

First Results from SMAUG: Insights into star formation conditions from spatially-resolved ISM properties in TNG50

BHAWNA MOTWANI,^{1,2} SHY GENEL,^{2,3} GREG L. BRYAN,^{1,2} CHANG-GOO KIM,⁴ VIRAJ PANDYA,^{5,2} RACHEL S. SOMERVILLE,^{2,6} MATTHEW C. SMITH,⁷ EVE C. OSTRICKER,⁴ DYLAN NELSON,⁸ ANNALISA PILLEPICH,⁹ JOHN C. FORBES,² RÜDIGER PAKMOR,¹⁰ AND LARS HERNQUIST⁷

¹*Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA*

²*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA*

³*Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA*

⁴*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*

⁵*UCO/Lick Observatory, Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA*

⁶*Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA*

⁷*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA*

⁸*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany*

⁹*Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany*

¹⁰*Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, 69118 Heidelberg, Germany*

(Received; Revised; Accepted)

Submitted to ApJ

ABSTRACT

Physical properties of the interstellar medium (ISM) at sub-galactic ($\sim$ kpc) scales play an indispensable role in controlling the ability of gas to form stars. As part of the SMAUG (Simulating Multiscale Astrophysics to Understand Galaxies) project, in this paper, we use the TNG50 cosmological simulation of the IllustrisTNG project to explore the physical parameter space comprised of 8 resolved ISM properties in star-forming regions across a wide range of host galaxy mass ($M_{\star} = 10^{7-11} M_{\odot}$) and redshift ($0 \leq z \leq 3$). We deconstruct our simulated galaxies into regions of 1 kpc in size, and measure the gas/stellar surface densities, gas metallicity, vertical stellar velocity dispersion, epicyclic frequency and dark-matter volumetric density representative of each region, and study them in the context of their star formation activity and galactic environment (radial galactocentric location). By examining the star formation rate weighted distributions of these properties, we show that stars primarily form in two distinct regimes, which are brought about by an underlying bi-component radial star formation rate profile in galaxies. We examine how the relative prominence of these two regimes depends on host galaxy mass and cosmic time, and compare our findings with observations from the MaNGA IFU survey. Further, using principal component analysis, we characterise the aforementioned parameter space to reveal a high-degree of multicollinearity in relationships between ISM properties that drive the distribution of star formation at kiloparsec scales. Based on this, we find that a reduced three dimensional representation underpinned in essence by a multi-variate radius relationship is sufficient to capture most of the variance in the original 8D space.

Keywords: Star forming regions (1565), Star formation (1569), Interstellar medium (847), Dimensionality reduction (1943), Galaxy physics (612), Galactic and extragalactic astronomy (563)

1. INTRODUCTION

Local characteristics of the ISM drive a complex interplay of sub-galactic-scale processes which regulate the star formation and feedback in a galaxy. Gravitational instabilities on kpc-scales heavily influence the lifecycle and properties of

giant molecular clouds, thereby setting up star formation and controlling its efficiency in different regions of the galaxy (Semenov et al. 2019; Kruijssen et al. 2019; Bigiel et al. 2008). Feedback on these scales influences the dynamical state of the gas by limiting its evolution toward high densities and dispersing the clouds (e.g., Chevance et al. 2020; Rahner et al. 2019; Semenov et al. 2018, 2017; Kim et al. 2013; Ostriker et al. 2010). Additionally, hydrodynamical interaction between the hot and cold ISM on kpc-scales facilitates the acceleration of galactic outflows that subsequently

suppress star formation through the depletion of cold gas, and prevention of future gas cooling and accretion (Kim & Ostriker 2018a; Rodríguez-Puebla et al. 2016; Forbes et al. 2014a). As such, by virtue of controlling the incidence and effects of star formation, physical properties of the interstellar medium on kpc-scales play a vital role in modulating the overall baryon cycle of galaxies.

Deciphering the link between star formation and galactic structure is a multi-scale problem, ranging from ~ 100 pc scale of molecular cloud collapse to the 10^{5-6} pc scale of the circumgalactic medium. Due to the steep computational challenge of simulating the vast range of scales involved, modern simulations of galaxy formation must implement only approximate *sub-grid* treatments of small-scale physical processes, smoothing over much of the complexity at or below cloud-scales (Naab & Ostriker 2017). Consequently, while results from contemporary large-scale cosmological simulations have been shown to match the integrated stellar mass abundances and star formation rates (SFRs) of observed galaxies (see Somerville & Davé 2015, for a review), they invoke simplified treatments of the underlying, often less-understood, small-scale processes such as star formation and ISM physics, making the treatment less realistic. Understanding the conditions that govern the onset of star formation, and their interrelationships, are therefore, of fundamental importance to the modern theory of galaxy formation and evolution (Morselli et al. 2019; Trayford & Schaye 2019; Chruslinska & Nelemans 2019; Orr et al. 2018). Where and when stars form is a crucial ingredient that must be firmly constrained in order to improve our grasp on the phenomena of star formation and feedback, and their impact on the fate of the baryons over cosmic time.

Ameliorating some of the deficiencies of cosmological simulations, such as in spatial resolution, intermediate-scale simulations (with horizontal box sizes of $\sim \text{kpc}^2$) capable of resolving supernova feedback (for e.g., the TIGRESS suite by Kim & Ostriker 2018b; see also Joung et al. 2009; Hill et al. 2012; Gatto et al. 2017; Kannan et al. 2020) are now helping to close the gap between stellar ($\lesssim \text{pc}$) and cosmological scales ($\gtrsim \text{Mpc}$). In general, these are vertically stratified “tall-box” simulations representative of specific local star-forming regions within a galaxy with a domain size of usually a few kpc and $\sim \text{pc}$ -resolution. In such simulations, the adopted models allow for a more comprehensive and self-consistent evolution of a self-gravitating multiphase ISM with an explicit treatment of star formation and supernova feedback and very few *a-priori* assumptions. The ISM content and disc gravity within such a framework are parameterized by a set of physical properties (e.g., gas/stellar surface density, gas metallicity, stellar scale height/vertical velocity dispersion) which are representative of the patch of the disc galaxy being simulated, and play a central role in governing the process of star formation. As an example, while surface densities dictate the availability (or lack) of pressure support for the collapsing gas, metals control gas cooling and provide sites for formation of dust/molecules in the ISM. Hence, in order to gain a thorough insight into how diverse galac-

tic environments influence the formation of stars, and consequently the outflow properties, a systematic exploration of this hyper-parameter-space in these simulations is warranted. Unfortunately, given the size of the parameter space, and the immense computational cost involved in conducting a sweep, this problem does not lend itself well to a brute-force approach and requires a better sampling scheme to pick out the most essential initial conditions.

Properties of the birth sites of stars can differ substantially between regions not only within a galaxy, but also between galaxies of different global phenotypes. Across large parts of a galaxy, star-forming gas can show a range of densities and metallicities, often correlated with the environment and changing with galactocentric radial location (e.g., Heyer & Dame 2015; Koda et al. 2009). Past observations capable of directly probing cloud-scale quantities like velocity dispersions and surface densities (such as those with ALMA and NOEMA) have indeed revealed sizeable deviations amongst the properties of molecular clouds in different sites, namely, the Galactic centre (Kauffmann et al. 2017; Shetty et al. 2012), outer parts of our Galaxy (Miville-Deschênes et al. 2016; Rice et al. 2016), and local star-forming galaxies (Sun et al. 2018; Schinnerer et al. 2013). Along the same lines, merging and starburst galaxies are known to boast higher densities, line-widths and diluted metallicities (Cortijo-Ferrero et al. 2017; Elmegreen et al. 2016; Irwin 1994) than most other environments. Nevertheless, even with these findings, our knowledge of the properties of stellar birth environments at small scales is limited to a handful of physical parameters, in our own and a number of nearby galaxies (Faesi et al. 2018; Druard et al. 2014; Wong et al. 2011; Bolatto et al. 2008).

To address the issues brought up heretofore, we undertake this study - as part of the SMAUG project¹ - to generate and survey a multi-dimensional parameter space of local physical properties of star-forming sites in a large-volume cosmological simulation. Specifically, we use the IllustrisTNG50 simulation to do a coarse-grained ($\sim \text{kpc}$ -scale) exploration of the gas surface density, gas metallicity, stellar surface density, stellar vertical velocity dispersion, epicyclic frequency, and dark-matter volumetric density in a statistically significant sample of galaxies across a wide range of galaxy mass as well as redshift. While these properties may not potentially be exhaustive in describing the conditions of star formation - either because other properties have a less obvious relevance to star formation or because the timescales for their evolution are short compared to the star formation time - we have aimed to examine the most commonly used local quantities (Table 1 summarises the list). We study these physical parameters in the context of ongoing star formation (via star formation surface density) and the galactic location/environment (radial galacto-centric distance) of each site. In doing this

¹ <https://www.simonsfoundation.org/flatiron/center-for-computational-astrophysics/galaxy-formation/smaug/papersplash1>

study, our goal is to uncover the region(s) of the parameter space over which most star-forming environments exist, or in other words, identify the birth-conditions in which most stars in the universe form. Alongside this, our work will enable the recognition of meaningful initial conditions for future local/tall-box simulations, and help devise an optimal strategy for their exploration. Doing so would enable us to understand the link between star formation and outflows directly, which is a key element of the SMAUG project's larger goal. Lastly, the results of our work will also be a useful research tool for pursuing in detail future lines of inquiry including but not limited to scaling laws connecting local properties of the ISM, and the correspondence between global and local properties of galaxies.

This paper is organised as follows: Section 2 provides a brief description of several features of the IllustrisTNG simulations that are most pertinent to this study. Section 3 describes our detailed methodology for generating the ISM physical parameter space. In Section 4, we present a picture of the parameter space in one dimension, focusing mainly on the distributions of all properties, as well as their dependence on galaxy masses and cosmic time. Section 5 portrays a multi-dimensional view of all of the parameters, where we briefly explore resolved scaling relationships amongst properties, and later, describe the results of dimensionality reduction conducted on the hyper-parameter-space. In Section 6, we make a qualitative comparison of our simulation results with resolved observations from the MaNGA IFU survey and report our findings. Finally, we summarise the conclusions of our work in Section 7.

2. THE TNG SIMULATIONS

The IllustrisTNG project is a suite of gravito-magnetohydrodynamic simulations (Marinacci et al. 2018; Nelson et al. 2018; Springel et al. 2018; Naiman et al. 2018; Pillepich et al. 2018a; Nelson et al. 2019a) consisting of three separate cosmological volumes (TNG300, TNG100, and TNG50) run using the moving-mesh code AREPO (Springel et al. 2001) at distinct mass resolutions. The physical model employed is described in Pillepich et al. 2018b, and includes subgrid treatments of star formation (Springel & Hernquist 2003), metal enrichment from stellar evolution (Naiman et al. 2018), ideal magneto-hydrodynamics (Pakmor & Springel 2013), SN-winds and AGN feedback (Weinberger et al. 2017), and has been shown to yield results that agree with observations over a diverse range of galaxy properties. For this work, we primarily use the highest resolution realisation of the smallest volume box in the suite, i.e. TNG50-1. Here, we present a brief summary of the key features of TNG50-1 (hereafter identified as TNG50; Pillepich et al. 2019; Nelson et al. 2019b), and refer the interested reader to the papers mentioned in this section for further details on the TNG suite.

TNG50 has a uniformly-sampled domain volume of roughly 50^3 Mpc^3 with 2×2160^3 initial resolution elements, and mass resolution of 8.5×10^4 and 4.5×10^5 solar masses for baryons and dark matter respectively. The co-moving

gravitational softening lengths for dark matter and collisionless star particles is 290 pc, while the gas gravitational softening length is adaptive and set by its cell size, with a floor at 74 pc. Both of these values are considerably smaller than the analysis scale we are interested in (i.e. $\sim \text{kpc}$), and hence ensure ample resolving power. At these values, TNG50 provides an exceptional combination of volume alongside resolution that allows us to meaningfully investigate spatially-resolved star formation in this study.

2.1. Star Formation in TNG50

The process of star formation from dense gas in all TNG simulations is governed by an updated Springel & Hernquist (2003) (hereafter SH03) sub-grid model of star formation, which uses a specified density threshold as a criterion for star formation to set in. Gas less dense than the threshold value $n_{\text{th}} = 0.13 \text{ cm}^{-3}$ (in physical units) is considered non star-forming, and its behavior is driven purely by hydrodynamics (in addition to gravity) based on an ideal-gas equation of state. Whereas, above this density value, the model treats the inter-stellar medium as an admixture of two phases of gas i) a cold, dense star-forming cloud phase, and ii) a hot, ionised phase. Of these, the star-forming phase stochastically turns into stars on a timescale set by the local cold gas density. This conversion of gas into stars is then regulated by the heating of the ISM due to supernova feedback (assumed to be instantaneous in this case) resulting from the death of a fraction of the formed stars. Finally, in this self-regulated regime, the model describes the bulk properties of this multiphase high-density gas, such as its pressure and temperature, in terms of an effective equation of state as a function of density.

The key way in which the model incorporated in the TNG simulations differs from the original SH03 model is in its use of a different initial mass function (Chabrier 2003 instead of Salpeter 1955) as well as a softer equation of state. Additionally, on account of the numerical challenge of forming stars due to the very high resolution and consequently extremely short MHD timesteps in the TNG50 simulation, the model implemented therein employs a steeper relationship between the equilibrium star formation timescale and gas density ($t_{\star} \propto n^{-1}$) for the densest gas in the simulation ($\gtrsim 230 n_{\text{th}}$) as opposed to the canonical scaling ($\propto n^{-1/2}$) set by the observed Kennicutt-Schmidt relation (Schmidt 1959; Kennicutt 1998) used in other TNG boxes. This change was made for numerical efficiency reasons alone, and was found to have no significant impact on the overall gas content or galaxy properties in the simulation (for more details, see Nelson et al. 2019b).

Utilizing the TNG simulations, a number of studies seeking to understand star formation and its implications have reported promising outcomes. E.g., Tacchella et al. (2019) investigate the connection between galaxy morphology and star formation, finding that the morphology of galaxies exhibits only a weak correlation with their star formation activity, which matches the observed correlation. Additionally, Donnari et al. (2019) have characterised the star formation activity of simulated galaxies, and shown that the slope and

Table 1. Constituents of the resolved ISM parameter space measured in this work, their physical importance, and their units of measurement.

Property	Physical Role (in tall-box models)	Units
Gas surface density, Σ_{gas}	Self-gravity	$M_{\odot} \text{ kpc}^{-2}$
Stellar surface density, $\Sigma_{\star}$	External gravity	$M_{\odot} \text{ kpc}^{-2}$
Dark matter volumetric density, ρ_{DM}	External gravity	$M_{\odot} \text{ kpc}^{-3}$
Stellar vertical velocity dispersion, $\sigma_{\star}$	External gravity	km s^{-1}
Epicyclic frequency, κ	Gravitational shear	$\text{km s}^{-1} \text{ kpc}^{-1}$
Gas metallicity, Z_{gas}	Gas cooling + heating	dimensionless
Star formation rate surface density, Σ_{sfr}	–	$M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}$
Galactocentric radius, R	–	kpc

NOTE—All spatial quantities here are expressed in physical units (not co-moving), and property roles are defined in the context of their contribution to the tall-box simulation physics.

normalisation of the star-forming main sequence in TNG are in excellent agreement with observations at $z = 0$. Using an improved framework to estimate neutral gas abundances, [Diemer et al. \(2019\)](#) demonstrated a good agreement between the simulated and observed galaxy HI size-mass relationship as well as the overall gas fractions. And most recently, [Nelson et al. \(2019b\)](#) have described how the subgrid input parameters in TNG successfully give rise to a realistic multi-phase structure and diverse properties of feedback-driven galactic outflows. These studies, combined with the multitude of results on other aspects of galaxy formation, provide a solid empirical validation to the TNG model.

3. METHODS AND ANALYSIS

3.1. Galaxy Sample Selection

The large volume of TNG50 provides a statistically sizeable sample of galaxies at a resolution capable of discerning the internal structure of galaxies. Haloes and their substructure are identified in the simulation using the standard Friends of Friends (FoF; [Turner & Gott 1976](#)) and SUBFIND ([Springel et al. 2001](#); [Dolag et al. 2009](#)) algorithms. The FoF algorithm identifies collective groups of dark matter particles (aka halos) based on their physical proximity, whereafter the SUBFIND algorithm identifies gravitationally self-bound associations of all resolution elements combined within each halo (aka subhalos). The gas, stars, black holes and dark matter associated with the most massive subhalo within a FoF halo are considered as belonging to the *central* galaxy, while the rest of the self-gravitating substructures, when present, are classified as *satellite* galaxies. At the present epoch, the simulation contains a total of $\sim 96,000$ galaxies with $M_{\star} \gtrsim 10^5 M_{\odot}$ with a variety of morphologies, sizes and formation histories (e.g., [Pillepich et al. 2019](#); [Joshi et al. 2020](#); [Pulsoni et al. 2020](#)). For our specific analysis, we construct a sample consisting of both centrals and satellites as follows:

1. We select galaxies with stellar masses $M_{\star}$ in the range $10^{7-11} M_{\odot}$ at redshifts $z = \{0, 0.5, 1, 2, 3\}$. The lower

limit of $10^7 M_{\odot}$ is chosen to ensure measurement robustness in view of the mass resolution of the simulation. In this mass range, we find a range of morphologies, from the extended, actively star-forming discs to quenched (non star-forming) elliptical systems, similar to those observed in the real universe.

2. For the selected galaxies, we implement a minimum cut on the number of particles contained within the galaxy to be at 100 each for the gas, stars and dark matter, such that all three components of the galaxies are sufficiently resolved. At the standard TNG50 (i.e. TNG50-1) resolution, this corresponds to a minimum $M_{\star} \simeq 10^{6.9} M_{\odot}$ and $M_{\text{dm}} \simeq 10^{7.7} M_{\odot}$. We also impose a minimum star formation rate threshold of $5 \times 10^{-4} M_{\odot} \text{ yr}^{-1}$ inside twice the 3D stellar half-mass radius ($R_{1/2}^{\star}$) of each galaxy. Even though the latter measure leads to an exclusion of galaxies that are fully or almost-fully quenched, it does not adversely impact our study given its largely star formation driven context.
3. Lastly, due to our inclusion of satellite galaxies in the sample, we take precaution to remove any misidentified subhalos such as clumps or fragments in the outer parts of a halo that did not arise from standard cosmological processes of structure-formation and collapse, using the SubhaloFlag in the simulation (see [Nelson et al. 2019a](#) for more details).

Finally, our sample contains 10394, 13806, 16663, 21039 and 20630 galaxies respectively at $z = \{0, 0.5, 1, 2, 3\}$. In totality, our selected sample at each of the redshifts encompasses $>80\%$ of the total instantaneous star formation occurring in the simulation volume. For most of the analysis in this paper, we will focus on $z = 0$ galaxies with stellar mass in the range $10^9-10^{10} M_{\odot}$, but also explore the variation in resulting trends with galaxy stellar mass and redshift.

3.2. Galaxy Data Processing

Since our goal is to characterise the birth-environments of stars in the context of their corresponding ISM properties, we set up a multi-dimensional physical parameter space by conducting a coarse-grained measurement of these properties from individual star-forming regions within our galaxy sample. In a nutshell, we divide each galaxy into spatially-resolved regions by projecting it onto a 2-dimensional image grid (in the manner of [Diemer et al. 2018](#)), where each pixel represents an ISM patch, and subsequently extract the physical parameters of interest as column-integrated quantities from the image pixels so produced. We further elaborate on the principal facets of our analysis below.

Particle Smoothing: Due to the discretised nature of the simulation volume and a finite mass resolution, each cell/particle in it represents an unresolved entity that should ideally be spatially distributed. For investigating spatially resolved properties, as we do in this paper, it is thus important that we alleviate the effects of this coarse sampling so as to avoid biasing our quantitative analysis or making it dependent on the choice of the analysis scale adopted.

In this work, we utilise a smoothing scheme where we re-sample each star, gas and dark matter particle such that it gives rise to a finitely extended distribution of *sub-particles*. The smoothing length σ , which governs the spatial extent within which the sub-particles originating from a parent particle are distributed, is determined adaptively for each particle based on the local density of the corresponding particle type. This translates into the radius encompassing a fixed number of nearest neighbours N_{ngb} , which we choose to be 32 and 64 for stars and dark matter respectively. For gas, we use one-third of the distance to the 32nd nearest neighbour. We note that our choice of N_{ngb} here is determined based on a visual conciliation between noticeable pixelation effect and fading of resolved structure in galaxies. Using this smoothing length as the standard deviation, we then convolve each particle with a discrete 2D Gaussian kernel of size 6σ on a side and resolution 1 kpc centred on its position, giving us a collection of sub-particles that inherit the physical properties of their parent particle in a manner that conserves the extensive properties (mass, energy, momentum) of the parent particle. From here on, we use these re-sampled sub-particles in lieu of the original particles for further analysis.

Galaxy rotation and coarse-graining: For each galaxy, we first compile the list of both the parent (original) gas/star/dark-matter and corresponding sub-particles comprised within it. We then calculate the angular momentum of the galaxy in its centre of mass rest frame, taken here to be that of all the star-forming *parent* gas particles within $2 \times R_{1/2}^*$, and use it to perform a rotation on the galaxy such that the cartesian z-axis is aligned with the direction of the calculated angular momentum vector. This operation transforms the galaxies with(without) rotational symmetry to a face-on(random) orientation, which is then spatially binned

using a grid of predetermined size and resolution along the x-y plane to create an image representation for each of the desired physical quantities. In this study, we choose a fixed pixel scale (grid resolution) of 1 kpc, which closely emulates the domain size of local ISM simulations, as well as the sampling scale of modern IFU surveys. The overall size of the square image is given by $L_{\text{image}} = 2 \times \max(2R_{1/2}^*, R_{95,\text{sfr}})$, where $R_{95,\text{sfr}}$ is the radius within which 95% of all star-forming parent gas particles of the galaxy are enclosed. This criterion allows us to include most of the star formation in each galaxy while also accounting for a significant fraction of its visible stellar component. A discussion on convergence with different values of pixel scale and simulation resolution is presented in Appendix A.

3.3. Generation of Property Maps

To obtain the desired physical property maps for a galaxy, we utilise the correspondingly binned gas, star and dark matter sub-particles gravitationally bound to the subhalo (as determined by SUBFIND) that lie within a vertical column of height $z = \pm 20$ kpc (unless otherwise noted) relative to the projection plane. This value is arbitrarily chosen, and was selected to minimise the contamination from the hot gaseous corona as well as from halo stars, while preserving the contribution from the diffuse stellar component associated with disc galaxies as well as accommodating galaxies that do not have a well-defined rotation axis. Below, we provide the exact definitions used to calculate local ISM properties of individual pixels (also see Table 1 for a summary list):

Gas surface density Σ_{g} . Sum of the masses of all gas (sub-)particles contained within the column divided by the area of the pixel. For non star-forming gas, we only include the fraction of mass present as neutral hydrogen (although this generally includes $\text{HI} + \text{H}_2$, the simulation does not differentiate between the two). For gas cells below the density threshold, the simulation computes this fraction using the atomic network of ([Katz et al. 1996](#)) and a density-based self shielding prescription ([Rahmati et al. 2013](#)) in the presence of a time-dependent uniform ultraviolet background ([Faucher-Giguère et al. 2009](#)).

Stellar surface density $\Sigma_{\star}$. Sum of the masses of all star particles present within the column divided by the pixel area.

Gas metallicity Z_{g} . Mass-weighted mean metallicity of the same gas as used for the calculation of Σ_{g} .

Star formation surface density Σ_{sfr} . Sum of the star formation rate of all gas particles contained within the column divided by the area of the pixel.

Stellar vertical velocity dispersion $\sigma_{\star}$. Mass-weighted standard deviation of the z-velocity of all the star particles present within the column.




Figure 1. Projections of physical properties for an example disc galaxy at $z = 0$ with $M_* = 2.84 \times 10^{10} M_\odot$ in TNG50. The images are ≈ 43 kpc on a side, which is twice the radius encompassing 95% of all star-forming gas particles in this galaxy. The white circles denote twice the stellar half-mass radius at ≈ 8.5 kpc.

Galacto-centric radius R . Measured based on the number of star-forming gas sub-particles (N_{sfg}) inside the pixel. For pixels with $N_{\text{sfg}} \leq 1$, R is assigned to be the Euclidean distance between the galaxy centre and the centre of the pixel, whereas when $N_{\text{sfg}} > 1$, R is the mean euclidean distance between the galaxy centre and the SFR-weighted mean 2D position coordinates of the gas sub-particles in the pixel.

Epicyclic frequency κ . Calculated using the following expression (simplified from Eq. 3.83-84 in [Binney & Tremaine 2008](#)) involving the galactocentric radius of the pixel R and the circular velocity at that location $v_c(R)$:

$$\kappa(R) = \sqrt{2 \left(\frac{v_c(R)}{R} \right)^2 \left(1 + \frac{R}{v_c(R)} \frac{\partial v_c}{\partial R} \right)} \quad (1)$$

Here, $v_c(R)$ is due to the total mass enclosed within a spherical volume of radius R , and defined as $\sqrt{GM(R)/R}$.

Dark-matter volumetric density ρ_{dm} . Sum of the masses of all dark-matter sub-particles within a column of height $z = \pm h_{*,z}$ divided by the volume of the column. Here, $h_{*,z}$ denotes the stellar half-mass height associated with the pixel.

As an example, we show in Figure 1 images of the aforementioned physical properties for a galaxy with high gas-mass fraction. After generating such images for our entire sample of galaxies, we record the values for all pixels obey-

ing $\Sigma_{\text{sfr}} > 0$ (hereafter dubbed ‘star-forming regions’) to obtain the final 8D parameter-space. We note that due to the presence of out-of-equilibrium and merging galaxies in the simulation volume, not all star-forming pixels in our sample are inherited from dynamically stable discs. Nevertheless, we do not exclude such pixels from our analysis as we are interested in exploring all types of star-forming environments in this study. We also, once again, remind the reader that there are potentially additional local properties that may be important in describing star formation, but in choosing the aforementioned quantities, we have attempted to identify those that are commonly discussed in the context of star formation on kpc-scales. It is indeed expected that not all of these quantities would be mutually independent and encode unique information, and we address this aspect in a later section of this paper (see §5.1).

4. PHYSICAL PROPERTIES OF STAR-FORMING REGIONS

One of the primary goals of our study is to understand the dominant regime of star formation in the universe by way of characterising the underlying physical properties of the ISM. A starting point for doing so would be to summarise the statistical characteristics of the properties themselves measured from our large sample of galaxies. Thereafter, one can glean insight into the physical processes driving the shapes of the distribution functions, as well as parse the degeneracies be-




Figure 2. Independently normalised 1D probability distributions for the measured physical properties of all star-forming pixels from the sample of ≈ 1900 galaxies with $M_* = 10^{9-10} M_\odot$ at $z = 0$ in TNG50. Grey curves depict the unweighted distributions for each property using only the pixels/regions with finite amount of star formation in them while black curves depict the corresponding Σ_{sfr} -weighted distributions. For completeness, we also show in blue the unweighted property distributions obtained from the full pixel dataset inclusive of non-star-forming pixels.

tween them, by exploring how the distributions evolve as galaxies evolve in the redshift and stellar mass space.

Thus, in this section, we begin by presenting probability distributions for the measured physical properties weighted by star formation and highlight their salient features in §4.1. We then explore how these distributions change as a function of stellar mass and redshift of the parent system in §4.2 and §4.3. We also present composite distributions of ISM conditions that have given rise to stars throughout cosmic time in §4.4. Lastly, in §4.5, we investigate the underlying origin of the bimodality seen in the distribution functions presented in §4.1.

4.1. Distribution Functions for Low-redshift Galaxies

As noted in our introduction, by virtue of its complex hydrodynamical, radiative, gravitational and magnetic interactions, the ISM can substantially differ in its physical and chemical properties from region to region, galaxy to galaxy. Since these processes also control the rate of star formation, the diversity in these properties in turn reflects on the way star formation is distributed not just spatially, but also across the spectrum of different physical environments in the universe.

In order to discern which regions of the ISM physical parameter space support most of the overall amount of star formation, it is instructive to look at the distributions of these parameters weighted by the star formation rate of the corresponding pixels. Due to the fixed physical size of all pixels, this is equivalent to weighting by Σ_{sfr} . In Figure 2, we show independently-normalised one-dimensional distributions of the properties of all star-forming regions belonging to our fiducial sample (which are galaxies with $M_* = 10^{9-10} M_\odot$ at

$z = 0$). The grey curve in each panel depicts the *unweighted* distribution, and in black we show the same distributions weighted by Σ_{sfr} . We observe that the weighted radius distribution prefers lower values while all other distributions are shifted towards higher values compared to their corresponding unweighted counterparts. This trend is reflective of the intuitive notion that the denser, inner regions of galaxies are more conducive to the formation of stars on account of the gas being dense enough to cool and collapse. More notably, we find that unlike the unweighted distributions, many of the weighted distributions – all except gas surface density and stellar vertical velocity dispersion – exhibit a strong bimodality. This suggests that star formation in our sample of galaxies is neither agnostic to the properties of the ISM nor does it favour a specific range of values, but instead preferentially occurs in two distinct environmental regimes. Later in this section, we investigate the origin of this feature from radial star formation surface density profiles of the overall population (see §4.5).

4.2. Dependence on Galaxy Stellar Mass

In view of the large dynamic range of galaxy masses present in the simulation, we now look at how properties of the ISM in star-forming regions differ between galaxies with different stellar masses. Figure 3 shows the Σ_{sfr} -weighted distributions of ISM properties of regions drawn from present day ($z = 0$) galaxies in four equal-sized bins of galaxy stellar mass. In each panel, the curves become darker with increasing stellar mass from 10^7 to $10^{11} M_\odot$. We find three broad features to be apparent.




Figure 3. Progression in the Σ_{sfr} -weighted probability distributions of all physical quantities with the stellar mass of the parent galaxy at $z = 0$ in TNG50. Since the curves in each panel are independently normalised, the respective areas under the curve here do not convey the relative amount of star formation occurring in different mass galaxies.

First, we notice that the distributions of gas surface density, stellar surface density, and SFR surface density all exhibit a subtle shift towards higher densities and SFR values for higher mass galaxies. While this trend is mostly manifest in the tails (particularly, on the high- $\Sigma_{\text{g}/\star/\text{sfr}}$ end), the peak values of the distributions do not significantly vary amongst galaxies of different masses. Given that we are solely looking at actively star-forming regions of the galaxies, the concurrence in the behaviours of Σ_{g} and Σ_{sfr} is consistent with, and ensues from the fact that the rate of star formation in galaxies is fundamentally governed by the density of gas on sub-galactic scales. In line with this, the two are directly related by construction in the star formation model of IllustrisTNG (§2.1). Additionally, the lack of strong variation in the resolved star formation rate (alongside gas and stellar density) distributions with galaxy mass confirms that star formation, being an inherently small-scale process, is rather impervious to the overall gravitational potential of the galaxy, but is instead strongly influenced by the local gravity set by Σ_{g} and $\Sigma_{\star}$.

Second, the mean stellar vertical velocity dispersion of star-forming regions monotonically increases as a function of galaxy mass, with the distributions themselves progressively broadening. This dependence can be ascribed to the fact that galaxies with a more massive stellar component require larger dispersions to maintain vertical dynamical equilibrium against the deepening of their gravitational potential wells. Our finding in this case is also corroborated by the results by Pillepich et al. 2019, where they show that the median 3D velocity dispersion and scale height of stellar discs of star-forming TNG50 galaxies indeed increase as a function of peak mass regardless of the sample redshift. Similarly, the peak gas metallicity also shows an increasing trend

with stellar mass (Tremonti et al. 2004), albeit with a gradual translation of the distribution as a whole to higher values. Star-forming gas in higher mass galaxies is expected to be on average more enriched than in lower-mass galaxies owing to a longer corresponding history of star formation and deeper potential, and consequentially, greater metal production and retention. In contrast, lower mass galaxies have a higher gas fraction relative to their stellar material, thus making the metal content more dilute compared to their high-mass counterparts. Interestingly, the variation in the shapes of Z_{g} distributions closely follows as those of the corresponding $\Sigma_{\star}$ distributions, indicating that the increase in metallicity is systematically linked to the bias towards higher stellar densities in more massive galaxies, hence explaining the presence of a *local* mass-metallicity relationship (cf. §5.1).

Finally, from the quantities exhibiting bimodally-shaped distributions, we find that at a given redshift (here, $z = 0$), star formation almost exclusively takes place in the high-DM density innermost regions of low stellar mass galaxies, while in the higher mass galaxies, a gradual suppression of this concentrated star formation paves way for relatively more diffuse star formation in lower density regions. These two regimes are roughly equally populated for galaxies with $M_{\star} \simeq 10^9 M_{\odot}$, above and below which star formation is prevalent in separate sets of parameters. This trend appears due in part to galaxies being more extended at larger masses, hence availing more area for star formation to happen at large radii. This size increase effect is reflected in the translation occurring in the peak positions of the R distribution towards larger values for larger masses. However, other factors could also potentially be at play, namely, an increasing prevalence of central AGN-feedback in more massive galaxies (Bongiorno et al. 2016; Kauffmann et al. 2003; Wang & Kauffmann 2008)




Figure 4. Evolution of Σ_{sfr} -weighted probability distributions of all physical quantities with redshift for the $M_* = 10^{10-11} M_\odot$ galaxy sample in TNG50. All curves shown here are independently normalised.

as well as mass-dependent secular transformation processes leading to a decline in the central gas supply that give rise to a quiescent dense centre surrounded by a more extended gas-rich envelope where star formation mainly occurs (Forbes et al. 2014b; Kormendy & Bender 2012).

4.3. Evolution with Redshift

Having looked at how the local property distributions transform with host galaxy stellar mass, we now explore how these properties vary in similar-mass galaxies as a function of cosmic time. Figure 4 shows independently-normalised property distributions for star-forming patches from galaxies at five different epochs from $z = 3$ to present with the host mass fixed in the range $10^{10-11} M_\odot$ at each epoch. The curves get darker towards lower values of redshift. The panels show that the bimodally distributed quantities favour lower density regions at later times compared to regions within similarly massive hosts at higher redshifts, albeit maintaining a similar overall range of values. At fixed mass, galaxies at higher redshifts are more compact potentially giving rise to the aforementioned trend.

Another notable feature is that in the case of stellar velocity dispersion, there occurs a discernible overall shift in the distributions from higher to lower values, while their width remains mostly unaffected. As expected, this trend can be ascribed to the fact that galaxies at low redshifts, especially star-forming, have a higher degree of rotational support and relatively thinner discs.

Lastly, we see that the shape of the metallicity distribution mildly changes from being bimodal for galaxies at high redshifts to unimodal at low redshifts. As evident from the corresponding Σ_* distributions, this pattern can be ascribed to the removal of high-density, high-metallicity gas from the central regions of galaxies. However, the range of local metallicity

values for our galaxy sample does not substantially evolve with redshift. Considering that the total gas fraction of galaxies appreciably varies with redshift at fixed stellar mass (Santini et al. 2014), the constancy in local metallicity distributions demonstrates that the chemical evolution in galaxies is predominantly driven through outflows and less so via gas inflow and stellar evolution (Torrey et al. 2019). The redshift dependence of the integrated mass-metallicity relationship (MZR) must thus result from galaxy populations losing gas while sampling from an underlying unevolving distribution of local gas metallicity.

4.4. Star formation across Cosmic Time

In the preceding subsections, we looked at the full expanse of conditions under which star formation occurs in the universe at fixed time, and explored how these conditions depend on galaxy mass and epoch. It is then natural to ask: what are the distributions of stellar birth conditions that have collectively given rise to all the stars that have ever been formed? To answer this, we now look at resolved ISM property distributions of star-forming regions across a very wide window of cosmic time. In Figure 5, black curves indicate independently-normalised *stellar-mass-weighted* distributions obtained from a composite dataset of pixels from galaxies with $M_* = 10^{7-11} M_\odot$ at 18 different snapshots - with a roughly uniform spacing of 0.1 in scale factor - between $z = 0$ and 10. Each set of property values (corresponding to a pixel) is weighted by the mass of newly-formed stars contributed by the associated star-forming region. Assuming that the distribution of star-forming region properties varies weakly enough with time, we calculate the stellar-mass contribution for each pixel to be its instantaneous star formation rate (same as Σ_{sfr} due to unit pixel size) times the inter-snapshot duration Δt_{sf} . More precisely, for snapshot i ,




Figure 5. Stellar-mass-weighted distributions for the ISM birth conditions in TNG50 for all stars formed during $10 \geq z \geq 0$ in galaxies with stellar mass in the range $10^{7-11} M_\odot$ (black), compared with the corresponding Σ_{sfr} -weighted distributions of star-forming regions properties in galaxies with $M_* = 10^{10-11} M_\odot$ at $z = 1$ (pink; similar to that in Figure 3). All curves are independently-normalised.

$\Delta t_{\text{sf}}^i = (t^{i-1} - t^{i+1})/2$, where t^i is the lookback time associated with snapshot i . For the last snapshot (corresponding to $z = 0$), it is $(t^{i-1} - t^i)/2$.

The property distributions have an overall strong resemblance to the distributions we saw in §4.3 for galaxies with $M_* = 10^{10-11} M_\odot$ at $z = 1$ (also shown in Figure 5 in pink). In all of the panels barring σ_* , the similarity is apparent both in the locations of the peaks as well as their amplitudes signifying a connection between local ISM conditions contributing new stellar mass in the universe to the conditions sustaining star-formation in massive galaxies at $z = 1$. Previously published work on the global star formation histories of galaxies and the evolution of star formation efficiencies have shown that galaxies with halo masses in the range $10^{11.5} - 10^{12.2} M_\odot$ have the highest star formation efficiency at every epoch, and are responsible for making most of the stars in the universe (Behroozi et al. 2013a). From the stellar mass-halo mass relationship (Moster et al. 2010), this is roughly equivalent to the stellar mass range $10^{10} - 10^{11} M_\odot$, in keeping with our result. Moreover, it has been estimated from both observations and simulations that galaxies in this mass range build up $\sim 80-90\%$ of their stellar mass at $z \lesssim 2$ (mostly in their discs; Tacchella et al. 2019; Behroozi et al. 2013b) and have a mass-weighted mean stellar age of ~ 7 Gyrs, corresponding to $z \simeq 1$ (Behroozi et al. 2013b). This again, is well-reflected in our current findings. Nonetheless, due to the influence of stellar mass formed along the entire cosmic star formation history - which includes stars formed at earlier times and in galaxies with lower stellar masses - the overall distributions here are somewhat broader than the ones we saw in the preceding sections for fixed mass and time.

4.5. The Origin of Bimodality

To glean some insight into the source of bimodality in several of the ISM properties we have seen thus far, we now examine the spatial distribution of star formation. Specifically, since the bimodality in the distributions arises from weighting by Σ_{sfr} , we look into how the mutual variation between star formation and ISM properties can generate bimodality in the respective distributions of those properties. In the preceding sections, we saw a concomitant evolution of bimodalities in multiple quantities, hinting at a common underlying reason for its appearance. Based on this, we choose only one of the parameters - the galactocentric radius R - in this section for illustrative purposes, and expect our inferences to apply equivalently to the other bimodally distributed quantities, namely, Σ_* , Z_g , Σ_{sfr} , κ , and ρ_{dm} .

In Figure 6 we show the sample-averaged $\Sigma_{\text{sfr}}(R)$ profile constructed using the measured values of R and Σ_{sfr} of all star-forming pixels from our fiducial sample of galaxies². The profile appears to be comprised of two disparate components, which we demonstrate below to be responsible for the two distinct peaks in the distributions of quantities. To ascertain the exact shape of the profile, we fit it with a two-component analytical model consisting of an inner plus an outer exponential component as:

$$\langle \Sigma_{\text{sfr}}(R) \rangle = \Sigma_{\text{E1}} e^{-R/R_{\text{E1}}} + \Sigma_{\text{E2}} e^{-R/R_{\text{E2}}} \quad (2)$$

² It is important to clarify here that due to the exclusion of pixels with no finite amount of star formation, the profile shown in Figure 6 does not represent a true radially-averaged density profile, in a way that would be constructed using a stacked galaxy sample. Here, unlike the total pixel population in a galaxy, the number of star-forming-only pixels per radial bin does not scale with the bin radius in a linear fashion.

where $\Sigma_{E1}(\Sigma_{E2})$ are the normalisations, and $R_{E1}(R_{E2})$ are the scale lengths associated with the inner(outer) exponentials. To calculate the fit parameters, we utilise MATLAB’s ‘trust-region’ algorithm (a kind of non-linear least squares formulation) with bisquare weights. The data points used for fitting are obtained by binning the sample values at a uniform interval of 0.5 kpc, while our choice of the model itself is motivated by commonly used profiles in the literature to describe azimuthally-averaged radial distributions of HI gas and star formation rates in galaxies. We exclude from our fitting procedure any data points with relative standard error of mean values exceeding 10%. Figure 6 top panel depicts the fitting procedure results for the fiducial galaxy sample. The overall best fit profile is represented as a solid crimson line with the individual components plotted as yellow (inner exponential) and purple (outer exponential) curves.

As our next step, we seek to combine the two-component fit we obtained with the unweighted distribution (or the number distribution) of R corresponding to the pixel sample (cf. $\log R$ panel in Figure 2). For this purpose, we obtain a functional approximation of the unweighted distribution by fitting it to a log-normal distribution of the form

$$f_{uw}(R) \equiv \frac{1}{N_{\text{tot}}} \frac{dN(R)}{dR} = \frac{1}{R\sigma\sqrt{2\pi}} \exp \frac{-(\ln R - \mu)^2}{2\sigma^2} \quad (3)$$

Here, $N(R)$ denotes the number of pixels as a function of radius, N_{tot} is the total number of pixels, and μ and σ are the mean and standard deviation of the distribution respectively.

Equipped with a functional form for both the average radial Σ_{sfr} -profile as well the unweighted distribution, we then proceed to derive the resultant weighted distribution as

$$\begin{aligned} f_w(R) &\equiv \frac{1}{\Sigma_{\text{sfr,tot}}} \frac{d\Sigma_{\text{sfr}}(R)}{dR} \\ &= \frac{1}{\Sigma_{\text{sfr,tot}}} \frac{d\Sigma_{\text{sfr}}}{dN}(R) \frac{dN(R)}{dR} \\ &= \langle \Sigma_{\text{sfr}}(R) \rangle f_{uw}(R) \left(\frac{\Sigma_{\text{sfr,tot}}}{N_{\text{tot}}} \right)^{-1} \end{aligned} \quad (4)$$

where $\Sigma_{\text{sfr,tot}}$ is the sum of Σ_{sfr} values of all pixels in the fiducial dataset.

As shown in Figure 6 bottom panel, the distribution so obtained not only reproduces the double-peaked structure of the original weighted distribution, but also, the two individual modes present can be separately recovered from the convolution of the unweighted distribution with the ‘inner-’ and ‘outer-’ component of the $\Sigma_{\text{sfr}}(R)$ -profile respectively. The multiplicity of modes in the weighted distribution is therefore a natural outcome of the occurrence of multiple exponential scale lengths in the corresponding average $\Sigma_{\text{sfr}}(R)$ -profile, with the scale length values also governing the positions of the modes.

Having discerned the origins of bimodality in the ensemble property distributions, we now examine whether, and




Figure 6. Top: A two-component fit to the average radial star formation rate surface density profile (black markers) generated using the sample of pixels with $\Sigma_{\text{sfr}} > 0$ belonging to TNG50 galaxies in the mass range $10^{9-10} M_{\odot}$ at $z = 0$. The best fit curves are shown in yellow (inner exponential), purple (outer exponential) and crimson (total). Error-bars on the markers indicate relative standard error of mean, and the abscissa limits do not contain data points that were excluded from the fit. Bottom: The two components making up the average star formation rate profile result in a bimodal shape of the galactocentric radius distribution of star-forming regions (cf. panel 7 in Figure 2).

to what degree, the dichotomy in star formation conditions arises from two distinct populations of galaxies contributing exclusively to either peaks. In other words, could the bimodality result from an admixture of galaxies undergoing inside-out quenching producing the ‘outer’ mode with the




Figure 7. Histogram depicting the fractional amount of star formation within individual TNG50 galaxies that falls within the inner mode relative to the outer mode of the Σ_{sfr} -weighted ensemble R distribution of the fiducial stellar mass bin at $z = 0$.

ones with mostly central star formation manifesting as the “inner” mode? Or perhaps a population of small galaxies with steep exponential profiles mixed with a separate population of large galaxies with shallow profiles? To this end, we assimilate the pixels corresponding to each galaxy separately and calculate the relative fraction of the galaxy’s total star formation contained within the two modes of the ensemble distribution. To separate the two modes, we utilise the local minimum between the peaks as the separation radius, which for our fiducial sample turns out to be at $\log R \simeq 0.19$. We use this simplified criterion instead of a Gaussian mixture model (GMM) for peak separation as GMMs cannot be applied to density distributions when sample weights are in consideration. Figure 7 shows the distribution of the fractional amount of integrated star formation rate of each galaxy that is contained within the “inner” mode (generated by the inner exponential component) of the weighted ensemble distribution of Figure 6. As can be gleaned from the figure, the distribution spans the full range from 0 to 1 indicating that galaxies do not in fact fall into two separate classes, viz., star-forming cores and rings, contributing exclusively to either of the two modes of star formation at the ensemble level. While this is true, this does not deliver a clear conclusion about whether multiple components exist within the star-formation rate profiles of individual galaxies.

The existence of a sharp transition in the slope of the ensemble Σ_{sfr} -profile may arise from two plausible scenarios: i) a population of small galaxies with steep profiles contributing to the inner exponential mixed with a population of large galaxies with shallow profiles producing the outer exponential, or ii) the preponderance of galaxies among the entire sample having the presence of two different scale lengths individually. To ascertain this, we investigate the shape of the average star-formation rate profile for our fiducial galaxy




Figure 8. Average radial star formation rate surface density profiles for the TNG50 fiducial galaxy sample separated into three different star-formation-size bins containing roughly equal number of galaxies.

sample separated by their sizes as shown in Figure 8. Galaxies are split into three equally populated bins of *star formation size*, which we define to be the galactocentric radius of the farthest star-forming pixel in each galaxy. The panels confirm that galaxies of different sizes all have the presence of a broken star formation rate profile, albeit with a more pronounced transition (or elbow) in bigger galaxies on account of their outer components having greater scale lengths, and hence, shallower slopes relative to the inner component. Combined with the inference from the previous figure, this finding suggests that galaxies in general provide a finite contribution to the two modes of star-formation that are present in the ensemble distribution. Additionally, the overall skewness of the distribution in Figure 7 towards lower values implies that the vast majority of galaxies belonging to the fiducial sample exhibit a greater amount of star formation in their diffuse outskirts also at the individual galaxy level, which is analogous to and confirms the population-wide trend noted in §4.

5. A MULTI-DIMENSIONAL VIEW OF THE ISM PARAMETER SPACE

5.1. Resolved Galaxy Scaling Relations

Having thus far examined and discussed the statistical nature of our physical parameter space one quantity at a time, we now look into the mutual relationships amongst these spatially-resolved ISM properties in a pairwise fashion, more commonly known as *resolved scaling relationships*, predicted by TNG50.

In Figure 9, we present the joint distribution functions of all possible pairs of physical properties measured for star-forming regions belonging to the galaxies in our fiducial sample. The diagonal panels show one-dimensional Σ_{sfr} -weighted probability density distributions for the property labeled on the corresponding abscissa (same as in Figure 2),




Figure 9. Corner plot depicting scaling relations between properties of kiloparsec-sized star-forming regions in TNG50. Densities are derived using the collective data corresponding to all star-forming pixels from our fiducial sample. The diagonal panels show 1D Σ_{sfr} -weighted histograms for each quantity, while coloured contours represent Σ_{sfr} -weighted cumulative joint density fractions at 10, 30, 50, 70, 90 and 99% from innermost to the outermost.

while each of the off-diagonal panels show the Σ_{sfr} -weighted two-dimensional cumulative density contours (upto 99% represented by the outermost contour) for the property pair indicated by the corresponding axes labels. In the following discussion, our use of the term *linear* is meant to indicate linearity in log-space.

Our results give rise to resolved counterparts to several canonical global galaxy scaling relationships, namely, the mass-metallicity relation ($M_{\star} - Z_{\text{g}}$) (Tremonti et al. 2004), star formation main sequence ($M_{\star} - \Sigma_{\text{sfr}}$) (Noeske et al. 2007), and the Schmidt star-formation law ($\Sigma_{\text{g}} - \Sigma_{\text{sfr}}$)

(Schmidt 1959; Kennicutt 1998). Apart from these, the figure demonstrates a widespread presence of linear or near-linear relationships between multiple other quantities. The existence of some correlations such as those with galactocentric radius is intuitively expected on account of structural and dynamical considerations as most galaxies are known to have well-defined density and metallicity profiles. However, other correlations (for e.g. metallicity vs. stellar velocity dispersion) have seemingly less obvious physical origins and warrant detailed investigation in a separate future study (Torrey et al., in preparation).




Figure 10. The (unweighted) Σ_{sfr} -weighted distributions of principal component scores in descending order of explained variance shown in grey(black). The first PC captures the bimodality present in multiple properties constituting the original parameter space. Due to the log-transform applied before conducting the PCA, the scores are unitless.

We notice that some relationships, specifically the Schmidt law and mutual relations between stellar mass density, dark-matter density, radius and the epicyclic frequency are extremely tight with negligible scatter. In the case of the former, the low scatter indicates that in TNG50, star formation is not only very closely related to the mass of the gas (due to SH03), but is also well-sampled on kpc-scales across the entire range of Σ_{g} . The latter set of relationships simply reflect the empirically known $\Sigma_{\star}$ and ρ_{dm} radial profiles, as well as the definition used for κ in our analysis (§2). Contrary to this, relationships involving gas density (with the exception of the Schmidt law) exhibit little to no correlation. These characteristics also come up in our subsequent analysis in §5.2.

In a similar vein of contrasting features, we find that while a majority of the scaling relationships have a monotonic and manifestly linear shape, others (most notably, all panels representing stellar velocity dispersion) hint at a break or a turnover. Finally, in many of the two-dimensional distributions, the underlying bimodality in the physical quantities is manifested as two distinct clouds, which in some cases deviate from one another in terms of their slope, thereby giving rise to the broken profiles.

The ubiquity of linear correlations in Figure 9 points to a high degree of redundancy in the overall parameter space, where multiple parameters encode shared information amongst them and lessen the effective degrees of freedom or “axes of variance” available. This information, in principle, should therefore be accessible using a lower-dimensional representation of the same space. Motivated by this feature, we subsequently embark on a search for a reduced representation of the ISM hyper-parameter space using the commonly used technique of principal component analysis for dimensionality-reduction.

5.2. Characterising the Hyperspace of ISM Properties

Even though each star-forming region in our analysis is represented by a set of multiple physical parameters, not all of them are expected to be equally informative. Sometimes, relationships between parameters exist (as evident in §5.1) thereby lowering the degrees of freedom needed to account for the information contained in the original space, also known as the intrinsic dimensionality of the dataset. In order to better scrutinise what these relationships are, and their implications for the conditions in which star formation takes place, we now conduct a statistical characterisation of our measured 8D parameter space by means of lowering its dimensionality using the technique of principal component analysis (PCA). Through this exercise, we seek to answer the question: Is there a simplified meaningful representation of the underlying distribution of star formation in the ISM, and if yes, how many controlling parameters are required for such a representation to work?

5.2.1. Principal Component Analysis

Principal component analysis is a widely-used non-parametric analysis technique to reveal low-dimensional representations of structures underlying complex high-dimensional datasets. It does so by quantifying the importance of each original dimension for describing the information content (or variance) contained within the data. Mathematically, PCA aims to re-express a given dataset with a new set of orthogonal basis - constructed through a *linear combination* of its original basis - that minimises noise and whose axes (known as principal components or PCs) preserve most of the variance within the dataset. This essentially amounts to the diagonalisation of the covariance matrix of the dataset (which carries information about the re-

Table 2. Results of the weighted principal component analysis on the full dataset corresponding to the fiducial galaxy population. Each row describes the linear coefficients (loadings) associated with the 8 physical parameters (features) that make-up the corresponding principal component.

Component	Σ_g	Σ_*	Z_g	Σ_{sfr}	σ_*	κ	R	ρ_{dm}	% variance explained
PC ₀	0.3147	0.3867	0.3437	0.3526	0.2968	0.3803	-0.3769	0.3664	78.83
PC ₁	0.7034	-0.1131	-0.1561	0.5209	-0.1869	-0.1757	0.1961	-0.3041	9.13
PC ₂	0.0362	-0.0771	-0.3897	0.0114	0.8965	-0.0920	0.1558	-0.0656	6.29
PC ₃	0.0475	-0.1157	-0.7803	0.0088	-0.2322	0.2838	-0.3937	0.2933	3.20
PC ₄	0.1226	-0.0775	0.0169	0.0011	-0.0414	-0.5489	0.2223	0.7914	1.20
PC ₅	-0.0707	0.8617	-0.3048	-0.0204	-0.1254	-0.0250	0.3770	-0.0279	0.65
PC ₆	-0.6120	-0.0209	-0.0522	0.7605	-0.0149	-0.1900	-0.0868	-0.0153	0.36
PC ₇	-0.0903	-0.2635	0.0306	0.1594	-0.0412	0.6305	0.6642	0.2359	0.34

NOTE—The 8 x 8 matrix of shaded values constitutes the transpose of the coefficient matrix ($\mathbf{C}^T$; with rows corresponding to PCs and columns corresponding to features), and can be used for the reconstruction of the original parameter space from the principal component values (ref. Appendix B for the exact procedure).

dundancy between parameters and overall noise in the data) using eigenvalue decomposition. The PCs thus obtained are naturally uncorrelated, and are traditionally expressed as a rank-ordered set based on their corresponding variances. Accordingly, the leading component PC₀ is aligned with the direction of largest variance in the data, followed by PC₁, PC₂ and so forth. A lower (say, k) dimensional *hyperplane* approximation to the initial space is then achieved by defining a threshold variance and keeping only the first k PCs required to capture that amount of variance.

A common practice in the application of PCA is to transform and standardise the data to bring all variables on the same footing in terms of their magnitude and scales. Given the vastly different units of measurements and dynamic ranges associated with ISM properties within our dataset, we apply a logarithmic transformation to our entire dataset. In doing so, we alleviate the impact of skewness of the distributions in the linear space, and attain a more practically useful sampling resolution across the dynamic range of all the properties. In addition to this transformation, we also standardise our data such that each original dimension is rescaled to have a mean of zero and unit variance. This is done to avoid variables with larger scales from dominating the covariance structure of the dataset and biasing the directions of the PCs. Taking this data standardisation step then makes the PCA procedure equivalent to diagonalising the correlation matrix instead of the covariance matrix.

In our study, we use the correlation-matrix-based PCA implementation available in MATLAB, which uses a singular-value decomposition algorithm in lieu of the more traditional eigenvalue decomposition to compute the principal component axes. Through the rest of this section, we will refer to the individual star-forming pixels as *samples*, and the corresponding ISM properties as *features*. The dataset derived from our fiducial sample of galaxies and analysed herein con-




Figure 11. Percentage of variance in the data explained by each of the principal components.

stitutes a total of 589,822 samples each associated with an 8-dimensional feature vector.

5.2.2. PCA Results

The principal components we get from our analysis represent a new set of axes (obtained by a generalised rotation of the initial basis) onto which the original variables can be projected. By virtue of being a simple transformation of the original coordinates, the principal component space is expected to capture the most important characteristics (such




Figure 12. Loading plots from the Σ_{sfr} -weighted principal component analysis for pixels corresponding to our fiducial sample. Coloured lines with markers indicate the coefficients of various physical quantities associated with the first and second principal components for the entire sample (top), the outer-mode (bottom-left), and the inner-mode (bottom-right).

as linear correlations) of the feature space. It is therefore instructive to look at the distributions of these transformed coordinates, known as component “scores”, shown here in Figure 10. As evident from the figure, the weighted distribution of the leading component PC_0 has a bimodal shape while the other PCs have distributions unimodal in nature. This indicates that PC_0 by itself is able to capture the distinction between the two star formation regimes exhibited in multiple dimensions in the original feature space. Taking advantage of this fact, we proceed to split our full data into two separate datasets corresponding to the aforementioned regimes (cf. §4) by making a cut along the PC_0 axis, which is deter-

mined by the point of minimum between the two peaks. We conduct PCA separately on these two subsets of the dataset so as to allow us to better understand which physical correlations govern the dispersion within each of the two regimes. Hereafter, the ‘outer-mode’ stands for the low-value peak of PC_0 corresponding to the high-radius low-density star-forming region population, whereas the ‘inner-mode’ represents the low-radius (central) high-density population.

In Figure 11, we show the percentage of the total variance explained by each PC obtained from the decomposition of the entire fiducial dataset, and the two modes separately. For the combined data, the first component (PC_0) alone cap-

tures $\sim 80\%$ of the total variance in the dataset, the second one (PC_1) $\sim 9\%$, the third (PC_2) $\sim 6\%$, and the values drop sharply thereafter with the last few PCs presumably depicting noise. The widths of the distributions in Figure 10 also portray this trend. That the first two PCs collectively account for $\sim 90\%$ of the total variability means that a 2D projection of the feature space could indeed provide a reasonable characterisation of the complete 8D dataset. Furthermore, the bimodal shape of the PC_0 distribution alongside its high value of associated variance suggests that most variability in the data is manifested as peak-to-peak variance arising from the mix of two different sample populations (i.e., bimodality) in the dataset. In fact, this variance far exceeds the sample-to-sample variability within the peaks themselves. On the other hand, in case of the separated modes, the explained variance curve declines rather gradually and the number of PCs needed to capture 90% of the variance goes from 2 up to 5, with the respective first components explaining far less variance (60 and 45%) than their counterpart for the combined data.

Next, we turn towards quantitatively examining the relationship between the feature space and the resulting principal component space. In other words, since PCs are a linear combination of the ISM properties, we can determine the contribution of each of those properties to the PCs, also known as “loading factors” (see Table 2). Figure 12 depicts the eight ISM properties as 2D vectors plotted in the PC_0 (abscissa) vs. PC_1 (ordinate) plane for the full data (upper panel) as well as the two modes separately (lower panels). The (x, y) coordinates of each property are their loading factors along the corresponding PC direction. Several key attributes emerge:

From the top panel, we see that all properties contribute with roughly equal weights to the first principal component and, with the exception of galactocentric radius, have a similarly positive sign. This conveys a correlated equal variation of all quantities with this component to first order, and explains why PC_0 of the full dataset captures the bimodality. Simply put - at low values of PC_0 , samples would have high values of R and low values of all the other quantities thereby belonging to the high- R /low density regime i.e., the outer mode. By the same token, samples that have a high value of PC_0 would belong to the inner mode. A similar trend is seen in the PC_0 loadings associated with the outer mode, albeit with a reduced contribution of Σ_g and Σ_{sfr} . Thus, PC_0 in this case highlights processes that modulate the environment within stellar discs, which is primarily governed by the dynamical influence of stars and dark-matter. In the inner mode, the pattern is significantly different and PC_0 loses most of its correlation with σ_* and ρ_{dm} , hence tracing a more complex dynamical environment driven mainly by gas- and stellar-gravity.

In addition to the composition of PCs, we can draw insights from Figure 12 pertaining to the correlations existing between the features themselves. In the top panel, we observe a clear clustering amongst ISM properties hinting at a high degree of multi-collinearity in the system. In particular, the quantities $\{\Sigma_*, \kappa, Z_g, \sigma_*, \rho_{dm}\}$ lie in a tight cluster indi-




Figure 13. Evolution of the probability density distribution of the first principal component with parent galaxy stellar mass (top) and redshift (bottom). All curves correspond to galaxies at $z = 0$ in the top panel, and the mass bin $10^{10-11} M_{\odot}$ in the bottom panel (consistent with Figures 3 and 4 respectively).

cating that they have a strong positive correlation³ amongst themselves, and a strong negative correlation with the galactocentric radius R . Σ_g and Σ_{sfr} also have a positive correlation between them, which is an expected result based on the canonical Schmidt law. These observations are well in line with our interpretation of Figure 9. In the case of Σ_g , the apparent overall lack of linear correlation with other parameters noted in the previous subsection also emerges in the results of our PCA analysis. This observation is perhaps suggestive of the fact that the variance associated with the Σ_g scaling laws either comes from predominantly non-linear dependencies, or is identified as noise and hence not captured by the high-ranking PCs shown in Figure 11. The figure also affords us the additional clarity that this behaviour is almost entirely

³ the strength of the correlation goes as $\cos(\theta)$, where θ is the angle between the vectors. This means that $\theta = 0(180^\circ)$ depicts perfect linear correlation(anti-correlation).




Figure 14. A reconstruction of the original 8D physical parameter space using the first 3 principal components. See Appendix B for details on the exact reconstruction procedure.

on account of the more dominant *outer* mode. In the case of the outer mode - which we previously saw to be representative of a stellar disc environment - the same correlations stand as in the case of the full dataset. However, in the *inner* mode, most correlations barring the $R - \kappa - Z_g$ and $\Sigma_g - \Sigma_{\text{sfr}}$ relationships are appreciably weakened. Strikingly, in these central star-forming regions, σ_* is no longer linked with the properties of stars/dark-matter but is more tightly coupled to the gas instead. Furthermore, σ_* does not show a strong correlation with the radius, signalling the near-flatness or lack of a well-defined vertical velocity-dispersion profile in those regions.

It is worth keeping in mind here that PCA, being a linear method, is not expected to fully capture the non-linear aspects of relationships between features, if present. Given that, in order to describe a non-linear relationship fully, one cannot rely on a single principal component. Rather, in such a case, a group of PCs is needed, of which one would provide the best linear-approximation to the underlying relationship while the others would encompass variances in the directions of deviations to non-linearity. By using a logarithmic transform, however, we are not strictly confined to the linear regime and are able to additionally capture power-law relationships. Although non-linear generalisations of PCA (such as autoencoder neural networks and kernel-PCA) as well as more advanced manifold-learning methods exist, they are not conducive to the kind of analysis we have conducted in this study i.e., one that requires taking individual sample weights into account.

While so far we stepped through the PCA analysis results for our fiducial sample of galaxies, we now briefly consider how these results change when we conduct PCA on samples representing different galaxy stellar mass ranges and redshifts (akin to our approach in §4.2, 4.3). In Figure 13, we

show the distribution of PC_0 for star-forming region datasets derived from galaxies in different stellar mass bins at $z = 0$, and for datasets corresponding to galaxies with $M_* = 10^{10-11} M_\odot$ at different redshifts. Darker curves represent higher stellar masses and lower redshifts. The variation in the relative amplitude of the two peaks as a function of M_* is reproduced well (compared to those observed in §4.2 and 4.3 for the full space), as is the trend with redshift, where the preference for star formation gradually shifts towards low density, disc-like extended environments for lower redshifts and higher stellar masses. This finding reinforces our understanding from previous results that the leading component fully captures the bimodality signature in the original features as well as the inherent physical relationships that correlate them.

Lastly, in Figure 14, we compare the actual distributions of the initial physical space parameters for our fiducial sample of galaxies with the reconstructed versions generated only by including the first three principal components. We find that the agreement between the original and the low-dimensional version is excellent. For all the physical parameters where there are clearly two distinct peaks, we recover their respective locations and the position of the minimum between them. On the other hand, in terms of the relative peak heights between the two modes, the reconstructed representations are somewhat discrepant from the original distributions. As such, it is possible to reproduce the presence of bimodality in quantity distributions by using PC_0 alone, however, capturing their exact attributes and meaningful variations within individual peaks requires additional components. In our case, the first three PCs account for $\sim 94\%$ of the information present in the dataset (see Table 2). Hence, by retaining a substantial fraction of the information, a three-dimensional embedding of our dataset can serve as a practically useful av-

enue for sampling a family of star-forming regions for further specific investigations, such as tall-box simulations to study feedback and gas dynamics. As a supplementary data product, we provide the PCA loading factors, variances explained by all of the PCs, and joint distributions of the first three PC scores for star-forming regions from galaxies with $M_\star = 10^{7-8}, 10^{8-9}, 10^{9-10}, 10^{10-11} M_\odot$ at $z = \{0, 0.5, 1, 2, 3\}$ at <https://github.com/bhawnamotwani/smaug>.

6. PRELIMINARY COMPARISON WITH OBSERVATIONS

To assess some of the results obtained from our simulations in the context of observations, we now proceed to examine the distributions of properties from resolved observations of nearby star-forming galaxies from the MaNGA integral field unit (IFU) survey data (Bundy et al. 2015). Given the resolution of TNG50 and the scale chosen for our analysis, MaNGA offers a suitable dataset for a qualitative comparison against our results. However, due to the unavailability of an observational counterpart to several of the ‘local’ properties we have worked with, we limit the comparison in this section only to a handful of quantities, namely, $\Sigma_\star$, Σ_{sfr} , and R .

6.1. Survey Description

The MaNGA survey is one of the three programs undertaken as part of the fourth installment of the Sloan Digital Sky Survey (SDSS-IV) aimed at observing resolved kinematic structure of $\sim 10,000$ nearby galaxies through integral field spectroscopy. The imaging and spectroscopy for the galaxies is conducted using the 2.5m telescope at the Apache Point Observatory (APO) alongside specialised IFUs and the BOSS spectrograph with coverage in the 3600-10300Å range and resolving power $R \sim 2000$. In this work, we use the latest public release DR15 (Aguado et al. 2019) consisting of 3D data-cubes for a sample of 4824 unique galaxies uniformly sampled over the stellar mass range $\sim 10^{9-11} M_\odot$.

Raw data is reduced using the MaNGA’s internal data reduction pipeline (Law et al. 2016) and analyzed using the data analysis pipeline (Westfall et al. 2019; Belfiore et al. 2019). The local mass density of spaxels used in this work is computed as part of the Pipe3D pipeline through stellar population modelling using a single age and metallicity population for each spaxel and correcting for dust attenuation (using the Balmer decrement) prior to fitting (for full details, see Sánchez et al. 2016). The local star formation rates are derived from $H\alpha$ luminosity using the formula from Kennicutt (1998) and a Chabrier (2003) initial mass function. Due to an imposed threshold of $S/N = 3$ on $H\beta$ (used for extinction-correction) alongside distance and intrinsic luminosity constraints, the effective median sensitivity limit of Σ_{sfr} in MaNGA lies at $\sim 10^{-3} M_\odot \text{ yr}^{-1} \text{ kpc}^{-2}$. The spatial coverage of galaxies is expected to be at minimum upto 1.5 effective radii (R_e). Lastly, MaNGA observations have a median spatial resolution of $2.5''$ FWHM ($\simeq 1.8 \text{ kpc}$ at the median redshift of ≈ 0.03), and are sampled at a scale of $0.5''$ per spaxel in the final data cubes which translates into a physical scale of $\sim 1\text{-}2 \text{ kpc}$ per spaxel given the redshift range.




Figure 15. The size-mass distribution of TNG50 and MaNGA galaxy samples used for comparative analysis in this section. The size and mass definitions are as described in the axes labels (with the size defined in a 2D face-on projection), and detailed further in the text.

6.2. Galaxy and Spaxel Selection

For our analysis, we choose all galaxies from MaNGA DR15 with a given BPT classification (Kewley et al. 2001; Kauffmann et al. 2003) of either ‘cLIER’ (galaxies with kpc-scale low-ionisation emission regions in their centres accompanied by star-formation in the outskirts) or ‘star-forming’ (as determined by Belfiore et al. 2018). To select a comparison sample for both MaNGA and TNG, we implement a selection criterion that matches galaxies in the size-mass plane. Specifically, we use a metric for the two-dimensional effective radius of the galaxy and the total stellar mass enclosed within the effective radius calculated from the corresponding pixels or spaxels for each galaxy. For MaNGA, we adopt the effective radius to be the inclination-corrected Petrosian half-light radius (R_e), and the face-on projected 2D r -band half-light radius ($R_{1/2,r\text{-band}}$) for TNG50 galaxies⁴. We then stochastically sample both the simulated and observed galaxies so as to match their size-mass distributions in a binned fashion as shown in Figure 15. Following this procedure gives us a total of 147 galaxies in both MaNGA and TNG. Thereafter, for all the selected galaxies, we convolve our images with a MaNGA-like Gaussian PSF with FWHM $\simeq 1.8 \text{ kpc}$, and then draw the corresponding spaxel contributions from within 1.5 times R_e ($R_{1/2,r\text{-band}}$) in MaNGA(TNG) so as to record their $\Sigma_\star$, Σ_{sfr} , and R values. Like the case for TNG50 galaxies, the R values in MaNGA are de-projected to

⁴ The particular choice of mass inside $R_{1/2,r\text{-band}}$ for TNG50 is made in order to avoid uncertainties arising from the differences between the definitions of total mass in the simulation (all mass within the 3D virial radius) and in observation (mass calculated from light within twice the 2D Petrosian radius).




Figure 16. Top: Probability density distributions for Σ_* (left), Σ_{sfr} (middle) and R (right) for all well-defined star-forming spaxels in MaNGA (orange) at $z = 0 - 0.02$ compared with those from TNG50 (black) at $z = 0$, derived in both cases from galaxies with total stellar mass tentatively in the $[10^9, 10^{10}] M_{\odot}$ range. Dotted lines represent the initial unweighted distributions while the solid curves are weighted by Σ_{sfr} . Bottom: Binned average radial star formation rate surface density profiles for MaNGA and TNG50, with the errorbars indicating relative error of mean in each bin. Vertical bars in grey and orange denote the analysis scale used for TNG50 (1 kpc) and the spatial resolution of MaNGA at median redshift (≈ 1.8 kpc) respectively.

be in the face-on orientation. Any bad spaxels such as those with ill-defined stellar mass and undetected star formation rate are excluded from the comparison study. Additionally, to mimic the MaNGA detection thresholds in the simulation data, we limit our analysis to only the subset of all spaxels that obey $\Sigma_* \geq 2 \times 10^7 M_{\odot} \text{ kpc}^{-2}$ and $\Sigma_{\text{sfr}} \geq 10^{-3} M_{\odot} \text{ yr}^{-1} \text{ kpc}^{-2}$.

6.3. Inference and Discussion

In Figure 16, we present the results of our comparative analysis between the observational and theoretical datasets. The top panels illustrate the individually normalised unweighted and Σ_{sfr} -weighted distributions of $\log \Sigma_*$, $\log \Sigma_{\text{sfr}}$, and $\log R$ as dotted and solid curves respectively. As the figure indicates, property distributions in MaNGA cover approximately the same range as TNG50 both in their original unweighted as well as the star formation-weighted forms. As a consequence, we also achieve an overall good agreement between the two in terms of peak height. Interestingly, due to the observational mocking steps involved, the TNG50 distributions no longer show a discernible bimodal shape, keeping in line with their observational counterparts. Specifically, applying the MaNGA point spread function on the simulation results reduces the Σ_{sfr} values of the inner pixels (at $R \lesssim 2$ kpc) by spreading star formation across multiple surrounding pixels. This in turn leads to an increase in the scale length of the otherwise steep inner exponential part of the TNG50

radial Σ_{sfr} -profile and the smoothening of the elbow following it. This change thereby not only causes the suppression of the ‘inner’ mode (c.f. §4.5), but also fades the prominent separation between the two modes by decreasing the difference between the slopes of the inner and outer exponential components of the profile.

Notwithstanding the general conformity, property distributions in TNG50 exhibit a few minor departures from that of MaNGA. The weighted MaNGA distributions of Σ_* and Σ_{sfr} seemingly favour peak values that are slightly greater in comparison with those exhibited by the TNG50 curves, a behaviour that is not perceptible in the case of their unweighted distributions. Another marginal disparity is apparent between the two R -distributions in that the TNG50 curve features a minute overall shift (≈ 0.1 dex) towards higher values both in the unweighted and weighted forms relative to MaNGA. Lastly, the imposed hard Σ_* -cut in the case of TNG50 expectedly manifests as a sharp lower limit of the unweighted distribution against a much softer, tapered edge in MaNGA. Curiously, while the peak of the weighted R distribution in MaNGA approximately lines up with the outer (large-radius) mode in TNG, the peak in Σ_* is more compatible with the inner higher-density mode of the corresponding simulation-derived distribution.

The bottom panel of Figure 16 depicts the average Σ_{sfr} profiles for the two datasets, constructed in a similar fashion to the curve shown in Figure 6. Barring a slight offset

in the normalisations at large radii (≈ 0.1 dex at ≥ 7 kpc), we find that the two profiles exhibit an excellent agreement with one another. In contrast with the 2-component profile composition discussed in §4.5, both the observed MaNGA and ‘mock’ TNG50 profiles have a smooth shape that can be best described by a single continuous component. Indeed, if we fit the two curves, we find that the MaNGA(TNG) profile can be well-represented by a single exponential of scale length $\simeq 1.3(0.7)$ kpc plus a constant. We therefore deem the lack of two disparate scale lengths in this case to be of direct consequence to the shapes of the weighted distributions, explaining the absence of a bimodality. Finally, we notice that in spite of being largely akin to observations, the TNG-profile near the centre (at $R < 1$ kpc) is somewhat shallower and has a value that is a factor of ≈ 2 lower compared to the inferred value in MaNGA. This difference could be partially responsible for the suppressed weighted probability i.e., star formation, at small-radii in TNG50 relative to MaNGA. Although, in light of our chosen simulation sampling scale and the spatial resolution of MaNGA ($\simeq 1$ and 1.8 kpc respectively), we note that the nature of both profiles and their mutual comparison at these small radii may be inadequate to draw any robust conclusions. At large radii, the TNG50 curve has a steeper decline as indicated by the smaller best-fit scale length, while MaNGA retains a roughly constant value of star-formation albeit with a relatively noisier profile due to a significant dearth of star-forming spaxels in the outskirts.

Despite a careful galaxy sample selection and the application of MaNGA-like detection limits on simulation-derived data, the presence of minor dissimilarities between TNG50 and MaNGA suggests that the true extent of the selection effects inherent to the survey is perhaps not fully captured in the hard thresholds that we have used in our analysis. The resulting discrepancies are conspicuously reflected in the (difference between) unweighted distributions of Σ_* and R , as well as the $\Sigma_{\text{sfr}}(R)$ profiles. As an example, the difference between the average Σ_{sfr} -profile slopes at large radii and in central regions of galaxies is symptomatic of that fact that low- Σ_{sfr} regions are preferentially underrepresented in the observational data set that we have utilised. Such regions of low star formation are typically found in galaxy outskirts as well as in the centres of galaxies undergoing inside-out quenching (the latter would be expected due to our inclusion of cLIER galaxies). Considering that the properties of TNG galaxies such as sizes, star formation rates and specific star formation rates have previously been shown to be consistent with observed galaxy properties (e.g., [Stevens et al. 2019](#); [Hwang et al. 2019](#); [Genel et al. 2018](#)), we hypothesize that the dearth of these high-radius/central low-SFR pixels is borne out of vulnerability to the exact nature of detection limits of the survey, particularly in Σ_{sfr} . An exhaustive future study conducted using carefully generated mock-MaNGA observations from TNG will allow us to test this hypothesis. Nonetheless, we hope that with future resolved spectroscopic surveys pushing the detection limits, the trends emerging from our study could be directly tested against ISM

property distributions in resolved observations, and provide the much needed clarity in this picture.

7. SUMMARY AND OUTLOOK

In this work, we use the TNG50 cosmological simulation volume to generate and statistically survey the multi-dimensional parameter space of resolved ISM properties across a wide range of galaxy masses and redshifts. Specifically, we select star-forming galaxies in the mass range $10^7 - 10^{11} M_\odot$ at $z = \{0, 0.5, 1, 2, 3\}$. This sample accounts for more than 80% of the total star formation in the simulation at each of the utilised snapshot. By dividing the galaxy into kpc-sized regions, we conduct a coarse-grained measurement of gas/stellar surface densities, gas metallicity, stellar vertical velocity dispersion, disc epicyclic frequency, star formation rate density and dark-matter volumetric density representative of each region (see §3.2 for further details). We present a synopsis of the main findings of our analyses below.

- The distributions of all ISM properties, with the exception of stellar vertical velocity dispersion and gas surface density, exhibit bimodally-shaped distribution functions when weighted by star formation (Figure 2), indicating that star formation in galaxies takes place in two separate environmental regimes. Star formation is most favoured to occur in the outer low-density low-metallicity regions for high-mass galaxies (above $\gtrsim 10^9 M_\odot$), while being localised to the central high-density regions for lower mass galaxies (Figure 3). For galaxies in a fixed mass bin, the preference for star formation in the outer diffuse regions is greater at lower redshifts (Figure 4). Additionally, our results show that most of the star formation in the universe takes place in galaxies with $M_* = 10^{10-11} M_\odot$ at $z \leq 2$ (Figure 5).
- The presence of a bimodality in property distributions results from an underlying bi-component average radial star formation rate profile for the galaxy sample. By fitting this profile with a combination of two exponentials, we demonstrate that the two peaks in the weighted distributions can be individually reproduced from the “inner” and “outer” exponential components separately (Figure 6). We also find that almost all galaxies sustain a finite amount of star formation in both modes, albeit with varying degrees of relative contributions to them (Figure 7, 8).
- We investigate the 2D joint density distributions between parameters (Figure 9), and find a very high degree of multicollinearity (aka redundancy) in the 8D space. Through linear dimensionality-reduction via principal component analysis, we find that almost all of the intrinsic variance of the parameter space can be well captured via a transformed 3-dimensional representation (Figure 11, 14). Moreover, the leading principal component alone also captures the multi-parameter bimodality signature present in the original space (Figure 10, 13). This signature is manifested in the form of a “radius-relation” (Figure 12), i.e., the anti-

correlation of galactocentric radius with the rest of the bimodally distributed ISM parameters.

- We conduct a preliminary comparison of our 1D property distributions and star formation rate profiles with those obtained from the MaNGA IFU survey for $z = 0$, with both galaxy samples selected to represent similar size-mass distributions (Figure 15, 16). Upon the application of observational detection limits, the Σ_{sfr} -weighted distributions in TNG50 lose their bimodal shape showing concordance with the shape of the observed resolved property distributions. The comparison reveals an overall good match between TNG50 and MaNGA for the spread and peak locations of the parameter distributions, as well as the underlying average radial Σ_{sfr} profiles below $\simeq 7$ kpc. We argue that some of the minor deviations between MaNGA and TNG50 results possibly arise from our inability to fully capture the nuances of the observational detection limits necessary to make a maximally unbiased comparison.

We envision that the results from our study will provide impetus for the construction of new heuristic star formation/ISM prescriptions in the near-future that are driven by fewer free parameters compared to currently used subgrid models. Given that our analysis is based on a dataset acquired from a fully cosmological, high-resolution, large volume simulation, we provide strong constraints for any model that endeavours to physically describe star formation and ISM physics in galaxies. By facilitating an optimal sampling of realistic initial conditions for future high-resolution tall-box ISM simulations (e.g. TIGRESS), our characterisation will allow us to bridge the gap between stellar and galactic

scales by establishing a direct link between small-scale ISM conditions and large-scale outflow properties. Finally, our work will provide avenues for meaningful comparison with similar measurements conducted with other large-scale cosmological simulations, as well as detailed quantitative patterns emerging from future high-precision spatially-resolved observations, especially at high-redshifts.

B.M. is grateful to Francesco Belfiore for providing the MaNGA data catalogue used for the observational comparison presented in this paper, and for useful advice. B.M. also thanks Christopher C. Hayward and Tjitske Starkenburg for thoughtful suggestions that helped improve this manuscript. This work was carried out as part of the SMAUG project. SMAUG gratefully acknowledges support from the Center for Computational Astrophysics at the Flatiron Institute, which is supported by the Simons Foundation. The computational analyses presented in this paper were conducted in part on the Iron Cluster at the Flatiron Institute and the Popeye-Simons System at the San Diego Supercomputing Centre. B.M. thanks the scientific computing staff at the Flatiron Institute for their technical assistance. The work of C.-G.K. was partly supported by a grant from the Simons Foundation (CCA 528307, E.C.O.) and NASA ATP grant No. NNX17AG26G. E.C.O. acknowledges support received under the Simons Investigator grant number 510940. TNG50 was run on the Hazel Hen supercomputer at the High Performance Computing Center Stuttgart using compute time granted by the Gauss Centre for Supercomputing (GCS) under GCS Large-Scale Project GCS-DWAR (CoPIs: Nelson, Pillepich).

REFERENCES

- Aguado, D. S., Ahumada, R., Almeida, A., et al. 2019, *ApJS*, 240, 23, doi: [10.3847/1538-4365/aaf651](https://doi.org/10.3847/1538-4365/aaf651)
- Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013a, *ApJL*, 762, L31, doi: [10.1088/2041-8205/762/2/L31](https://doi.org/10.1088/2041-8205/762/2/L31)
- . 2013b, *ApJ*, 770, 57, doi: [10.1088/0004-637X/770/1/57](https://doi.org/10.1088/0004-637X/770/1/57)
- Belfiore, F., Maiolino, R., Bundy, K., et al. 2018, *MNRAS*, 477, 3014, doi: [10.1093/mnras/sty768](https://doi.org/10.1093/mnras/sty768)
- Belfiore, F., Westfall, K. B., Schaefer, A., et al. 2019, *AJ*, 158, 160, doi: [10.3847/1538-3881/ab3e4e](https://doi.org/10.3847/1538-3881/ab3e4e)
- Bigiel, F., Leroy, A., Walter, F., et al. 2008, *AJ*, 136, 2846, doi: [10.1088/0004-6256/136/6/2846](https://doi.org/10.1088/0004-6256/136/6/2846)
- Binney, J., & Tremaine, S. 2008, *Galactic Dynamics: Second Edition*
- Bolatto, A. D., Leroy, A. K., Rosolowsky, E., Walter, F., & Blitz, L. 2008, *ApJ*, 686, 948, doi: [10.1086/591513](https://doi.org/10.1086/591513)
- Bongiorno, A., Schulze, A., Merloni, A., et al. 2016, *A&A*, 588, A78, doi: [10.1051/0004-6361/201527436](https://doi.org/10.1051/0004-6361/201527436)
- Bundy, K., Bershady, M. A., Law, D. R., et al. 2015, *ApJ*, 798, 7, doi: [10.1088/0004-637X/798/1/7](https://doi.org/10.1088/0004-637X/798/1/7)
- Chabrier, G. 2003, *ApJL*, 586, L133, doi: [10.1086/374879](https://doi.org/10.1086/374879)
- Chevance, M., Kruijssen, J. M. D., Hygate, A. P. S., et al. 2020, *MNRAS*, 493, 2872, doi: [10.1093/mnras/stz3525](https://doi.org/10.1093/mnras/stz3525)
- Chruslinska, M., & Nelemans, G. 2019, *MNRAS*, 488, 5300, doi: [10.1093/mnras/stz2057](https://doi.org/10.1093/mnras/stz2057)
- Cortijo-Ferrero, C., González Delgado, R. M., Pérez, E., et al. 2017, *MNRAS*, 467, 3898, doi: [10.1093/mnras/stx383](https://doi.org/10.1093/mnras/stx383)
- Diemer, B., Stevens, A. R. H., Forbes, J. C., et al. 2018, *ApJS*, 238, 33, doi: [10.3847/1538-4365/aae387](https://doi.org/10.3847/1538-4365/aae387)
- Diemer, B., Stevens, A. R. H., Lagos, C. d. P., et al. 2019, *MNRAS*, 487, 1529, doi: [10.1093/mnras/stz1323](https://doi.org/10.1093/mnras/stz1323)
- Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, *MNRAS*, 399, 497, doi: [10.1111/j.1365-2966.2009.15034.x](https://doi.org/10.1111/j.1365-2966.2009.15034.x)
- Donnari, M., Pillepich, A., Nelson, D., et al. 2019, *MNRAS*, 485, 4817, doi: [10.1093/mnras/stz712](https://doi.org/10.1093/mnras/stz712)
- Druard, C., Braine, J., Schuster, K. F., et al. 2014, *A&A*, 567, A118, doi: [10.1051/0004-6361/201423682](https://doi.org/10.1051/0004-6361/201423682)
- Elmegreen, B. G., Kaufman, M., Bournaud, F., et al. 2016, *ApJ*, 823, 26, doi: [10.3847/0004-637X/823/1/26](https://doi.org/10.3847/0004-637X/823/1/26)
- Faesi, C. M., Lada, C. J., & Forbrich, J. 2018, *ApJ*, 857, 19, doi: [10.3847/1538-4357/aad60](https://doi.org/10.3847/1538-4357/aad60)

- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *ApJ*, 703, 1416, doi: [10.1088/0004-637X/703/2/1416](https://doi.org/10.1088/0004-637X/703/2/1416)
- Forbes, J. C., Krumholz, M. R., Burkert, A., & Dekel, A. 2014a, *MNRAS*, 438, 1552, doi: [10.1093/mnras/stt2294](https://doi.org/10.1093/mnras/stt2294)
- . 2014b, *MNRAS*, 443, 168, doi: [10.1093/mnras/stu1142](https://doi.org/10.1093/mnras/stu1142)
- Gatto, A., Walch, S., Naab, T., et al. 2017, *MNRAS*, 466, 1903, doi: [10.1093/mnras/stw3209](https://doi.org/10.1093/mnras/stw3209)
- Genel, S., Nelson, D., Pillepich, A., et al. 2018, *MNRAS*, 474, 3976, doi: [10.1093/mnras/stx3078](https://doi.org/10.1093/mnras/stx3078)
- Heyer, M., & Dame, T. M. 2015, *ARA&A*, 53, 583, doi: [10.1146/annurev-astro-082214-122324](https://doi.org/10.1146/annurev-astro-082214-122324)
- Hill, A. S., Joung, M. R., Mac Low, M.-M., et al. 2012, *ApJ*, 750, 104, doi: [10.1088/0004-637X/750/2/104](https://doi.org/10.1088/0004-637X/750/2/104)
- Hwang, H. S., Shin, J., & Song, H. 2019, *MNRAS*, 489, 339, doi: [10.1093/mnras/stz2136](https://doi.org/10.1093/mnras/stz2136)
- Irwin, J. A. 1994, *ApJ*, 429, 618, doi: [10.1086/174349](https://doi.org/10.1086/174349)
- Joshi, G. D., Pillepich, A., Nelson, D., et al. 2020, arXiv e-prints, arXiv:2004.01191. <https://arxiv.org/abs/2004.01191>
- Joung, M. R., Mac Low, M.-M., & Bryan, G. L. 2009, *ApJ*, 704, 137, doi: [10.1088/0004-637X/704/1/137](https://doi.org/10.1088/0004-637X/704/1/137)
- Kannan, R., Marinacci, F., Simpson, C. M., Glover, S. C. O., & Hernquist, L. 2020, *MNRAS*, 491, 2088, doi: [10.1093/mnras/stz3078](https://doi.org/10.1093/mnras/stz3078)
- Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *ApJS*, 105, 19, doi: [10.1086/192305](https://doi.org/10.1086/192305)
- Kauffmann, G., Heckman, T. M., Tremonti, C., et al. 2003, *MNRAS*, 346, 1055, doi: [10.1111/j.1365-2966.2003.07154.x](https://doi.org/10.1111/j.1365-2966.2003.07154.x)
- Kauffmann, J., Pillai, T., Zhang, Q., et al. 2017, *A&A*, 603, A90, doi: [10.1051/0004-6361/201628089](https://doi.org/10.1051/0004-6361/201628089)
- Kennicutt, Robert C., J. 1998, *ApJ*, 498, 541, doi: [10.1086/305588](https://doi.org/10.1086/305588)
- Kewley, L. J., Dopita, M. A., Sutherland, R. S., Heisler, C. A., & Trevena, J. 2001, *ApJ*, 556, 121, doi: [10.1086/321545](https://doi.org/10.1086/321545)
- Kim, C.-G., & Ostriker, E. C. 2018a, *ApJ*, 853, 173, doi: [10.3847/1538-4357/aaa5ff](https://doi.org/10.3847/1538-4357/aaa5ff)
- . 2018b, *ApJ*, 853, 173, doi: [10.3847/1538-4357/aaa5ff](https://doi.org/10.3847/1538-4357/aaa5ff)
- Kim, C.-G., Ostriker, E. C., & Kim, W.-T. 2013, *ApJ*, 776, 1, doi: [10.1088/0004-637X/776/1/1](https://doi.org/10.1088/0004-637X/776/1/1)
- Koda, J., Scoville, N., Sawada, T., et al. 2009, *ApJL*, 700, L132, doi: [10.1088/0004-637X/700/2/L132](https://doi.org/10.1088/0004-637X/700/2/L132)
- Kormendy, J., & Bender, R. 2012, *ApJS*, 198, 2, doi: [10.1088/0067-0049/198/1/2](https://doi.org/10.1088/0067-0049/198/1/2)
- Kruijssen, J. M. D., Schrubba, A., Chevance, M., et al. 2019, *Nature*, 569, 519, doi: [10.1038/s41586-019-1194-3](https://doi.org/10.1038/s41586-019-1194-3)
- Law, D. R., Cherinka, B., Yan, R., et al. 2016, *AJ*, 152, 83, doi: [10.3847/0004-6256/152/4/83](https://doi.org/10.3847/0004-6256/152/4/83)
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, 480, 5113, doi: [10.1093/mnras/sty2206](https://doi.org/10.1093/mnras/sty2206)
- Miville-Deschênes, M. A., Duc, P. A., Marleau, F., et al. 2016, *A&A*, 593, A4, doi: [10.1051/0004-6361/201628503](https://doi.org/10.1051/0004-6361/201628503)
- Morselli, L., Popesso, P., Cibinel, A., et al. 2019, *A&A*, 626, A61, doi: [10.1051/0004-6361/201834559](https://doi.org/10.1051/0004-6361/201834559)
- Moster, B. P., Somerville, R. S., Maulbetsch, C., et al. 2010, *ApJ*, 710, 903, doi: [10.1088/0004-637X/710/2/903](https://doi.org/10.1088/0004-637X/710/2/903)
- Naab, T., & Ostriker, J. P. 2017, *ARA&A*, 55, 59, doi: [10.1146/annurev-astro-081913-040019](https://doi.org/10.1146/annurev-astro-081913-040019)
- Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 477, 1206, doi: [10.1093/mnras/sty618](https://doi.org/10.1093/mnras/sty618)
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 475, 624, doi: [10.1093/mnras/stx3040](https://doi.org/10.1093/mnras/stx3040)
- Nelson, D., Springel, V., Pillepich, A., et al. 2019a, *Computational Astrophysics and Cosmology*, 6, 2, doi: [10.1186/s40668-019-0028-x](https://doi.org/10.1186/s40668-019-0028-x)
- Nelson, D., Pillepich, A., Springel, V., et al. 2019b, *MNRAS*, 490, 3234, doi: [10.1093/mnras/stz2306](https://doi.org/10.1093/mnras/stz2306)
- Noeske, K. G., Weiner, B. J., Faber, S. M., et al. 2007, *ApJL*, 660, L43, doi: [10.1086/517926](https://doi.org/10.1086/517926)
- Orr, M. E., Hayward, C. C., Hopkins, P. F., et al. 2018, *MNRAS*, 478, 3653, doi: [10.1093/mnras/sty1241](https://doi.org/10.1093/mnras/sty1241)
- Ostriker, E. C., McKee, C. F., & Leroy, A. K. 2010, *ApJ*, 721, 975, doi: [10.1088/0004-637X/721/2/975](https://doi.org/10.1088/0004-637X/721/2/975)
- Pakmor, R., & Springel, V. 2013, *MNRAS*, 432, 176, doi: [10.1093/mnras/stt428](https://doi.org/10.1093/mnras/stt428)
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018a, *MNRAS*, 475, 648, doi: [10.1093/mnras/stx3112](https://doi.org/10.1093/mnras/stx3112)
- Pillepich, A., Springel, V., Nelson, D., et al. 2018b, *MNRAS*, 473, 4077, doi: [10.1093/mnras/stx2656](https://doi.org/10.1093/mnras/stx2656)
- Pillepich, A., Nelson, D., Springel, V., et al. 2019, *MNRAS*, 490, 3196, doi: [10.1093/mnras/stz2338](https://doi.org/10.1093/mnras/stz2338)
- Pulsoni, C., Gerhard, O., Arnaboldi, M., et al. 2020, arXiv e-prints, arXiv:2004.13042. <https://arxiv.org/abs/2004.13042>
- Rahmati, A., Pawlik, A. H., Raičević, M., & Schaye, J. 2013, *MNRAS*, 430, 2427, doi: [10.1093/mnras/stt066](https://doi.org/10.1093/mnras/stt066)
- Rahner, D., Pellegrini, E. W., Glover, S. C. O., & Klessen, R. S. 2019, *MNRAS*, 483, 2547, doi: [10.1093/mnras/sty3295](https://doi.org/10.1093/mnras/sty3295)
- Rice, T. S., Goodman, A. A., Bergin, E. A., Beaumont, C., & Dame, T. M. 2016, *ApJ*, 822, 52, doi: [10.3847/0004-637X/822/1/52](https://doi.org/10.3847/0004-637X/822/1/52)
- Rodríguez-Puebla, A., Primack, J. R., Behroozi, P., & Faber, S. M. 2016, *MNRAS*, 455, 2592, doi: [10.1093/mnras/stv2513](https://doi.org/10.1093/mnras/stv2513)
- Salpeter, E. E. 1955, *ApJ*, 121, 161, doi: [10.1086/145971](https://doi.org/10.1086/145971)
- Sánchez, S. F., Pérez, E., Sánchez-Blázquez, P., et al. 2016, *RMxAA*, 52, 21. <https://arxiv.org/abs/1509.08552>
- Santini, P., Maiolino, R., Magnelli, B., et al. 2014, *A&A*, 562, A30, doi: [10.1051/0004-6361/201322835](https://doi.org/10.1051/0004-6361/201322835)
- Schinnerer, E., Meidt, S. E., Pety, J., et al. 2013, *ApJ*, 779, 42, doi: [10.1088/0004-637X/779/1/42](https://doi.org/10.1088/0004-637X/779/1/42)
- Schmidt, M. 1959, *ApJ*, 129, 243, doi: [10.1086/146614](https://doi.org/10.1086/146614)
- Semenov, V. A., Kravtsov, A. V., & Gnedin, N. Y. 2017, *ApJ*, 845, 133, doi: [10.3847/1538-4357/aa8096](https://doi.org/10.3847/1538-4357/aa8096)

- , 2018, *ApJ*, 861, 4, doi: [10.3847/1538-4357/aac6eb](https://doi.org/10.3847/1538-4357/aac6eb)
- , 2019, *ApJ*, 870, 79, doi: [10.3847/1538-4357/aaf163](https://doi.org/10.3847/1538-4357/aaf163)
- Shetty, R., Beaumont, C. N., Burton, M. G., Kelly, B. C., & Klessen, R. S. 2012, *MNRAS*, 425, 720, doi: [10.1111/j.1365-2966.2012.21588.x](https://doi.org/10.1111/j.1365-2966.2012.21588.x)
- Somerville, R. S., & Davé, R. 2015, *ARA&A*, 53, 51, doi: [10.1146/annurev-astro-082812-140951](https://doi.org/10.1146/annurev-astro-082812-140951)
- Springel, V., & Hernquist, L. 2003, *MNRAS*, 339, 289, doi: [10.1046/j.1365-8711.2003.06206.x](https://doi.org/10.1046/j.1365-8711.2003.06206.x)
- Springel, V., Yoshida, N., & White, S. D. M. 2001, *NewA*, 6, 79, doi: [10.1016/S1384-1076\(01\)00042-2](https://doi.org/10.1016/S1384-1076(01)00042-2)
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, 475, 676, doi: [10.1093/mnras/stx3304](https://doi.org/10.1093/mnras/stx3304)
- Stevens, A. R. H., Diemer, B., Lagos, C. d. P., et al. 2019, *MNRAS*, 490, 96, doi: [10.1093/mnras/stz2513](https://doi.org/10.1093/mnras/stz2513)
- Sun, J., Leroy, A. K., Schruba, A., et al. 2018, *ApJ*, 860, 172, doi: [10.3847/1538-4357/aac326](https://doi.org/10.3847/1538-4357/aac326)
- Tacchella, S., Diemer, B., Hernquist, L., et al. 2019, *MNRAS*, 487, 5416, doi: [10.1093/mnras/stz1657](https://doi.org/10.1093/mnras/stz1657)
- Torrey, P., Vogelsberger, M., Marinacci, F., et al. 2019, *MNRAS*, 484, 5587, doi: [10.1093/mnras/stz243](https://doi.org/10.1093/mnras/stz243)
- Trayford, J. W., & Schaye, J. 2019, *MNRAS*, 485, 5715, doi: [10.1093/mnras/stz757](https://doi.org/10.1093/mnras/stz757)
- Tremonti, C. A., Heckman, T. M., Kauffmann, G., et al. 2004, *ApJ*, 613, 898, doi: [10.1086/423264](https://doi.org/10.1086/423264)
- Turner, E. L., & Gott, J. R., I. 1976, *ApJS*, 32, 409, doi: [10.1086/190403](https://doi.org/10.1086/190403)
- Wang, L., & Kauffmann, G. 2008, *MNRAS*, 391, 785, doi: [10.1111/j.1365-2966.2008.13907.x](https://doi.org/10.1111/j.1365-2966.2008.13907.x)
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, 465, 3291, doi: [10.1093/mnras/stw2944](https://doi.org/10.1093/mnras/stw2944)
- Westfall, K. B., Cappellari, M., Bershad, M. A., et al. 2019, *AJ*, 158, 231, doi: [10.3847/1538-3881/ab44a2](https://doi.org/10.3847/1538-3881/ab44a2)
- Wong, T., Hughes, A., Ott, J., et al. 2011, *ApJS*, 197, 16, doi: [10.1088/0067-0049/197/2/16](https://doi.org/10.1088/0067-0049/197/2/16)

APPENDIX

A. CONVERGENCE WITH SIMULATION RESOLUTION AND PIXEL SIZE

We test the convergence of our results by way of comparing the average star formation rate profiles (shown in Fig. 16) in Figure 17. We use the two lower resolution, same volume counterparts of the standard TNG50-1 run, namely TNG50-2 and TNG50-3. TNG50-2(TNG50-3) has a mass-resolution of 6.8(54.2) $M_\odot$ for gas, and 36.3(290.4) $M_\odot$ for dark matter, making it roughly a factor of 8(64) coarser compared to TNG50-1. From the left panel in the figure, we observe that using a coarser-resolution box has marginal influence on our results both in terms of slope and normalisation, except at $R \geq 20$ kpc. Given the mass range we are working with, we expect these regions to correspond to the far-outskirts of the galaxies with very few star-forming gas particles, thus making the profile noisier due to poorer sampling.

In the right panel of the figure, we show a comparison of the same profile in TNG50-1, but this time for different choices of pixel size (image resolution). The profiles do not appear to have a dependence on pixel size for pixels in the central parts where the simulation resolution elements are smaller due to higher densities. In these regions, star formation is adequately resolved with a few 10s of particles contributing to each pixel on average. At large radii, we start to see a weak variation of the slope with pixel size, such that bigger pixels give rise to steeper profiles. This is because in galaxy outskirts, where star formation does not have a uniform coverage, larger pixels tend to smooth over small-scale spatial patterns hence acquiring increasingly lower values of Σ_{sfr} as we go farther out. Lastly, as these profiles are binned and only composed of star-forming spaxels, they become noisier in the outskirts due to arbitrarily low area coverage as well as worsening Poissonian statistics.




Figure 17. Σ_{sfr} profile variation with the simulation resolution (left) and choice of pixel size (right) for star-forming regions in our fiducial galaxy sample ($M_\star = 10^{9-10}$ at $z = 0$). In keeping with Figure 16, only star-forming pixels are used to construct these profiles.

B. RECONSTRUCTING THE ORIGINAL SPACE FROM PCA RESULTS

Let $\mathbf{X}$ be the data matrix corresponding to the initial space with n rows for the samples and m columns denoting features ($m = 8$ in our study). As described in §5.2.1, we standardise our dataset before conducting PCA by first subtracting the mean vector $\boldsymbol{\mu}$ from all rows and dividing them element wise by the standard deviation vector $\boldsymbol{\sigma}$. In our case, we made $\boldsymbol{\sigma}$ to be the Σ_{sfr} -weighted standard deviations. Doing this gives us the corresponding matrix of standardised data, also known as z-scores $\mathbf{Z}$. After the PCA analysis, we obtain our results in the form of a coefficient matrix $\mathbf{C}$, which is an $m \times m$ matrix whose columns are the m eigenvectors representing the directions of the principal components (PCs). Then, the principal component scores are nothing but a projection of our original space along each of the PC directions. These are given in matrix form as $\mathbf{P} = \mathbf{Z}\mathbf{C}$, with the same dimensionality as our original parameter space, i.e., $n \times m$.

Now, in order to reconstruct the original data from these scores, we apply the inverse operation such that $\mathbf{Z} = \mathbf{P}\mathbf{C}^{-1}$. Due to the orthonormality of $\mathbf{C}$, this is equivalent to $\mathbf{Z} = \mathbf{P}\mathbf{C}^T$. Finally, to obtain the reconstructed $\mathbf{X}$, we multiply each column of $\mathbf{Z}$

by the corresponding σ value and add the corresponding mean μ . To obtain an approximate reconstruction using only a few PCs, say k in number, one would only use the first k PC scores, keeping just the first k columns of $\mathbf{C}$ in the calculation above.