more finely resolved in general, but also in visibly apparent structure on smaller scales (Fig. 1). In our best resolved runs, we achieve a median spatial resolution of $\sim$75 pc at the inner boundary of the CGM ($0.15\,R_{200,\,\mathrm{c}}$), which coarsens with increasing distance to $\sim$700 pc at the virial radius. The cold phase of gas, however, is better spatially resolved at all distances (Fig. 2).

(ii) Integrated (i.e. global) properties of the halo that we explore are well converged at the resolutions we achieve (Fig. 4 and Table 2). These include the CGM gas mass, cold gas fraction, temperature, and volume fractions of cold, warm, and hot phases. Time-averaged, sample-averaged PDFs of gas properties – temperature, metallicity, magnetic field strength, and velocity – are also well converged (Fig. 5). The notable exception is the abundance of cold CGM gas clouds, the number of which continues to increase with increasing resolution.

(iii) While some key observables are well converged with resolution, others are not. For instance, the H I covering fraction of pLLS systems ($N_{\mathrm{H\,I}} > 10^{16}\,\mathrm{cm}^{-2}$) is only converged at resolutions $\lesssim 10^4\,\mathrm{M}_\odot$, and simulations run at a resolution of $\sim 10^5\,\mathrm{M}_\odot$ under predict covering fractions out to a distance of $\sim 0.8\,R_{200,\,\mathrm{c}}$. Covering fractions of more dense gas, that is, sub-DLA and DLA systems, are not converged at our resolutions (Fig. 6).

(iv) Similar trends are also observed for the covering fraction of Mg II and C IV, which are tracers of the cold- and warm-phase of gas, respectively. Alternatively, the diffuse, volume-filling hot phase traced by O VII is well converged at these resolutions (Fig. 7).

(v) As expected, better resolved simulations produce structures on smaller scales (Fig. 8). One way to quantify this is by characterizing the number and properties of cold, dense gas clouds. At better resolutions, the CGM has a larger number of such clouds, with the number increasing steeply as a power law. Importantly, the number of massive clouds, that is, clouds well above the resolution limit, are converged across resolution levels (Fig. 9).

(vi) This has an impact on predictions for example the number of pLLS absorbers along a given line of sight through the CGM. Better resolved simulations predict a larger number of absorbers at small impact parameters, although results approach convergence at larger impact parameters, where clouds are bigger. While the median number of absorbers in our highest resolution runs varies between 1 and 4, depending on impact parameter, absorption spectra along some sightlines may have more complex spectral morphologies, with possibly 6 or more distinct absorbers (Fig. 11).

(vii) We also quantify small-scale structure through the autocorrelation of the density field and VSF. Both metrics yield valuable information about features in CGM gas. The former shows a finite slope with the autocorrelation decreasing towards larger separations, signifying that smaller-separation pairs are typically dominated by colder, denser gas. The latter reveals a typical slope of $m_p \sim 0.45$, a value between that predicted by Kolmogorov and Burgers turbulence. (Fig. 12).

The GIBLE cosmological simulations aim to bridge the gap between large-volume cosmological and non-cosmological idealized setups. However, despite achieving a resolution similar to the (currently) best resolved cosmological zoom-in simulation of a Milky Way-like galaxy, we do not resolve pc-scale features in the CGM. Future extensions of GIBLE will aim to push the resolution significantly higher. Further, the current simulations omit potentially important physical processes such as cosmic rays or thermal conduction. Future simulations will aim to tackle these shortcomings.

ACKNOWLEDGEMENTS

RR and DN acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Group (grant number NE 2441/1-1). RR is a Fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). The authors thank Volker Springel for providing access to the private branch of the AREPO code, Annalisa Pillepich for a helpful discussion on selecting a diverse sample of Milky Way-like galaxies, Christoph Federrath for suggestions on the VSF analysis, and Cameron Hummels, Freeke van de Voort, and Robert J J Grand for sharing relevant data points. This work has made use of the VERA supercomputer of the Max Planck Institute for Astronomy (MPIA), the COBRA supercomputer, both operated by the Max Planck Computational Data Facility (MPCDF), and the HELIX supercomputer of the University of Heidelberg. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1597-1 FUGG. This research has made use of NASA's Astrophysics Data System Bibliographic Services.

5 DATA AVAILABILITY

Collaboration on new projects with the GIBLE simulation data is welcomed, and potential collaborators are encouraged to contact the authors directly. Data related to this publication are available upon reasonable request to the corresponding author.

REFERENCES

Adams E. A. K., Giovanelli R., Haynes M. P., 2013, ApJ, 768, L77
Anand A., Nelson D., Kauffmann G., 2021, MNRAS, 504, 65
Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, MNRAS, 426, 2046
Augustin R., Péroux C., Hamanowicz A., Kulkarni V., Rahmani H., Zanella A., 2021, MNRAS, 505, 6195

Balogh M. L., Baldry I. K., Nichol R., Miller C., Bower R., Glazebrook K., 2004, ApJ, 615, L101
Bauer A., Springel V., 2012, MNRAS, 423, 2558
Bernardi M., Meert A., Vikram V., Huertas-Company M., Mei S., Shankar F., Sheth R. K., 2014, MNRAS, 443, 874
Bordoloi R. et al., 2014, ApJ, 796, L136
Bradley L., Davé R., Cui W., Smith B., Sorini D., 2022, preprint (arXiv:2203.15055)
Bregman J. N., Lloyd-Davies E. J., 2007, ApJ, 669, L990
Bryan G. L. et al., 2014, ApJS, 211, 19
Burgers J., 1948, Adv. Appl. Mech., 1, 171
Cameron A. J. et al., 2021, ApJ, 918, L16
Cui W., Davé R., Peacock J. A., Anglés-Alcázar D., Yang X., 2021, Nature Astron., 5, 1069
Danforth C. W. et al., 2016, ApJ, 817, L111
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, L371
Donahue M., Voit G. M., 2022, Phys. Rep., 973, 1
Donnari M. et al., 2021, MNRAS, 500, 4004
Engler C. et al., 2021, MNRAS, 507, 4211
Faerman Y., Werk J. K., 2023, preprint (arXiv:2302.00692)
Feldmann R. et al., 2023, MNRAS
Ferland G. J. et al., 2013, RMxAA, 49, 137
Fielding D. B., Bryan G. L., 2022, ApJ, 924, L82
Fielding D., Quataert E., McCourt M., Thompson T. A., 2017, MNRAS, 466, 3810
Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, ApJ, 894, L24
Geha M., Blanton M. R., Yan R., Tinker J. L., 2012, ApJ, 757, L85
Genel S. et al., 2018, MNRAS, 474, 3976
Genel S. et al., 2019, ApJ, 871, L21
Gnedin N. Y., Kravtsov A. V., 2011, ApJ, 728, L88
Goldsmith K. J. A., Pittard J. M., 2016, MNRAS, 461, 578
Grand R. J. J. et al., 2021, MNRAS, 507, 4953
Grasha K. et al., 2017, ApJ, 840, L113
Grete P., O'Shea B. W., Beckwith K., 2023, ApJ, 942, L34
Gronke M., Oh S. P., 2020, MNRAS, 492, 1970
Hani M. H., Ellison S. L., Sparre M., Grand R. J. J., Pakmor R., Gomez F. A., Springel V., 2019, MNRAS, 488, 135
Heesen V. et al., 2023, A&A, 670, 23
Henden N. A., Puchwein E., Sijacki D., 2020, MNRAS, 498, 2114
Hillel S., Soker N., 2020, ApJ, 896, L104
Ho S. H., Martin C. L., Schaye J., 2020, ApJ, 904, L76
Hsu W. H., Putman M. E., Heitsch F., Stanimirović S., Peek J. E. G., Clark S. E., 2011, AJ, 141, 57
Hummels C. B. et al., 2019, ApJ, 882, L156
Hummels C. B., Rubin K. H. R., Schneider E. E., Fielding D. B., 2023, preprint (arXiv:2311.05691)
Joung M. R., Bryan G. L., Putman M. E., 2012, ApJ, 745, L148
Keihänen E. et al., 2019, A&A, 631, 73
Kolmogorov A. N., 1941, Akademiia Nauk SSSR Doklady, 32, 16
Kopenhafer C., O'Shea B. W., Voit G. M., 2023, ApJ, 951, L107
Kraft R. et al., 2022, preprint (arXiv:2211.09827)
Kulkarni V. P., Cashman F. H., Lopez S., Ellison S. L., Som D., Maureira M. J., 2019, ApJ, 886, L83
Lehner N., Howk J. C., Thom C., Fox A. J., Tumlinson J., Tripp T. M., Meiring J. D., 2012, MNRAS, 424, 2896
Li Y. et al., 2020, ApJ, 889, L1
Lochhaas C., Bryan G. L., Li Y., Li M., Fielding D., 2020, MNRAS, 493, 1461
Lochhaas C. et al., 2023, ApJ, 948, L43
Marinacci F. et al., 2018, MNRAS, 480, 5113
Marra R., Churchill C. W., Kacprzak G. G., Nielsen N. M., Trujillo-Gomez S., Lewis E., 2024, MNRAS, 527, 10522
Medling A. M. et al., 2018, MNRAS, 475, 5194
Mohapatra R., Jetti M., Sharma P., Federrath C., 2022, MNRAS, 510, 2327
Muller C. A., Oort J. H., Raimond E., 1963, Academie des Sciences Paris Comptes Rendus, 257, 1661
Naiman J. P. et al., 2018, MNRAS, 477, 1206
Nelson D. et al., 2018a, MNRAS, 475, 624
Nelson D. et al., 2018b, MNRAS, 477, 450
Nelson D. et al., 2019, MNRAS, 490, 3234
Nelson D. et al., 2020, MNRAS, 498, 2391
Pakmor R., Marinacci F., Springel V., 2014, ApJ, 783, L20
Peeples M. S. et al., 2019, ApJ, 873, L129
Péroux C., Nelson D., van de Voort F., Pillepich A., Marinacci F., Vogelsberger M., Hernquist L., 2020, MNRAS, 499, 2462
Pillepich A. et al., 2018a, MNRAS, 473, 4077
Pillepich A. et al., 2018b, MNRAS, 475, 648
Pillepich A. et al., 2019, MNRAS, 490, 3196
Pillepich A., Nelson D., Truong N., Weinberger R., Martin-Navarro I., Springel V., Faber S. M., Hernquist L., 2021, MNRAS, 508, 4667
Pillepich A. et al., 2023, preprint (arXiv:2303.16217)
Popping G. et al., 2019, ApJ, 882, L137
Predehl P. et al., 2020, Nature, 588, 227
Prochaska J. X. et al., 2017, ApJ, 837, L169
Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, MNRAS, 430, 2427
Ramesh R., Nelson D., Pillepich A., 2023a, MNRAS, 518, 5754
Ramesh R., Nelson D., Pillepich A., 2023b, MNRAS, 522, 1535
Ramesh R., Nelson D., Heesen V., Brüggen M., 2023c, MNRAS, 526, 5483
Richter P. et al., 2017, A&A, 607, 48
Rodriguez-Gomez V. et al., 2016, MNRAS, 458, 2371
Rubin K. H. R., Hennawi J. F., Prochaska J. X., Simcoe R. A., Myers A., Lau M. W., 2015, ApJ, 808, L38
Rubin K. H. R. et al., 2018, ApJ, 859, L146
Rudie G. C. et al., 2012, ApJ, 750, L67
Rudie G. C., Steidel C. C., Pettini M., Trainor R. F., Strom A. L., Hummels C. B., Reddy N. A., Shapley A. E., 2019, ApJ, 885, L61
Scannapieco E., Brüggen M., 2015, ApJ, 805, L158
Schmidt W., Schmidt J. P., Grete P., 2021, A&A, 654, 115
Schneider E. E., Robertson B. E., 2018, ApJ, 860, L135
Schroetter I. et al., 2021, MNRAS, 506, 1355
Shen S., Mo H. J., White S. D. M., Blanton M. R., Kauffmann G., Voges W., Brinkmann J., Csabai I., 2003, MNRAS, 343, 978
Smith M. C. et al., 2024, MNRAS, 527, 1216
Springel V., 2010, MNRAS, 401, 791
Springel V., Hernquist L., 2003, MNRAS, 339, 289
Springel V. et al., 2008, MNRAS, 391, 1685
Springel V. et al., 2018, MNRAS, 475, 676
Springel V., Pakmor R., Zier O., Reinecke M., 2021, MNRAS, 506, 2871
Strateva I. et al., 2001, AJ, 122, 1861
Su M., Slatyer T. R., Finkbeiner D. P., 2010, ApJ, 724, L1044
Suresh J., Nelson D., Genel S., Rubin K. H. R., Hernquist L., 2019, MNRAS, 483, 4040
Tan B., Fielding D. B., 2024, MNRAS, 527, 9683
Thom C., Chen H.-W., 2008, ApJ, 683, L22
Tripp T. M., Sembach K. R., Bowen D. V., Savage B. D., Jenkins E. B., Lehner N., Richter P., 2008, ApJS, 177, 39
Truong N., Pillepich A., Nelson D., Werner N., Hernquist L., 2021, MNRAS, 508, 1563
Tumlinson J., Peeples M. S., Werk J. K., 2017, ARA&A, 55, 389
Tuominen T., Nevalainen J., Heinämäki P., Tempel E., Wijers N., Bonamente M., Aragon-Calvo M. A., Finoguenov A., 2023, A&A, 671, 103
van de Voort F., Springel V., Mandelker N., van den Bosch F. C., Pakmor R., 2019, MNRAS, 482, 85
van de Voort F., Bieri R., Pakmor R., Gómez F. A., Grand R. J. J., Marinacci F., 2021, MNRAS, 501, 4888
Vazza F., Jones T. W., Brüggen M., Brunetti G., Gheller C., Porter D., Ryu D., 2017, MNRAS, 464, 210
Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, Nature Rev. Phys., 2, 42
Wakker B. P., 1991, A&A, 250, 499
Wakker B. P., 2001, ApJS, 136, 463
Wakker B. P., van Woerden H., 1997, ARA&A, 35, 217
Wang C., Ruszkowski M., Pfrommer C., Oh S. P., Yang H. Y. K., 2021, MNRAS, 504, 898
Weinberger R. et al., 2017, MNRAS, 465, 3291

Weng S. et al., 2022, MNRAS, 512, 3638
Weng S. et al., 2023, MNRAS, 519, 931
Weng S., Péroux C., Ramesh R., Nelson D., Sadler E. M., Zwaan M., Bollo V., Casavecchia B., 2024, MNRAS, 527, 3494
Werk J. K. et al., 2016, ApJ, 833, L54
Wijers N. A., Schaye J., Oppenheimer B. D., Crain R. A., Nicastro F., 2019, MNRAS, 488, 2947
Williams R. J., Mulchaey J. S., Kollmeier J. A., 2013, ApJ, 762, L10
Wright R. J., Lagos C. d. P., Davies L. J. M., Power C., Trayford J. W., Wong O. I., 2019, MNRAS, 487, 3740
Zahedy F. S., Chen H.-W., Johnson S. D., Pierce R. M., Rauch M., Huang Y.-H., Weiner B. J., Gauthier J.-R., 2019, MNRAS, 484, 2257
Zhang H., Zaritsky D., 2022, ApJ, 941, L18
Zhou G., He H. Q., Wan W., 2020, ApJ, 899, L32
Zier O., Springel V., 2022, MNRAS, 515, 525

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 6

Zooming in on the Circumgalactic Medium with GIBLE: the Topology and Draping of Magnetic Fields around Cold Clouds

6.1 Statement of contribution

- Scientific Analysis: My contribution was central. In addition, the analysis benefited greatly from feedback provided by collaborators.
- Figures: I independently produced the figures based on the planned scientific analysis. The process also greatly benefited from iterative input by collaborators.
- Writing: I primarily wrote most of the manuscript text. The drafts were further refined based on feedback from collaborators.
- Code and Simulation Development: This paper made use of the new set of GIBLE simulations, which I set up and ran during the first half of my PhD. This also involved developing a 'CGM refinement' module into the AREPO code, which required extensive coding in C. The cumulative effort of Project GIBLE greatly benefited from input from my thesis supervisor, including, but not limited to, feedback on optimizing the code architecture and generating initial conditions for these new simulations. The code to analyze the simulation data and generate plots was written by me, but it also benefited from input provided by colleagues.

A&A, 684, L16 (2024)
https://doi.org/10.1051/0004-6361/202348786

Astronomy & Astrophysics

LETTER TO THE EDITOR

Zooming in on the circumgalactic medium with GIBLE

The topology and draping of magnetic fields around cold clouds

Rahul Ramesh[1], Dylan Nelson[1], Drummond Fielding[2,3], and Marcus Brüggen[4]

[1] Universität Heidelberg, Zentrum für Astronomie, ITA, Albert-Ueberle-Str. 2, 69120 Heidelberg, Germany
e-mail: rahul.ramesh@stud.uni-heidelberg.de
[2] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA
[3] Department of Astronomy, Cornell University, Ithaca, NY 14853, USA
[4] Hamburg University, Hamburger Sternwarte, Gojenbergsweg 112, 21029 Hamburg, Germany

Received 29 November 2023 / Accepted 27 March 2024

ABSTRACT

We used a cosmological zoom-in simulation of a Milky Way-like galaxy to study and quantify the topology of magnetic field lines around cold gas clouds in the circumgalactic medium (CGM). This simulation is a new addition to Project GIBLE, a suite of cosmological magnetohydrodynamic simulations of galaxy formation with preferential super-Lagrangian refinement in the CGM, reaching an unprecedented CGM gas mass resolution of ∼225 $M_\odot$. To maximize statistics and resolution, we focused on a sample of ∼200 clouds with masses of ∼10^6 $M_\odot$. The topology of magnetic field lines around clouds is diverse, from threading to draping, and there is large variation in the magnetic curvature (κ) within cloud-background interfaces. We typically find little variation of κ between upstream and downstream cloud faces, implying that strongly draped configurations are rare. In addition, κ correlates strongly with multiple properties of the interface and the ambient background, including cloud overdensity and relative velocity, suggesting that cloud properties impact the topology of interface magnetic fields.

Key words. galaxies: halos – galaxies: magnetic fields

1. Introduction

Observations and simulations suggest that galaxies are surrounded by a multiphase multiscale reservoir of gas. Termed the circumgalactic medium (CGM), this gaseous halo is believed to play a vital role in the growth and evolution of galaxies (see Donahue & Voit 2022 for a recent review of the CGM). While the volume of the CGM is dominated by a warm–hot component, it can also host small cold gas structures. The high-velocity clouds (HVCs) of the Milky Way are prototypical examples (e.g., Muller et al. 1963; Wakker & van Woerden 1997).

Despite having been first observed several decades ago, there remain a number of open questions regarding HVCs, and cold CGM clouds in general. Their expected lifetimes, and the nature of their growth and evolution, are uncertain. A number of idealized "cloud-crushing" simulations have explored these puzzles. While early studies suggested that cloud lifespans should be short (e.g., Klein et al. 1994; Mellema et al. 2002), certain physical mechanisms could enhance their survival. For instance, the Kelvin–Helmholtz instability may produce a warm interface layer between the cold cloud and the hot background, facilitating rapid cooling and cloud growth (e.g., Scannapieco & Brüggen 2015; Gronke & Oh 2018; Fielding et al. 2020).

In addition, nonthermal components including magnetic fields may be important. They can suppress fluid instabilities (e.g., Berlok & Pfrommer 2019; Sparre et al. 2020; Das & Gronke 2024), provide nonthermal pressure support (e.g., Girichidis 2021; Hidalgo-Pineda et al. 2024; Fielding et al. 2023), or enhance the Rayleigh–Taylor instability, thereby accelerating condensation (Grønnow et al. 2022). The direction and topology of magnetic field lines may also be important by influencing the amplification of magnetic energy density (Shin et al. 2008), the kinematics (Kwak et al. 2009), and the shape (Banda-Barragán et al. 2016; Brüggen et al. 2023) of clouds.

While these theoretical studies have advanced our understanding of cloud growth and survival, they have a fundamental limitation: they are all idealized noncosmological simulations. As a result, they must assume the existence of a preexisting cloud, and background, with particular properties. In the case of magnetic fields, the strength and orientation must be chosen ad hoc (i.e., freely explored). Cosmological simulations overcome this limitation by self-consistently evolving halo gas and magnetic fields over cosmic epochs, with the trade-off of coarser resolution. Recent cosmological simulations including TNG50 have been shown to realize small-scale cold gas structures (Nelson et al. 2020; Ramesh et al. 2023a), even at the limited resolution available in large uniform volumes.

Here we take a step forward by using Project GIBLE (Ramesh & Nelson 2024), a suite of cosmological zoom-in galaxy formation simulations with targeted, additional super-Lagrangian refinement of gas in the CGM. In particular, we present a new simulation of a Milky Way-like galaxy run to $z = 0$ with even higher resolution than our first GIBLE results. These simulations make it possible to better resolve and study small-scale phenomena in the full ΛCDM cosmological context, thereby bridging the gap between highly resolved idealized simulations, and more realistic cosmological runs at lower resolution. Building on our earlier work on the magnetothermal properties of the clumpy CGM in a cosmological context, we can now quantify, for the first time, the topology and draping of magnetic field lines around cold dense gas clouds in a self-consistent environment without the sensitivity to initial magnetic field geometries that limit the robustness of previous studies on this topic due to their idealized nature.

The paper is structured as follows. In Sect. 2 we describe Project GIBLE and our methodology. Our results are presented in Sect. 3, discussed in Sect. 4, and summarized in Sect. 5.

2. Methods

2.1. Simulation overview

For this paper we used Project GIBLE (Ramesh & Nelson 2024), a suite of cosmological magneto-hydrodynamical zoom-in simulations of Milky Way-like galaxies ($M_\star \sim 10^{10.9}\,M_\odot$, $M_{200c} \sim 10^{12.2}\,M_\odot$). In particular, we present a new RF4096 simulation, of a single halo, with preferential mass refinement that achieves a gas mass resolution of $\sim$225 $M_\odot$ in the CGM, defined as the region bounded between 0.15 R_{200c} and R_{200c} (virial radius). This is the latest addition to Project GIBLE, currently comprised of eight Milky Way-like galaxies each simulated at CGM gas mass resolutions of $\sim 10^3$, 10^4, and $10^5\,M_\odot$, labeled the RF512, RF64, and RF8 suites, respectively. In all cases, the galaxy is maintained at a resolution of $\sim 8.5 \times 10^5\,M_\odot$.

Project GIBLE uses the IllustrisTNG model (Weinberger et al. 2017; Pillepich et al. 2018), within the AREPO code (Springel 2010), to account for the physical processes that regulate galaxy formation and evolution. This includes radiative thermochemistry and metal cooling with the metagalactic radiation field, star formation, stellar evolution and enrichment, supermassive black hole (SMBH) formation, and feedback from stars (supernovae) and SMBHs (i.e., AGN; thermal, kinetic, and radiative modes).

The TNG model also includes ideal magneto-hydrodynamics (Pakmor et al. 2014). A uniform primordial field of 10^{-14} comoving Gauss is seeded at the start of the simulation, which is subsequently (self-consistently) amplified as a combined result of structure formation, small-scale dynamos, and feedback processes (Pakmor et al. 2020). While the initial field is divergence-free by construction, the Powell et al. (1999) cleaning scheme is used to maintain $\nabla \cdot \boldsymbol{B} = 0$ over time. We note that the relative divergence error is typically small ($\lesssim O(10^{-2})$), indicating that divergence errors are minimal (see also Pakmor & Springel 2013). The order of magnitude of field strengths predicted to be found in the gaseous halos around galaxies by the TNG model (Marinacci et al. 2018; Nelson et al. 2018; Ramesh et al. 2023b,c) is broadly consistent with recent indications for the existence of large-scale $\sim\mu$G magnetic fields in the CGM observed using Faraday rotation (Heesen et al. 2023; Böckmann et al. 2023).

2.2. Cloud and interface identification algorithm

Following Nelson et al. (2020) and Ramesh et al. (2023a), we defined and identified clouds as spatially contiguous sets of cold ($T \leq 10^{4.5}$ K) Voronoi cells (i.e., collections of cold gas cells that are Voronoi natural neighbors). Further, we considered only those gas cells that are not gravitationally bound to any satellite galaxies, as identified by the substructure identification algorithm SUBFIND (Springel et al. 2001). To maximize statistics and resolution throughout this work, we restricted the selection of clouds to those with masses in the range $[10^{5.8}, 10^{6.2}]\,M_\odot$, resulting in a sample size of 233. These clouds, on average, are resolved by $\sim$3700 gas cells, with more in the surrounding interfaces.

The interface layer of each cloud is defined as all gas cells that are not cold ($T > 10^{4.5}$ K), but share a face with a member cell of the cloud, determined using the Voronoi tessellation connectivity[1]. This interface layer is thus the cocoon of cells immediately surrounding the cloud. On average, the interface layers around our clouds are sampled by $\sim$2550 resolution elements.

We quantified the topology of magnetic field lines in the interface layer using measurements of magnetic curvature (κ), defined as (Shen et al. 2003; Pfrommer et al. 2022)

$$\kappa = |\boldsymbol{\kappa}| = |(\boldsymbol{b} \cdot \nabla)\boldsymbol{b}|, \tag{1}$$

where $\boldsymbol{b} = \boldsymbol{B}/B$ is the unit vector in the direction of the magnetic field $\boldsymbol{B}$. Vector $\boldsymbol{\kappa}$ points in the direction of the local center of curvature of $\boldsymbol{B}$, while κ corresponds to the inverse of the radius of curvature (Boozer 2005). A value of $\kappa = 0$ thus describes a straight field, while larger values denote field lines with increasing deviations from uniformity.

Following a calculation of $\kappa_{\rm cell}$ for each interface gas cell using Eq. (1), we compute the magnetic curvature around a cloud as the mean of all its interface gas cells. We denote this mean magnetic curvature as κ throughout the rest of the text.

3. Results

We begin with a visualization of the distribution of clouds in the CGM of our simulated halo in Fig. 1. The image, which shows one quadrant of the halo at $z = 0$[2], extends R_{200c} ($\sim$236 kpc) from edge-to-edge, and $\pm R_{200c}$ along the projection direction, with the center of the galaxy located at the top right corner. The color-coding shows the average (mass-weighted) temperature of gas along the line of sight. The small white circles give the positions of all clouds with masses greater than $10^5\,M_\odot$, with radii scaling with the mass of clouds. Of these many hundreds of clouds, the fiducial sample that we consider in this work, $10^{5.8} < M_{\rm cl}/M_\odot < 10^{6.2}$, is shown in gray. These cool clouds are embedded in the volume filling background CGM, which is $\sim$100 times hotter.

The inset shows a slice of the Voronoi mesh centered around a single cloud. We outlined the cells belonging to the cloud with white lines. The ratio of the physical scales of the two images is $\sim$60, and so the inset shows a highly zoomed-in region of the main image, but is still well resolved. Simulations of the kind shown here thus enable the study of small-scale cloud phenomena, including formation and evolution, mixing, among others, with clouds self-consistently evolved in the full cosmological context.

In Fig. 2 we show two examples of the topology of magnetic field lines around cool clouds. The two panels show slices of the Voronoi mesh centered around two distinct clouds. Both are oriented such that their velocities, computed as the mass-weighted mean velocity of all cloud member cells, point along the positive x-axis. Both clouds are infalling toward the center of the galaxy, and are at similar galactocentric distances ($\sim$50 kpc). Cells that belong to the clouds (interface layers) are outlined with black

[1] While interface gas cells are typically contiguous among themselves, there are rare cases where gaps may be present in the interface.

[2] We exclusively consider the $z = 0$ simulation snapshot to best connect with the observational Milky Way HVC community.

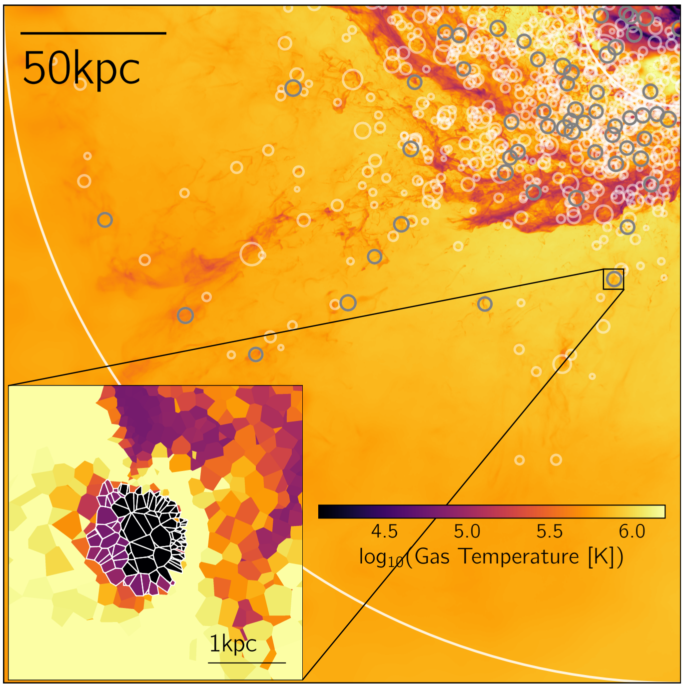

Fig. 1. Visualization of the distribution of clouds in a quadrant of our highest resolved GIBLE halo, a Milky Way-like galaxy at $z = 0$. The center of the galaxy is in the top right corner (main image). The image extends R_{200c} from edge-to-edge, and $\pm R_{200c}$ along the projection direction. The colors show mean mass-weighted gas temperatures in projection. The circles show the positions of the many hundred cold dense CGM gas clouds with masses greater than $10^5\,M_\odot$. Our fiducial sample with $M_{cl} \sim 10^6\,M_\odot$ is marked in gray. The inset, a highly zoomed-in region of the halo, shows a slice of the Voronoi mesh centered around a random cloud from our sample, with all member cells outlined by white lines. Despite their small sizes, the clouds (and their interface layers) were resolved with ~3700 (2550) gas cells, which enabled the study of small-scale phenomena self-consistently evolved in a cosmological context.

(white) lines. Streamlines show the direction of magnetic field lines in this plane, while the background color corresponds to the density of gas. We include three numbers at the top of each panel: the magnetic curvature averaged over all interface gas cells (κ), the ratio of the mean density of the cloud to that of the ambient background[3] (i.e., the density contrast, δ), and the modulus of the difference between the cloud velocity and that of the ambient background (i.e., the velocity contrast, v_{rel}).

The magnetic field lines around these clouds show contrasting structures. While the field is largely coherent in the left panel, quantified by a relatively low κ of $\sim$0.1 kpc^{-1}, the topology is more complex on the right with tangled and less laminar field lines ($\kappa \sim 1.87$ kpc^{-1}). The numbers above each panel show that κ correlates with, among other properties, contrasts in both density and velocity, which we consider further in Fig. 4.

In the main panel of Fig. 3 we explore the variation in values of κ across our sample. The black curve corresponds to the computation of κ as the mean over all interface gas cells (i.e., our fiducial definition). The other curves instead compute averages of a sub-selection of interface gas cells from ten bins, constructed based on their relative position in the direction of motion of the cloud, as shown by the colorbar. A relative position of 0.1 would thus correspond to the first 10% of interface gas cells in the direction of motion, 0.2 to the next 10%, and so on.

The black distribution peaks at a value of κ of $\lesssim$0.1 kpc^{-1}, is relatively flat out to $\kappa \sim 0.7$ kpc^{-1}, and drops steadily toward higher values. A significant fraction of clouds are thus predicted to have largely coherent fields around them in their interface regions, as in the left panel of Fig. 2. The colored curves show similar behavior, with little dependence on relative position. That is, the degree of curvature does not strongly change between the upstream and downstream interface regions. Previous studies with idealized simulations have shown that field line draping around clouds moving through an initially uniform magnetic field perpendicular to the motion of the cloud is more (less) efficient in the head (tail) (e.g., Jung et al. 2023), corresponding to larger (smaller) values of κ upstream (downstream) of the cloud. The insignificant difference in the distributions of κ between these regions in Fig. 3 suggests that such strong draping configurations are not common around our simulated clouds. We speculate that this is largely due to the background field lines upstream of the cloud not being oriented in perpendicular directions and as uniformly as is typically assumed in idealized setups (see also Sparre et al. 2020). We note here that the distributions explored in Fig. 3 are largely converged up to our RF512 simulations (i.e., eight times lower mass resolution, not shown explicitly).

As the value of κ corresponds to the inverse of the radius of curvature of the local magnetic field, it should scale inversely with the radius of the cloud for strongly draped configurations. However, we checked that the above results are qualitatively similar when values of κ are normalized[4] by $1/R$ (see also the large diversity of κ at fixed R in Fig. 4). The κ distribution therefore reflects physically different field geometries in the cloud interfaces.

In the inset of Fig. 3, we make a comparison to a simple model describing the expected field around a sphere of radius R moving through a homogeneous ambient medium with an initial uniform magnetic field oriented perpendicular to the cloud motion (Dursi & Pfrommer 2008). Given the assumptions of this model, the solution is not valid in the wake behind the sphere. We thus restrict our comparison to the region around the head of the cloud, which we define using the angular bounds $\theta = [-\pi/3, \pi/3]$ and $\phi = [-\pi/3, \pi/3]$. Although this choice is arbitrary, we note that adopting other angular ranges for θ and ϕ has no significant impact on the analysis that follows. We compute the theoretical estimate of the model, κ_{DP08}, for each cloud separately by setting R to the effective radius of the cloud.

The curve shows the distribution of the ratio $\kappa_{DP08}/\kappa_{GIBLE}$. The PDF peaks around -0.7, with a large spread, and the model typically underpredicts the curvature seen in the simulation. Although not shown here, we find a weak anti-correlation in this ratio with δ and v_{rel}[5], suggesting that the model works better for clouds with relatively small density and velocity contrasts. We speculate that this could be linked to the more efficient build-up of magnetic fields around objects of greater overdensities and velocity contrasts (Lyutikov 2006), thereby increasing the impact of magnetic back-reaction on the flow, which is not considered in the Dursi & Pfrommer (2008) model.

Finally, we return to the correlation between κ and properties of the interface and the ambient background. From left to right, Fig. 4 shows κ as a function of the density contrast (δ),

[3] We define the ambient background as being comprised of three layers of non-cold gas cells around clouds.

[4] Following Nelson et al. (2020), we define the effective radius by the volume equivalent sphere, $R = [3V_{cloud}/4\pi]^{1/3}$.

[5] The value of κ predicted by the Dursi & Pfrommer (2008) model only depends on R, and does not take δ and v_{rel} as input parameters.

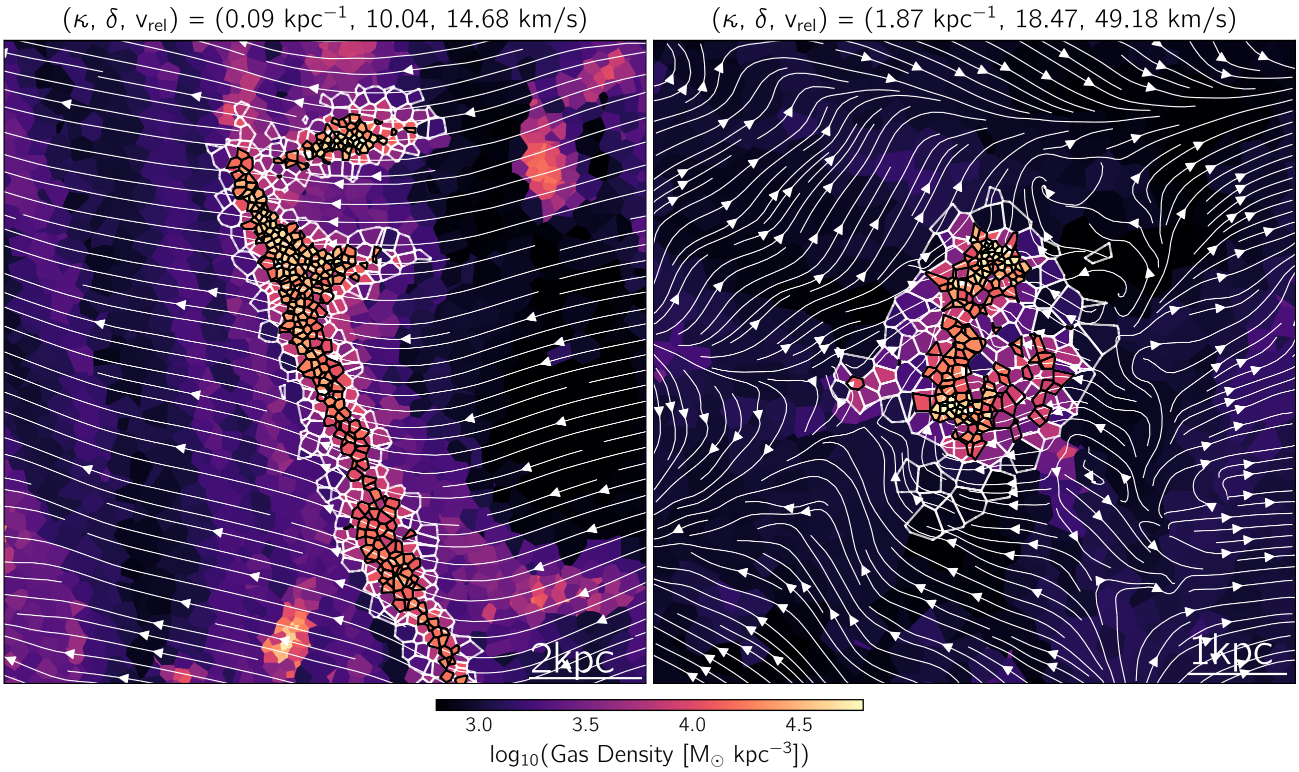

Fig. 2. Visualization of the diverse topologies of magnetic field lines around cold CGM clouds. The two panels show slices of the Voronoi mesh centered around two clouds, both oriented such that their mean velocity vector (i.e., the direction of motion) are to the right, the positive x-axis direction. Cells that belong to the cloud (interface layer) are outlined using black (white) lines. Streamlines show the direction of magnetic field lines. The three numbers at the top of each image correspond to the mean magnetic curvature of the cloud interface (κ), the overdensity (δ), and relative velocity ($v_{\rm rel}$) between the cloud and the interface layer. While the left cloud is threaded by magnetic fields in a region of the background CGM that has particularly uniform field orientation, the magnetic fields in the right panel begin to respond to the motion of the cloud. This diversity is captured by the different values of κ.

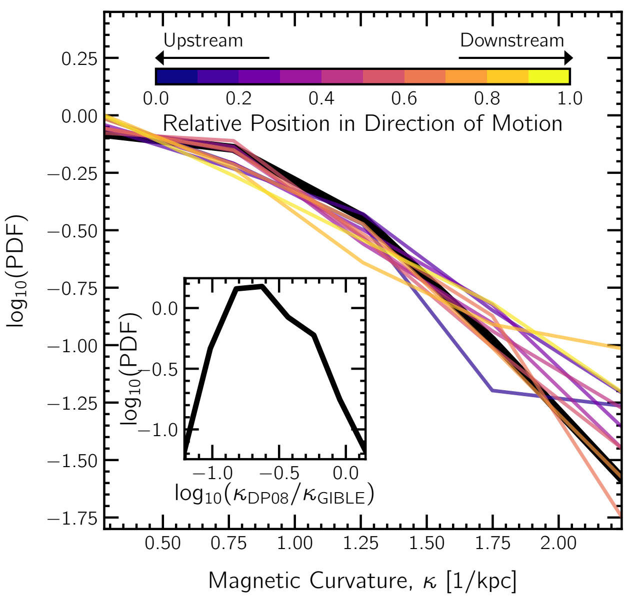

Fig. 3. Distribution of interface magnetic curvature values for our sample of $M_{\rm cl} \sim 10^6\,M_\odot$ clouds (main panel). The black curve is based on all interface gas, while the other curves show values derived using gas with different relative positions with respect to the direction of motion of the cloud. The purple curves show κ for the head or upstream regions, while the yellow curves show κ in the tail or downstream regions. The inset compares the upstream interface magnetic curvature from our simulations ($\kappa_{\rm GIBLE}$) to a simple theoretical model ($\kappa_{\rm DP08}$).

vorticity in the interface layer ($\omega = \nabla \times v$), the thermal-to-magnetic pressure ratio of the interface (β), and the radius of the cloud (R). The solid curve shows the median, while the colored points correspond to individual clouds, each colored by $v_{\rm rel}$. On average, κ increases almost linearly with δ. The scatter clearly correlates with relative velocity: larger values of κ have higher $v_{\rm rel}$ at fixed δ. A linear correlation is also present between κ and ω, indicating a possible connection between a disordered velocity field and a disordered magnetic field. Consistent with theoretical predictions for the case of an initial magnetic field that is coherent on scales larger than the cloud size (McCourt et al. 2015), we find that κ correlates with β. A least-squares fit yields $\kappa \propto \beta^{0.9}$, roughly in the ballpark of the predicted $\kappa \propto \beta$ trend by Schekochihin et al. (2004) for the case of small-scale turbulent dynamos. While κ decreases with increasing R, the drop is steeper than the $1/R$ trend expected for a strongly draped configuration (see discussion above), suggesting again that such configurations are not common in our sample. As before, we find that the results shown here are largely converged up to our RF512 level runs.

4. Discussion

The sizes, density contrasts, and kinematics of clouds may thus have an important impact on the structure of ambient field lines. The resulting magnetic field topology may in turn affect cloud growth and evolution. For example, the draped magnetic field layer increases the drag force by a factor of $\sim[1 + (v_{\rm A}/v_{\rm rel})^2] \equiv [1+(R\kappa)^{-1}]$, where $v_{\rm A}$ is the Alfven speed in the background. The

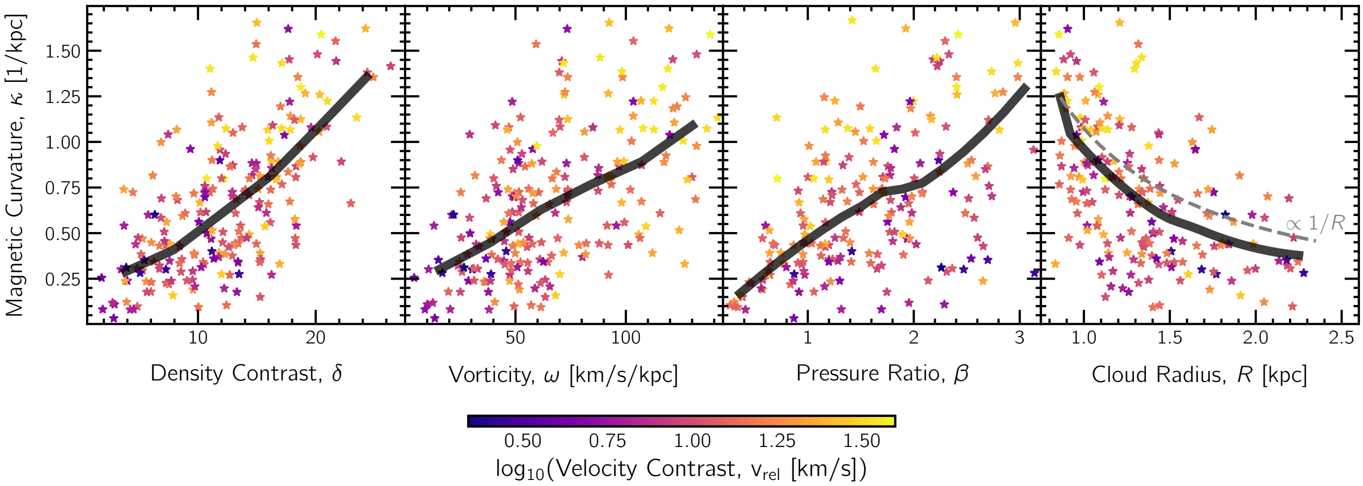

Fig. 4. Magnetic curvature (κ) as a function of (from left to right) density contrast between the cloud and its ambient background, vorticity in the interface layer, thermal-to-magnetic pressure ratio of the interface, and cloud radius. The solid curves show the median, while the scatter points correspond to individual clouds, each colored by the velocity contrast between the cloud and its ambient background. A strong trend of κ is seen in the median with respect to each of the properties considered, while the variation of κ at fixed abscissa is typically well captured by the diversity of velocity contrasts of clouds.

enhanced drag force decreases the "stopping distance" by $[1 + (R\kappa)^{-1}]^{-1}$ (i.e., the distance travelled by the cloud prior to achieving velocity equilibrium with its surroundings), thereby improving the chances of their survivability (McCourt et al. 2015). The diversity of cloud and interface properties portrayed by Fig. 4 suggests that the impact of field line topology on a population of clouds is expected to be varied. Specifically, at fixed R, clouds with lower δ, β and $v_{\rm rel}$ typically have lower values of $R\kappa$, and would thus experience a larger boost in their drag force compared to high $\delta/\beta/v_{\rm rel}$ counterparts[6]. We reiterate that cosmological simulations like GIBLE allow us to assess such predictions for an actual diverse cloud population since clouds, their interfaces, and their magnetic fields evolve self-consistently.

The magnetic field topology and draped layers may furthermore play a role in suppressing the impact of the Kelvin–Helmholtz instability along the surface of the cloud (e.g., Pfrommer & Dursi 2010). For instance, at fixed $B_{\rm wind}$, Sparre et al. (2020) showed that clouds with draped topologies in their interfaces are expected to survive longer. In addition, Jung et al. (2023) find that regions where field lines are inefficiently draped (i.e., low values of κ) fragment rapidly into smaller clumps, while regions of high κ are instead extended into long filamentary structures as a result of enhanced magnetic tension and effectively survive longer. However, it is important to note that the net impact of this suppression of mixing on the evolution of clouds may depend on the efficiency of radiative cooling in the interface layer (e.g., Gronke & Oh 2018). In agreement with idealized work, clouds in TNG50 have temperature gradients into their interfaces (i.e., they are surrounded by a mixed-phase layer of warm gas that rapidly cools onto the cloud; Nelson et al. 2020), consistent with theoretical local cooling flow models (Dutta et al. 2022). Moreover, the metal content of clouds and their interfaces can vary significantly (Nelson et al. 2020; Ramesh et al. 2023a), possibly affecting the rate at which the gas in the interface condenses. Future work will quantify the resulting impact on cloud growth and survival in our simulations with Lagrangian tracers (Ramesh et al., in prep.).

While we find that clouds are roughly in pressure balance with their interface layers (Ramesh et al. 2023a), thereby preventing them from being crushed and dissolved, the TNG model does not include thermal conduction. The inclusion of this component may contribute to cloud evaporation (e.g., Marcolini et al. 2005; Vieser & Hensler 2007), although certain configurations of magnetic fields may partially suppress this effect (Ettori & Fabian 2000; Brüggen et al. 2023). Future simulations that include conduction can explore its role in cosmological cloud evolution. This will require that we adequately resolve the Field length (Field 1965) to avoid spurious numerical effects (Koyama & Inutsuka 2004). For example, for 10% Spitzer conduction (see, e.g., Brüggen et al. 2023) the Field length would be ~120 pc for interface gas cells, requiring a spatial resolution of $\lesssim$40 pc in this region[7] (i.e., 2–4× better spatial resolution than our current RF4096 run; see Fig. 2 of Ramesh & Nelson 2024).

5. Summary

In this paper we used a cosmological zoom-in galaxy formation simulation with additional CGM refinement to study and explore the complex and diverse topology of magnetic field lines around cold, dense clouds in the CGM. At an average baryonic mass resolution of $\sim 225\,M_\odot$, the interface layers around our sample of $10^{5.8} < M_{\rm cl}/M_\odot < 10^{6.2}$ clouds are resolved by over 2000 resolution elements, allowing the study of interface phenomena in a cosmological context.

We quantified the structure of magnetic field lines around clouds (i.e., in interface layers) by the magnetic curvature κ. We find that values of κ vary significantly, reflecting the diversity in field line topologies around clouds. There is no significant difference in the distribution of κ between the regions upstream and downstream of the cloud, suggesting that strong draping configurations are rare in our sample. However, curvature correlates

[6] This only describes the enhancement factor of the drag force as a result of draping. The total drag force experienced by the cloud ($\sim\rho^2_{\rm interface}v^2_{\rm rel}R^2[1 + (R\kappa)^{-1}]$) depends on other properties of the cloud and of the interface.

[7] This Field condition, that spatial resolution is better than the Field length by at least a factor of 3, was derived using one-dimensional simulations with isotropic conduction (Koyama & Inutsuka 2004).

strongly with cloud-background contrasts in density and velocity: greater contrasts correspond to higher κ, on average. In addition, κ also correlates with other interface properties, including vorticity and the thermal-to-magnetic pressure ratio.

This study provides a first perspective from the point of view of cosmological simulation regarding the topology of magnetic field lines around cold clouds. However, there are several clear avenues to extend this work. In particular, we can assess the impact of cloud motion on the immediate interface layer, and on the broader local gaseous environment of clouds. With Lagrangian tracers we can also quantify the impact of magnetic fields on the lifetime, survival, and evolution of clouds.

Acknowledgements. RR and DN acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Group (grant number NE 2441/1-1). RR is a Fellow of the International Max Planck Research School for Astronomy and Cosmic Physics at the University of Heidelberg (IMPRS-HD). MB acknowledges support from the Deutsche Forschungsgemeinschaft under Germany's Excellence Strategy – EXC 2121 "Quantum Universe" – 390833306 and from the BMBF ErUM-Pro grant 05A2023. This work has made use of the VERA supercomputer of the Max Planck Institute for Astronomy (MPIA), and the COBRA supercomputer, both operated by the Max Planck Computational Data Facility (MPCDF), and of NASA's Astrophysics Data System Bibliographic Services.

References

Banda-Barragán, W. E., Parkin, E. R., Federrath, C., Crocker, R. M., & Bicknell, G. V. 2016, MNRAS, 455, 1309
Berlok, T., & Pfrommer, C. 2019, MNRAS, 489, 3368
Böckmann, K., Brüggen, M., Heesen, V., et al. 2023, A&A, 678, A56
Boozer, A. H. 2005, Rev. Mod. Phys., 76, 1071
Brüggen, M., Scannapieco, E., & Grete, P. 2023, ApJ, 951, 113
Das, H. K., & Gronke, M. 2024, MNRAS, 527, 991
Donahue, M., & Voit, G. M. 2022, Phys. Rep., 973, 1
Dursi, L. J., & Pfrommer, C. 2008, ApJ, 677, 993
Dutta, A., Sharma, P., & Nelson, D. 2022, MNRAS, 510, 3561
Ettori, S., & Fabian, A. C. 2000, MNRAS, 317, L57
Field, G. B. 1965, ApJ, 142, 531
Fielding, D. B., Ostriker, E. C., Bryan, G. L., & Jermyn, A. S. 2020, ApJ, 894, L24
Fielding, D. B., Ripperda, B., & Philippov, A. A. 2023, ApJ, 949, L5
Girichidis, P. 2021, MNRAS, 507, 5641
Gronke, M., & Oh, S. P. 2018, MNRAS, 480, L111
Grønnow, A., Tepper-García, T., Bland-Hawthorn, J., & Fraternali, F. 2022, MNRAS, 509, 5756
Heesen, V., O'Sullivan, S. P., Brüggen, M., et al. 2023, A&A, 670, L23
Hidalgo-Pineda, F., Farber, R. J., & Gronke, M. 2024, MNRAS, 527, 135
Jung, S. L., Grønnow, A., & McClure-Griffiths, N. M. 2023, MNRAS, 522, 4161
Klein, R. I., McKee, C. F., & Colella, P. 1994, ApJ, 420, 213
Koyama, H., & Inutsuka, S.-I. 2004, ApJ, 602, L25
Kwak, K., Shelton, R. L., & Raley, E. A. 2009, ApJ, 699, 1775
Lyutikov, M. 2006, MNRAS, 373, 73
Marcolini, A., Strickland, D. K., D'Ercole, A., Heckman, T. M., & Hoopes, C. G. 2005, MNRAS, 362, 626
Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, MNRAS, 480, 5113
McCourt, M., O'Leary, R. M., Madigan, A.-M., & Quataert, E. 2015, MNRAS, 449, 2
Mellema, G., Kurk, J. D., & Röttgering, H. J. A. 2002, A&A, 395, L13
Muller, C. A., Oort, J. H., & Raimond, E. 1963, Acad. Sci. Paris Comptes Rendus, 257, 1661
Nelson, D., Pillepich, A., Springel, V., et al. 2018, MNRAS, 475, 624
Nelson, D., Sharma, P., Pillepich, A., et al. 2020, MNRAS, 498, 2391
Pakmor, R., & Springel, V. 2013, MNRAS, 432, 176
Pakmor, R., Marinacci, F., & Springel, V. 2014, ApJ, 783, L20
Pakmor, R., van de Voort, F., Bieri, R., et al. 2020, MNRAS, 498, 3125
Pfrommer, C., & Dursi, J. 2010, Nat. Phys., 6, 520
Pfrommer, C., Werhahn, M., Pakmor, R., Girichidis, P., & Simpson, C. M. 2022, MNRAS, 515, 4229
Pillepich, A., Springel, V., Nelson, D., et al. 2018, MNRAS, 473, 4077
Powell, K. G., Roe, P. L., Linde, T. J., Gombosi, T. I., & De Zeeuw, D. L. 1999, J. Comput. Phys., 154, 284
Ramesh, R., & Nelson, D. 2024, MNRAS, 528, 3320
Ramesh, R., Nelson, D., & Pillepich, A. 2023a, MNRAS, 522, 1535
Ramesh, R., Nelson, D., Heesen, V., & Brüggen, M. 2023b, MNRAS, 526, 5483
Ramesh, R., Nelson, D., & Pillepich, A. 2023c, MNRAS, 518, 5754
Scannapieco, E., & Brüggen, M. 2015, ApJ, 805, 158
Schekochihin, A. A., Cowley, S. C., Taylor, S. F., Maron, J. L., & McWilliams, J. C. 2004, ApJ, 612, 276
Shen, C., Li, X., Dunlop, M., et al. 2003, J. Geophys. Res. (Space Phys.), 108, 1168
Shin, M.-S., Stone, J. M., & Snyder, G. F. 2008, ApJ, 680, 336
Sparre, M., Pfrommer, C., & Ehlert, K. 2020, MNRAS, 499, 4261
Springel, V. 2010, MNRAS, 401, 791
Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, MNRAS, 328, 726
Vieser, W., & Hensler, G. 2007, A&A, 472, 141
Wakker, B. P., & van Woerden, H. 1997, ARA&A, 35, 217
Weinberger, R., Springel, V., Hernquist, L., et al. 2017, MNRAS, 465, 3291

Chapter 7

Summary and Outlook

The circumgalactic medium (CGM), the tenuous reservoir of gas around galaxies is not only multi-scale and multi-phase, but also multi-physics, as it is impacted by, and its evolution is governed by, a multitude of processes. While this leads to a number of interesting phenomena, it also makes accurately modeling and simulating the CGM, and galaxies in general, an arduous task.

Cosmological (magneto)hydrodynamical simulations over the past decade have become powerful tools for studying the growth and evolution of large-scale structure over cosmic time. This progress has been made possible not only by advancements in computing technologies, but also by the development of new numerical techniques and models. Although these simulations rely on several assumptions regarding the physics implemented, they allow for the self-consistent evolution of galaxies and their CGMs in the full ΛCDM context.

In this thesis, we use simulations run with the IllustrisTNG model to explore various questions regarding the CGM. In Chapter 2, we begin by analyzing the physical properties of gas in the CGM of 132 MW-likes in the TNG50 simulation [186, 19], which evolves a volume of $\sim (50\,\mathrm{Mpc})^3$ at a baryonic mass resolution of $\sim 10^5\,\mathrm{M}_\odot$. We find significant diversity across the sample in all quantities explored, and that feedback from the SMBH plays a key role: energy from SMBH-driven kinetic winds not only generates high-velocity outflows ($\gtrsim 500 - 2000$ km s^{-1}), but also heats gas to super-virial temperatures ($> 10^{6.5-7}$ K) and regulates the net balance of inflows versus outflows in otherwise quasi-static gaseous halos. In contrast, thermal-mode feedback has a less pronounced impact, and high-temperature/velocity tails of the distribution are absent in galaxies where the SMBH is inactive in the kinetic state.

In Chapter 3, we focus exclusively on cold clouds in the CGM of these 132

MW-like galaxies, studying their radial and size distribution, physical properties, and contrast in properties with the ambient background. We also make a crude zeroth-order comparison with available observational data of clouds in the *real* Milky Way, finding that the distribution of cloud radial velocities from TNG50 matches those of observations to a very high degree. As before, SMBH-driven feedback may play a role, with galaxies hosting larger SMBHs having systematically fewer clouds in their CGMs. In Chapter 4, we take a final look at the impact of feedback on the CGM by studying the angular anisotropy of CGM magnetic field strengths at a fixed impact parameter. In qualitative agreement with recent observational results, we find that gas along the minor axes is relatively more magnetized, as a result of SN-driven outflows preferentially flowing along these directions.

In Chapter 5, we introduce a new suite of simulations called GIBLE, a set of cosmological zoom-ins with added refinement in the CGM. In our primary runs, we achieve a CGM gas mass resolution of $\sim 10^3$ M$_\odot$, a level comparable to some of the best cosmological zoom-in simulations of MW-likes run to date. In Chapter 6, we present the next addition to GIBLE, where we achieve an ultra-high resolution of $\sim 10^2$ M$_\odot$ in the CGM, and focus on the topology of magnetic field lines around cold clouds.

While all the above studies led to interesting results, there are clear avenues for future work and improvements, particularly with the physics model employed. As mentioned in § 1.3.3, the IllustrisTNG model is calibrated to statistically reproduce a set of integrated galaxy properties, and everything else derived from these simulations is a free prediction of the model. While certain other galaxy and halo properties are in decent agreement with available observational constraints, there is no clear consensus regarding many other quantities.

For instance, other large-volume cosmological simulations such as EAGLE [187] and SIMBA [188], which employ different models – particularly in the implementation of galactic feedback – also statistically reproduce a given set of $z = 0$ galaxy properties, i.e., those against which they are calibrated, among others. However, many properties of the CGM differ significantly between these different simulations, such as the length scale over which baryons are distributed around galaxies [55].

Another example is the strength of CGM magnetic fields. As discussed in Chapters 2 and 4, CGMs in TNG are typically more magnetized than those in the FIRE-2 simulations [189], although both predict much weaker fields than available observational constraints. The reason for this remains unclear, but it is very likely a combination of the physics implemented, the resolution achieved, and the numerical techniques used to evolve the simulation forward in time. Given that magnetic fields may play a role in the evolution of galaxies and their surrounding CGMs, it is crucial that more progress is made in this direction. This seems a promising future avenue, especially with facilities such as ASKAP [190] and SKA [191] soon to begin operations, thus providing more observational constraints to advance existing models.

The IllustrisTNG model also does not explicitly account for contributions by various processes. For instance, the injection and evolution of cosmic rays (CRs) and their coupling to gas are ignored. As discussed in § 1.2.6, CRs may play a role in the evolution of galaxies and their CGMs, and including these components is an important step toward building a more complete model. In [134], we recently made a stride in this direction by implementing a simple CR model into AREPO. However, we exclusively used SNe as CR sources, neglecting possible contributions from SMBHs, and thus focused only on the low halo mass regime. Moreover, despite testing many variations, we were unable to produce a realistic sample of galaxies, at least with CR transport parameters that fall under available constraints. This was likely due to the simplicity of the model, and explorations with more sophisticated CR schemes are required to build a better and more complete understanding of the impact of CRs on cosmic growth.

Another interesting scientific direction is to better understand the origin and evolution of cold clouds in cosmological simulations, a topic we made a first attempt at in [141] using Project GIBLE. A key scientific result of this work was that a large fraction of clouds ($\sim 40-60$ %,, depending on mass) originate through the condensation of hot halo gas. However, it remains unclear how this condensation is initiated: theory suggests that it may occur either through the upliftment of low-entropy gas to higher altitudes or due to local thermal instabilities in the presence of strong seed perturbations [192]. Future efforts will aim to explore this further with the current simulations, while also expanding the GIBLE sample to more massive halos, where CGM gas temperatures are significantly higher. In this regard, including the physics of thermal conduction will likely lead to interesting effects: in addition to suppressing thermal instabilities [102] and offsetting radiative cooling [193], it may also lead to the evaporation of cold clouds [194]. Future iterations of the physics model will make it possible to probe such phenomena for clouds evolved self-consistently in a full ΛCDM context, as well as the ultimate impact on galaxy evolution.

References

[1] N. Aghanim et al., *Astronomy & Astrophysics*, 641, A6, 2020.
[2] P. J. E. Peebles, *ApJ*, 147, p. 859, 1967.
[3] P. J. E. Peebles. *The large-scale structure of the universe*. 1980.
[4] K. G. Begeman, *AAP*, 223, pp. 47–60, 1989.
[5] P. F. de Salas et al., *JCAP*, 2019, p. 037, 2019.
[6] S. M. Faber and R. E. Jackson, *ApJ*, 204, pp. 668–683, 1976.
[7] P. L. Schechter, *AJ*, 85, pp. 801–811, 1980.
[8] P. Schneider et al., *AAP*, 333, pp. 767–778, 1998.
[9] X.-P. Wu et al., *MNRAS*, 301, pp. 861–871, 1998.
[10] G. Münch and H. Zirin, *ApJ*, 133, p. 11, 1961.
[11] M. Schmidt, *Nature*, 197, p. 1040, 1963.
[12] J. N. Bahcall and L. Spitzer Jr., *ApJL*, 156, p. L63, 1969.
[13] W. L. W. Sargent et al., *ApJS*, 42, pp. 41–81, 1980.
[14] H.-W. Chen et al., *ApJ*, 498, pp. 77–94, 1998.
[15] S. S. McGaugh et al., *ApJL*, 708, pp. L14–L17, 2010.
[16] P. S. Behroozi, C. Conroy, and R. H. Wechsler, *ApJ*, 717, pp. 379–403, 2010.
[17] B. D. Oppenheimer and R. Davé, *MNRAS*, 387, pp. 577–600, 2008.
[18] C. Péroux and J. C. Howk, *ARAA*, 58, pp. 363–406, 2020.
[19] D. Nelson et al., *MNRAS*, 490, pp. 3234–3261, 2019.
[20] J. J. Davies et al., *MNRAS*, 491, pp. 4462–4480, 2020.
[21] D. Kereš et al., *MNRAS*, 363, pp. 2–28, 2005.
[22] D. Nelson et al., *MNRAS*, 429, pp. 3353–3370, 2013.
[23] L. Mayer et al., *MNRAS*, 369, pp. 1021–1038, 2006.
[24] E. Rohr et al., *MNRAS*, 524, pp. 3502–3525, 2023.
[25] J. Binney, C. Nipoti, and F. Fraternali, *MNRAS*, 397, pp. 1804–1815, 2009.
[26] F. Fraternali et al., *MNRAS*, 447, pp. L70–L74, 2015.
[27] C. A. Muller, J. H. Oort, and E. Raimond, *Academie des Sciences Paris Comptes Rendus*, 257, pp. 1661–1662, 1963.

[28] J. E. G. Peek et al., *ApJ*, 692, pp. 827–838, 2009.
[29] D. B. Henley et al., *ApJ*, 723, pp. 935–953, 2010.
[30] P. Kaaret et al., *Nature Astronomy*, 4, pp. 1072–1077, 2020.
[31] P. Predehl et al., *Nature*, 588, pp. 227–231, 2020.
[32] J. H. Black, *MNRAS*, 197, pp. 553–563, 1981.
[33] N. Katz, D. H. Weinberg, and L. Hernquist, *ApJS*, 105, p. 19, 1996.
[34] R. P. C. Wiersma, J. Schaye, and B. D. Smith, *MNRAS*, 393, pp. 99–107, 2009.
[35] B. Smith, S. Sigurdsson, and T. Abel, *MNRAS*, 385, pp. 1443–1454, 2008.
[36] M. J. Rees, *MNRAS*, 218, 25P–30P, 1986.
[37] G. Efstathiou, *MNRAS*, 256, 43P–47P, 1992.
[38] C.-A. Faucher-Giguère et al., *ApJ*, 703, pp. 1416–1443, 2009.
[39] F. Haardt and P. Madau, *ApJ*, 746, p. 125, 2012.
[40] S. Woosley and T. Janka, *Nature Physics*, 1, pp. 147–154, 2005.
[41] H. T. Janka et al., *Physics Reports*, 442, pp. 38–74, 2007.
[42] S. Chandrasekhar, *ApJ*, 74, p. 81, 1931.
[43] E. H. Lieb and H.-T. Yau, *ApJ*, 323, p. 140, 1987.
[44] J. C. Wheeler. *Cosmic Catastrophes*. 2000.
[45] W. Hillebrandt and J. C. Niemeyer, *ARAA*, 38, pp. 191–230, 2000.
[46] P. A. Mazzali et al., *Science*, 315, p. 825, 2007.
[47] M. Modjaz, C. P. Gutiérrez, and I. Arcavi, *Nature Astronomy*, 3, pp. 717–724, 2019.
[48] A. Nyholm et al., *A&A*, 637, A73, 2020.
[49] V. Springel and L. Hernquist, *MNRAS*, 339, pp. 289–311, 2003.
[50] P. F. Hopkins et al., *MNRAS*, 522, pp. 2936–2950, 2023.
[51] R. Davé, *MNRAS*, 385, pp. 147–160, 2008.
[52] A. Katsianis, E. Tescari, and J. S. B. Wyithe, *MNRAS*, 448, pp. 3001–3021, 2015.
[53] D. Fielding et al., *MNRAS*, 466, pp. 3810–3826, 2017.
[54] V. Pandya et al., *MNRAS*, 508, pp. 2979–3008, 2021.
[55] M. Ayromlou, D. Nelson, and A. Pillepich, *MNRAS*, 524, pp. 5391–5410, 2023.
[56] A. Muzzin et al., *ApJ*, 777, p. 18, 2013.
[57] B. Darvish et al., *ApJ*, 837, p. 16, 2017.
[58] N. M. Förster Schreiber et al., *ApJ*, 787, p. 38, 2014.
[59] R. Genzel et al., *ApJ*, 796, p. 7, 2014.

[60] J. R. Oppenheimer and G. M. Volkoff, *Physical Review*, 55, pp. 374–381, 1939.

[61] F. Hoyle and R. A. Lyttleton, *Proceedings of the Cambridge Philosophical Society*, 35, p. 405, 1939.

[62] H. Bondi and F. Hoyle, *MNRAS*, 104, p. 273, 1944.

[63] F. Pacucci and A. Loeb, *ApJ*, 895, p. 95, 2020.

[64] X.-B. Wu et al., *Nature*, 518, pp. 512–515, 2015.

[65] R. L. Larson et al., *ApJL*, 953, p. L29, 2023.

[66] P. Madau and M. J. Rees, *ApJL*, 551, pp. L27–L30, 2001.

[67] A. Heger et al., *ApJ*, 591, pp. 288–300, 2003.

[68] M. C. Begelman, M. Volonteri, and M. J. Rees, *MNRAS*, 370, pp. 289–298, 2006.

[69] J. L. Johnson et al., *ApJ*, 771, p. 116, 2013.

[70] B. Carr and F. Kühnel, *Annual Review of Nuclear and Particle Science*, 70, pp. 355–394, 2020.

[71] J. Kormendy and D. Richstone, *ARAA*, 33, p. 581, 1995.

[72] J. Kormendy and L. C. Ho, *ARAA*, 51, pp. 511–653, 2013.

[73] D. Lynden-Bell and J. E. Pringle, *MNRAS*, 168, pp. 603–637, 1974.

[74] S. A. Balbus and J. F. Hawley, *ApJ*, 376, p. 214, 1991.

[75] N. I. Shakura and R. A. Sunyaev, *A&A*, 24, pp. 337–355, 1973.

[76] I. D. Novikov and K. S. Thorne. "Astrophysics of black holes." In: *Black Holes (Les Astres Occlus)*. Ed. by C. Dewitt and B. S. Dewitt. Jan. 1973, pp. 343–450.

[77] R. Narayan and R. Popham, *Nature*, 362, pp. 820–822, 1993.

[78] M. Jaroszynski, M. A. Abramowicz, and B. Paczynski, *Acta Astronomica,* 30, pp. 1–34, 1980.

[79] M. A. Abramowicz et al., *ApJ*, 332, p. 646, 1988.

[80] Y. Dubois et al., *A&A*, 651, A109, 2021.

[81] R. Penrose and R. M. Floyd, *Nature Physical Science*, 229, pp. 177–179, 1971.

[82] R. D. Blandford and D. G. Payne, *MNRAS*, 199, pp. 883–903, 1982.

[83] R. D. Blandford and R. L. Znajek, *MNRAS*, 179, pp. 433–456, 1977.

[84] M. Sikora, Ł. Stawarz, and J.-P. Lasota, *ApJ*, 658, pp. 815–828, 2007.

[85] N. J. McConnell and C.-P. Ma, *ApJ*, 764, p. 184, 2013.

[86] G. L. Granato et al., *ApJ*, 600, pp. 580–594, 2004.

[87] B. R. McNamara and P. E. J. Nulsen, *ARAA*, 45, pp. 117–175, 2007.

[88] D. Kawata and B. K. Gibson, *MNRAS*, 358, pp. L16–L20, 2005.

[89] D. Sijacki et al., *MNRAS*, 380, pp. 877–900, 2007.

[90] T. Okamoto, R. S. Nemmen, and R. G. Bower, *MNRAS*, 385, pp. 161–180, 2008.

[91] R. S. Somerville et al., *MNRAS*, 391, pp. 481–506, 2008.

[92] S. L. Shapiro, A. P. Lightman, and D. M. Eardley, *ApJ*, 204, pp. 187–199, 1976.

[93] S. Ichimaru, *ApJ*, 214, pp. 840–855, 1977.

[94] R. Weinberger et al., *MNRAS*, 465, pp. 3291–3308, 2017.

[95] M. S. Turner and L. M. Widrow, *PRD*, 37, pp. 2743–2754, 1988.

[96] K. Subramanian, *Reports on Progress in Physics*, 79, p. 076901, 2016.

[97] L. Biermann, *Zeitschrift Naturforschung Teil A*, 5, p. 65, 1950.

[98] M. Lazar et al., *Journal of Plasma Physics*, 75, pp. 19–33, 2009.

[99] R. Pakmor et al., *MNRAS*, 469, pp. 3185–3199, 2017.

[100] F. Marinacci et al., *MNRAS*, 480, pp. 5113–5139, 2018.

[101] R. Pakmor et al., *MNRAS*, 498, pp. 3125–3137, 2020.

[102] P. Sharma, I. J. Parrish, and E. Quataert, *ApJ*, 720, pp. 652–665, 2010.

[103] D. Nelson et al., *MNRAS*, 498, pp. 2391–2414, 2020.

[104] S. Ji, S. P. Oh, and M. McCourt, *MNRAS*, 476, pp. 852–867, 2018.

[105] M. Sparre, C. Pfrommer, and K. Ehlert, *MNRAS*, 499, pp. 4261–4281, 2020.

[106] K. Kwak, R. L. Shelton, and E. A. Raley, *ApJ*, 699, pp. 1775–1788, 2009.

[107] W. E. Banda-Barragán et al., *MNRAS*, 455, pp. 1309–1333, 2016.

[108] M. Brüggen, E. Scannapieco, and P. Grete, *ApJ*, 951, p. 113, 2023.

[109] M. McCourt et al., *MNRAS*, 449, pp. 2–7, 2015.

[110] F. van de Voort et al., *MNRAS*, 501, pp. 4888–4902, 2021.

[111] U. P. Steinwandel et al., *MNRAS*, 494, pp. 4393–4412, 2020.

[112] F. Vazza et al., *MNRAS*, 445, pp. 3706–3722, 2014.

[113] M. Ruszkowski and C. Pfrommer, *ARAA*, 31, p. 4, 2023.

[114] V. L. Ginzburg and S. I. Syrovatskii. *The Origin of Cosmic Rays*. 1964.

[115] D. P. Cox, *ARAA*, 43, pp. 337–385, 2005.

[116] T. Naab and J. P. Ostriker, *ARAA*, 55, pp. 59–109, 2017.

[117] A. V. Ivlev et al., *ApJ*, 855, p. 23, 2018.

[118] M. Padovani et al., *SSR*, 216, p. 29, 2020.

[119] V. N. Zirakashvili et al., *A&A*, 311, pp. 113–126, 1996.

[120] P. Girichidis et al., *ApJL*, 816, p. L19, 2016.

[121] C. Pfrommer et al., *MNRAS*, 465, pp. 4500–4529, 2017.

[122] I. S. Butsky et al., *ApJ*, 903, p. 77, 2020.

[123] M. Brüggen and E. Scannapieco, *ApJ*, 905, p. 19, 2020.

[124] I. S. Butsky et al., *ApJ*, 935, p. 69, 2022.

[125] K. Kotera and A. V. Olinto, *ARAA*, 49, pp. 119–153, 2011.

[126] P. Girichidis et al., *MNRAS*, 527, pp. 10897–10920, 2024.

[127] R. Kulsrud and W. P. Pearce, *ApJ*, 156, p. 445, 1969.

[128] J. Wiener, C. Pfrommer, and S. P. Oh, *MNRAS*, 467, pp. 906–921, 2017.

[129] T. Buck et al., *MNRAS*, 497, pp. 1712–1737, 2020.

[130] M. Jubelgas et al., *A&A*, 481, pp. 33–63, 2008.

[131] L. Armillotta et al., *ApJ*, 964, p. 99, 2024.

[132] M. Gronke et al., *ApJL*, 862, p. L7, 2018.

[133] S. Ji et al., *MNRAS*, 496, pp. 4221–4238, 2020.

[134] R. Ramesh, D. Nelson, and P. Girichidis, *arXiv e-prints*, arXiv:2409.18238, 2024.

[135] M. J. Rees and J. P. Ostriker, *MNRAS*, 179, pp. 541–559, 1977.

[136] S. D. M. White and M. J. Rees, *MNRAS*, 183, pp. 341–358, 1978.

[137] A. Dekel and Y. Birnboim, *MNRAS*, 368, pp. 2–20, 2006.

[138] Y. Birnboim and A. Dekel, *MNRAS*, 345, pp. 349–364, 2003.

[139] D. Kereš and L. Hernquist, *ApJL*, 700, pp. L1–L5, 2009.

[140] A. Dekel et al., *Nature*, 457, pp. 451–454, 2009.

[141] R. Ramesh et al., *arXiv e-prints*, arXiv:2407.00172, 2024.

[142] S. Chandrasekhar, *ApJ*, 97, p. 255, 1943.

[143] M. Boylan-Kolchin, C.-P. Ma, and E. Quataert, *MNRAS*, 383, pp. 93–101, 2008.

[144] A. R. Wetzel and M. White, *MNRAS*, 403, pp. 1072–1088, 2010.

[145] G. Kauffmann, C. Li, and T. M. Heckman, *MNRAS*, 409, pp. 491–499, 2010.

[146] J. C. Mihos and L. Hernquist, *ApJL*, 425, p. L13, 1994.

[147] P. J. Quinn, L. Hernquist, and D. P. Fullagar, *ApJ*, 403, p. 74, 1993.

[148] L. Hernquist, *ApJ*, 409, p. 548, 1993.

[149] J. Binney and S. Tremaine. *Galactic dynamics*. 1987.

[150] J. E. Gunn and J. R. Gott III, *ApJ*, 176, p. 1, 1972.

[151] C. A. Olano, *A&A*, 485, pp. 457–473, 2008.

[152] E. O'Sullivan et al., *A&A*, 618, A126, 2018.

[153] A. Dekel and Y. Birnboim, *MNRAS*, 383, pp. 119–138, 2008.

[154] M. Gronke and S. P. Oh, *MNRAS*, 480, pp. L111–L115, 2018.

[155] Helmholtz, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 36, pp. 337–346, 1868.

[156] J. A. Fejer, *Physics of Fluids*, 7, pp. 499–503, 1964.

[157] D. H. Sharp, *Physica D Nonlinear Phenomena*, 12, 3,IN1,11–10,IN10,18, 1984.

[158] E. Scannapieco and M. Brüggen, *ApJ*, 805, p. 158, 2015.

[159] D. B. Fielding et al., *ApJL*, 894, p. L24, 2020.

[160] B. Tan and D. B. Fielding, *MNRAS*, 527, pp. 9683–9714, 2024.

[161] J. Barnes and P. Hut, *Nature*, 324, pp. 446–449, 1986.

[162] L. Greengard and V. Rokhlin, *Journal of Computational Physics*, 73, pp. 325–348, 1987.

[163] L. Hernquist, *ApJS*, 64, p. 715, 1987.

[164] A. A. Klypin and S. F. Shandarin, *MNRAS*, 204, pp. 891–907, 1983.

[165] R. Teyssier, *A&A*, 385, pp. 337–364, 2002.

[166] J. S. Bagla, *Journal of Astrophysics and Astronomy*, 23, pp. 185–196, 2002.

[167] K. G. Powell et al., *Journal of Computational Physics*, 154, pp. 284–309, 1999.

[168] J. J. Monaghan and J. C. Lattanzio, *A&A*, 149, pp. 135–143, 1985.

[169] V. Springel and L. Hernquist, *MNRAS*, 333, pp. 649–664, 2002.

[170] R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002.

[171] P. F. Hopkins, *MNRAS*, 450, pp. 53–110, 2015.

[172] V. Springel, *MNRAS*, 401, pp. 791–851, 2010.

[173] A. Pillepich et al., *MNRAS*, 473, pp. 4077–4106, 2018.

[174] G. Chabrier, *PASP*, 115, pp. 763–795, 2003.

[175] M. Schmidt, *ApJ*, 129, p. 243, 1959.

[176] M. Vogelsberger et al., *MNRAS*, 436, pp. 3031–3067, 2013.

[177] A. Rahmati et al., *MNRAS*, 430, pp. 2427–2445, 2013.

[178] D. Nelson et al., *MNRAS*, 475, pp. 624–647, 2018.

[179] M. Donnari et al., *MNRAS*, 500, pp. 4004–4024, 2021.

[180] S. Genel et al., *MNRAS*, 474, pp. 3976–3996, 2018.

[181] D. N. Spergel et al., *ApJS*, 148, pp. 175–194, 2003.

[182] Y. B. Zel'dovich, *A&A*, 5, pp. 84–89, 1970.

[183] D. Munshi, V. Sahni, and A. A. Starobinsky, *ApJ*, 436, p. 517, 1994.

[184] R. E. Angulo et al., *MNRAS*, 426, pp. 2046–2062, 2012.

[185] O. Hahn and T. Abel, *MNRAS*, 415, pp. 2101–2121, 2011.

[186] A. Pillepich et al., *MNRAS*, 490, pp. 3196–3233, 2019.

[187] J. Schaye et al., *MNRAS*, 446, pp. 521–554, 2015.

[188] R. Davé et al., *MNRAS*, 486, pp. 2827–2849, 2019.

[189] P. F. Hopkins et al., *MNRAS*, 492, pp. 3465–3498, 2020.

[190] B. M. Gaensler et al. “Survey Science with ASKAP: Polarization Sky Survey of the Universe’s Magnetism (POSSUM)”. In: *American Astronomical Society Meeting Abstracts #215*. Vol. 215. American Astronomical Society Meeting Abstracts. Jan. 2010, 470.13, p. 470.13.

[191] G. Heald et al., *Galaxies*, 8, p. 53, 2020.

[192] G. M. Voit et al., *ApJ*, 845, p. 80, 2017.

[193] G. M. Voit, *ApJ*, 740, p. 28, 2011.

[194] W. Vieser and G. Hensler, *A&A*, 472, pp. 141–153, 2007.