<https://helda.helsinki.fi>

Interacting galaxies in the IllustrisTNG simulations - I : Triggered star formation in a cosmological context

Patton, David R.

2020-06

Patton , D R , Wilson , K D , Metrow , C J , Ellison , S L , Torrey , P , Brown , W , Hani , M H , McAlpine , S , Moreno , J & Woo , J 2020 , ' Interacting galaxies in the IllustrisTNG simulations - I : Triggered star formation in a cosmological context ' , Monthly Notices of the Royal Astronomical Society , vol. 494 , no. 4 , pp. 4969-4985 . <https://doi.org/10.1093/mnras/staa913>

<http://hdl.handle.net/10138/320658>

<https://doi.org/10.1093/mnras/staa913>

---

cc_by_nc_sa

publishedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Interacting galaxies in the IllustrisTNG simulations - I: Triggered star formation in a cosmological context

David R. Patton¹,¹★ Kieran D. Wilson,¹ Colin J. Metrow,¹ Sara L. Ellison²,
Paul Torrey³, Westley Brown,¹ Maan H. Hani,² Stuart McAlpine⁴,⁴ Jorge Moreno⁵
and Joanna Woo²

¹Department of Physics and Astronomy, Trent University, 1600 West Bank Drive, Peterborough ON K9L 0G2, Canada

²Department of Physics and Astronomy, University of Victoria, Finnerty Road, Victoria BC V8P 1A1, Canada

³Department of Astronomy, University of Florida, 211 Bryant Space Science Center, Gainesville, FL 32611, USA

⁴Department of Physics, University of Helsinki, Gustaf Hållströmin katu 2a P.O. Box 64, FI-00014 Helsinki, Finland

⁵Department of Physics and Astronomy, Pomona College, Claremont, CA 91711, USA

Accepted 2020 March 19. Received 2020 February 21; in original form 2019 September 24

ABSTRACT

We use the IllustrisTNG cosmological hydrodynamical simulations to investigate how the specific star formation rates (sSFRs) of massive galaxies ($M_* > 10^{10} M_\odot$) depend on the distance to their closest companions. We estimate sSFR enhancements by comparing with control samples that are matched in redshift, stellar mass, local density, and isolation, and we restrict our analysis to pairs with stellar mass ratios of 0.1 to 10. At small separations (~ 15 kpc), the mean sSFR is enhanced by a factor of 2.0 ± 0.1 in the flagship (110.7 Mpc)³ simulation (TNG100-1). Statistically significant enhancements extend out to 3D separations of 280 kpc in the (302.6 Mpc)³ simulation (TNG300-1). We find similar trends in the EAGLE and Illustris simulations, although their sSFR enhancements are lower than those in TNG100-1 by about a factor of two. Enhancements in IllustrisTNG galaxies are seen throughout the redshift range explored ($0 \leq z < 1$), with the strength of the enhancements decreasing with increasing redshift for galaxies with close companions. In order to more closely compare with observational results, we separately consider 2D projected distances between galaxies in IllustrisTNG. We detect significant sSFR enhancements out to projected separations of 260 kpc in TNG300-1, with projection effects diluting the size of the enhancements by about 20 per cent below 50 kpc. We find similar sSFR enhancements in TNG100-1 and Sloan Digital Sky Survey galaxies, with enhancements extending out to projected separations of about 150 kpc for star-forming galaxies at $z < 0.2$. Finally, by summing over all separations, we estimate that the presence of closest companions boosts the average sSFR of massive galaxies in TNG100-1 by 14.5 per cent.

Key words: methods: data analysis – methods: numerical – galaxies: evolution – galaxies: interactions – galaxies: star formation – galaxies: statistics.

1 INTRODUCTION

Interactions and mergers of galaxies play an important role in the formation and evolution of galaxies within the Λ CDM paradigm. Idealized binary merger simulations predict that galaxy properties undergo profound changes throughout the merger sequence, with the strength and nature of these changes depending on the orbital properties and mass ratio of the galaxy pair, and on the underlying properties of the galaxies themselves (Di Matteo et al. 2007;

Cox et al. 2008; Lotz et al. 2008; Chilingarian et al. 2010; Moreno et al. 2015). Merger simulations predict that interactions can trigger gas inflows, which lead to enhanced star formation, diluted gas-phase metallicities, and increased AGN activity (Mihos & Hernquist 1996; Di Matteo, Springel & Hernquist 2005; Rupke, Kewley & Barnes 2010; Torrey et al. 2012; Hopkins et al. 2013; Blumenthal & Barnes 2018; Moreno et al. 2019). These changes are often most apparent during and shortly after coalescence, but they are also seen earlier in the merger sequence, following close encounters between the galaxies.

Observational studies of galaxies have provided general support for these predictions, particularly at low redshift. For example,

★ E-mail: dpatton@trentu.ca

studies of galaxy pairs have shown that galaxies with close companions have enhanced star formation (Barton, Geller & Kenyon 2000; Ellison et al. 2008; Woods et al. 2010; Scudder et al. 2012; Patton et al. 2013; Cao et al. 2016), diluted metallicities (Ellison et al. 2008; Kewley et al. 2010; Rupke, Kewley & Chien 2010; Scudder et al. 2012; Bustamante et al. 2020), increased structural asymmetries (Hernández-Toledo et al. 2005; Patton et al. 2005; De Propriis et al. 2007; Casteels et al. 2014; Patton et al. 2016) and higher AGN fractions (Alonso et al. 2007; Ellison et al. 2011; Satyapal et al. 2014; Weston et al. 2017; Ellison et al. 2019) when compared with control samples of relatively isolated galaxies.

However, the link between merger simulations and observations of galaxy pairs is indirect. Idealized binary merger simulations assume pre-selected orbits and orientations, and span a limited combination of stellar masses, gas fractions, bulge fractions, etc. In addition, most merger simulations do not include a circumgalactic medium or external gas accretion, nor do they include neighbouring galaxies or realistic group/cluster environments. The advent of relatively high resolution cosmological zoom-in simulations has allowed for some progress in studying mergers within a more realistic cosmological context (Sparre & Springel 2016; Bustamante et al. 2018); however, the number of merging systems studied in this manner is too small to cover a representative sample of interacting and merging galaxies.

Observationally, galaxy pairs have been found in a wide range of environments (Barton et al. 2007; McIntosh et al. 2008; Ellison et al. 2010; Lin et al. 2010; Patton et al. 2016). However, the interpretation of these galaxy pair samples is limited by projection effects, wherein the 3D separations and relative velocities of individual pairs are unknown (Kitzbichler & White 2008). As such, the orbital histories of most galaxy pairs are unconstrained. Moreover, with each pair seen at a single epoch, information about the evolutionary and orbital histories of the galaxies is very limited.

Fortunately, the availability of increasingly realistic cosmological hydrodynamical simulations provides a way forward on many of these issues. Simulations such as Illustris (Vogelsberger et al. 2014a), Horizon-AGN (Dubois et al. 2014), EAGLE (Schaye et al. 2015), MassiveBlack-II (Khandai et al. 2015), SIMBA (Davé et al. 2019), and IllustrisTNG (Nelson et al. 2019) yield thousands of massive galaxies within volumes spanning at least $(100 \text{ Mpc})^3$, with properties that are well matched to galaxy populations at low redshift. Galaxy interactions and mergers arise naturally within these simulations, with a full distribution of orbits, spin-orbit orientations, environments, and galaxy properties.

While these cosmological simulations are naturally of much lower mass resolution than typical merger simulations, the presence of tidal tails and star forming regions in synthetic images (Torrey et al. 2015; Pearson et al. 2019) suggests that they may nevertheless capture triggered star formation in interacting galaxies. Indeed, Martin et al. (2017) find that Horizon-AGN galaxies that are within 1 Gyr of a merger have SFRs that are enhanced by a factor of ~ 1.7 , while Rodríguez Montero et al. (2019) find that recently merged galaxies in the SIMBA simulations have sSFRs that are enhanced by a factor of 2–3.

In this study, we aim to bridge the gap between merger simulations and observations of galaxy pairs by identifying galaxies and their closest companions within the IllustrisTNG cosmological hydrodynamical simulations, focussing on the sSFRs of galaxies and the degree to which they are enhanced due to interactions. By identifying galaxy pairs from these simulations using both 3D and projected separation, our goal is to facilitate a direct comparison with observed galaxy pairs. In a companion paper (Hani et al.

2020), we extend this analysis to the sSFRs of post-merger galaxies in IllustrisTNG.

In the following section, we describe the simulations and the methodology used to identify the closest companion of each massive galaxy in the simulations. In Section 3, we show how the sSFR and its enhancement depends on the 3D distance to the closest companion in the highest resolution and largest volume IllustrisTNG simulations. In Section 4, we investigate how these results depend on resolution and volume within IllustrisTNG, and we carry out a direct comparison with the Illustris-1 and EAGLE simulations. We then show in Section 5 how IllustrisTNG sSFR enhancements depend on projected separation, and we make a direct comparison with sSFR enhancements in the Sloan Digital Sky Survey (hereafter SDSS). We discuss our findings in Section 6 and summarize our conclusions in Section 7.

2 SIMULATIONS AND METHODOLOGY

2.1 The IllustrisTNG simulations

Our study is primarily focused on the cosmological hydrodynamical simulations from the IllustrisTNG¹ suite. IllustrisTNG is the successor to the Illustris simulations (see below), with numerous updates that are described in the public data release (Nelson et al. 2019). Weinberger et al. (2018) and Pillepich et al. (2018a) provide a complete description of the physical models. Additional details about these simulations are given by Pillepich et al. (2018b), Springel et al. (2018), Nelson et al. (2018), Naiman et al. (2018), and Marinacci et al. (2018).

The IllustrisTNG simulations are publicly available for cubic periodic volumes of co-moving side length 110.7 and 302.6 Mpc. For each volume, baryonic, and dark-matter only simulations are available at three mass resolution levels. The baryonic simulations follow the evolution of dark matter, stars, gas, and supermassive black holes from a redshift of 127 to the present epoch using the AREPO moving-mesh code (Springel 2010). Star formation occurs above a threshold density criterion, and is regulated by a subgrid model for the interstellar medium (Springel & Hernquist 2003; Nelson et al. 2015; Pillepich et al. 2018a). Where needed for stellar evolution and mass return calculations, a Chabrier (2003) initial mass function is used.

Our primary focus is on the highest resolution run for the $(110.7 \text{ Mpc})^3$ volume (TNG100-1) and on the highest resolution run of the $(302.6 \text{ Mpc})^3$ volume (TNG300-1). However, we also use two of the lower resolution simulations (TNG100-2 and TNG100-3) to assess how our primary results depend on resolution. Subhalo catalogues are available at 100 redshifts, including 50 snapshots at $z < 1$. Following Nelson et al. (2019), we exclude SUBFIND subhaloes for which the SubhaloFlag parameter has been set to 0, as they are known to be sub-galactic clumps and are not considered to be bona fide galaxies. Properties of these simulations (and others described below) are summarized in Table 1.

2.2 The Illustris-1 simulation

While this paper focuses on the IllustrisTNG simulations, we also carry out a direct comparison with the original Illustris simulations (Vogelsberger et al. 2014a). Nelson et al. (2015) describe the Illustris public data release, while Genel et al. (2014) and Vogelsberger et al.

¹<http://www.tng-project.org>

Table 1. A summary of the simulations used in this analysis, including the co-moving side length of the simulation box (L_{box}), the baryonic particle mass (M_{baryon}), the average number of massive host galaxies ($M_* > 10^{10} M_{\odot}$) per snapshot ($\langle N^{\text{host}} \rangle$), the average time between snapshots (Δt_{snap}) at $z \lesssim 1$, and the 3D separation at which the crowding fraction reaches 20 per cent (r_{crowd}).

Simulation Name	L_{box} (Mpc)	M_{baryon} ($M_{\odot}$)	$\langle N^{\text{host}} \rangle$	Δt_{snap} (Myr)	r_{crowd} (kpc)
TNG100-1	110.7	1.4×10^6	5976	162	15.9
TNG100-2	110.7	1.1×10^7	4382	162	17.7
TNG100-3	110.7	9.0×10^7	2534	162	23.4
TNG300-1	302.6	1.1×10^7	90279	162	17.4
Illustris-1	106.5	1.6×10^6	5789	155	18.0
EAGLE	100.0	1.8×10^6	3521	121	16.3

(2014b) provide additional information about these simulations. As with IllustrisTNG, the Illustris simulations use the AREPO moving-mesh code (Springel 2010). The physical models are described by Vogelsberger et al. (2013). For this study, we consider only the flagship Illustris-1 simulation, which is the highest resolution (106.5 Mpc)³ hydrodynamical simulation from Illustris. As seen in Table 1, Illustris-1 is comparable to TNG100-1 in terms of both volume and resolution. Illustris-1 has 51 snapshots at $z < 1$, with time sampling that is very similar to IllustrisTNG.

2.3 The EAGLE simulation

Given the inherent connections between Illustris and IllustrisTNG, we carry out an additional comparison with the largest simulation from the *Evolution and Assembly of GaLaxies and their Environments* (EAGLE) project. This simulation, denoted Ref-L0100N1504, is comparable to both TNG100-1 and Illustris-1 in terms of volume and mass resolution (see Table 1). However, EAGLE is a smooth particle hydrodynamics code and does not use an adaptive mesh, and, more notably, adopts a considerably different subgrid scheme. Galaxy catalogues and particle properties are publicly available for 29 snapshots at $z < 20$. We instead use galaxy catalogues from a set of 201 ‘snipshots’, including 65 snipshots at $z < 1$. Following McAlpine et al. (2016), we exclude subhaloes for which the Spurious flag has been set to 1, as these objects are known to be dense concentrations within subhaloes and therefore are artefacts of the SUBFIND algorithm. For a full description of the simulation and calibration strategy, see Schaye et al. (2015) and Crain et al. (2015), respectively.

We briefly outline the model governing star formation here, as it is the most important model related to this study. Star formation occurs above the metallicity-dependent threshold of Schaye (2004), whereby ‘cold’ gas particles ($T \sim 10^4 \text{ K}$) are stochastically converted into star particles at a pressure-dependent rate that reproduces the observed Kennicutt–Schmidt star formation law (see Schaye & Dalla Vecchia 2008). Each newly formed star particle represents a simple stellar population with a Chabrier (2003) initial mass function. The feedback from massive stars and supernovae is injected thermally and stochastically by raising the surrounding gas to a fixed temperature increment as described in Dalla Vecchia & Schaye (2012).

2.4 Complete samples of galaxies with $M_* > 10^9 M_{\odot}$

For each of the simulations under consideration, we begin by constructing a sample of galaxies which is complete above a stellar

mass of $10^9 M_{\odot}$. Galaxies in this regime are resolved with more than 500 stellar particles per galaxy for the highest resolution 100 Mpc volumes (TNG100-1, Illustris-1, and EAGLE) and at least 90 stellar particles per galaxy for the lower resolution TNG300-1 simulation (see Table 1).

However, in cases where two galaxies are close to one another, it can be difficult to correctly disentangle their stellar mass. For all of the simulations analysed in this study, the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009) was used to identify galaxies. In cases of an overlapping galaxy pair, SUBFIND will often assign some of the stellar mass from the outskirts of the lower mass galaxy to the higher mass galaxy (Rodríguez-Gomez et al. 2015). This effect is often referred to as numerical stripping, and it can result in spurious mass estimates for close galaxy pairs. Moreover, galaxies which are close to our selected lower limit in stellar mass may be erroneously removed from the sample by numerical stripping.

While numerical stripping is largely irrelevant for most galaxies, it is particularly problematic for galaxies that are undergoing close encounters – a population that is of particular interest in our study. We therefore attempt to minimize the effects of numerical stripping on our mass-limited samples by identifying systems which are affected by crowding and, in some cases, by correcting for this effect.

To begin, we extract two estimates of stellar mass for each galaxy: the stellar mass at the epoch of interest (hereafter M_*^{now}) and the maximum recent stellar mass (hereafter M_*^{max}), including the epoch of interest. In cases where numerical stripping has occurred, M_*^{max} may provide an improved estimate of the galaxy’s true current mass. However, if one probes too far back in time, this approach can overestimate the current stellar mass if there has subsequently been physical stripping (e.g. the removal of stellar mass upon infall into a galaxy group or cluster).

In order to minimize the effects of numerical (but not physical) stripping, we elect to evaluate M_*^{max} within the past 0.5 Gyr.² Merger simulations suggest that this time frame should be long enough to catch the galaxies when they were clearly separated from one another (e.g. Patton et al. 2013). In addition, this time frame should be short enough that any physical stripping of stellar mass would be limited. Given the time sampling of the simulations under consideration (see Table 1), we trace each galaxy’s main progenitor branch back in time by 2–3 snapshots for IllustrisTNG, 2–4 snapshots for Illustris-1, and 3–5 snapshots for EAGLE. We note also that the main conclusions of this study are not sensitive to the precise time frame used for these calculations. For example, changing the time frame by $\pm 0.25 \text{ Gyr}$ has a negligible impact on our results.

We then quantify the degree of crowding in each galaxy’s immediate surroundings by measuring the relative separation between the galaxy and its most overlapping companion, taking into account the separation and size of the galaxy and its companions. In order to ensure that we do not miss any galaxies above our $10^9 M_{\odot}$ threshold, we initially consider all galaxies with $M_*^{\text{max}} > 10^9 M_{\odot}$ (thereby including some galaxies which have $M_*^{\text{now}} < 10^9 M_{\odot}$), and we search for all companions with M_*^{max} greater than 10 per cent of the host galaxy’s M_*^{now} . We calculate the relative separation

²Our approach is similar to that of Rodríguez-Gomez et al. (2015), who estimate the mass of both galaxies at the time that the secondary galaxy had its maximum progenitor mass. However, as they do not apply a restriction on lookback time, their approach avoids both numerical and physical stripping.

Figure 1. In the upper panel, the mean ratio of the current stellar mass (M_*^{now}) to the maximum stellar mass within the past 500 Myr (M_*^{max}) is plotted versus the relative separation (r_{sep}) between each galaxy and its closest companion in TNG100-1. In the lower panel, mean $M_*^{\text{now}}/M_*^{\text{max}}$ is plotted versus the 3D distance (r) between each galaxy and its closest companion for both the ‘Uncorrected’ sample (using the true M_*^{now} in each case) and the ‘Corrected’ sample (for which we replace M_*^{now} with M_*^{max} when $r_{\text{sep}} < 2$). For each set of data, the solid line depicts the mean, and the shaded region depicts the 2σ standard error in the mean.

(hereafter r_{sep}) of each neighbouring galaxy, as follows:

$$r_{\text{sep}} = \frac{r}{R_{1/2}^{\text{host}} + R_{1/2}^{\text{comp}}}, \quad (1)$$

where r is the 3D separation between the centres of the two galaxies and $R_{1/2}^{\text{host}}$ and $R_{1/2}^{\text{comp}}$ are the stellar half-mass radii of the host and companion respectively. For each host galaxy, we record the minimum r_{sep} of its neighbouring galaxies.

We investigate the dependence of stripping on the relative separation between TNG100-1 galaxies by plotting the mean $M_*^{\text{now}}/M_*^{\text{max}}$ versus r_{sep} in the upper panel of Fig. 1. At large r_{sep} , $M_*^{\text{now}}/M_*^{\text{max}}$ is close to unity, indicating that few galaxies have experienced a recent decline in their stellar mass. This ratio initially declines gradually as r_{sep} decreases, reaching 90 per cent at $r_{\text{sep}} \sim 2$. Given that galaxies are relatively well-separated at $r_{\text{sep}} > 2$, we interpret this trend as being driven primarily by physical stripping. $M_*^{\text{now}}/M_*^{\text{max}}$ declines sharply at smaller r_{sep} , dropping below 80 per cent at $r_{\text{sep}} \sim 0.5$. Given that there is substantial overlap between galaxies at these small relative separations, we attribute this more rapid decline as being caused by numerical stripping.

In order to identify the regime in which numerical stripping is prevalent, we carried out an extensive visual inspection of synthetic stellar images of galaxies from both TNG100-1 (Nelson et al. 2019) and Illustris-1 (Torrey et al. 2015). This inspection focused on systems with $r_{\text{sep}} < 3$, which corresponds to typical 3D separations of $r \lesssim 30$ kpc for TNG100-1 and $r \lesssim 60$ kpc for Illustris-1. At $r_{\text{sep}} < 1$ (i.e. for galaxy pairs with largely overlapping stellar distributions), we find that nearly every system appears to suffer from substantial numerical stripping, with a significant fraction of a galaxy’s stellar mass being assigned to a close neighbour (or vice versa). For these galaxies, we therefore set $M_* = M_*^{\text{max}}$. Conversely, at $r_{\text{sep}} > 2$, the

Figure 2. Synthetic stellar composite images of four $z = 0$ TNG100-1 host galaxies and their closest companions are shown, with one column corresponding to each pair. The relative separation of each pair is shown at the top. In the upper row, each image is centred on the host galaxy and shows all of the subhaloes within the specified volume (i.e. both the host and companion are shown). In the middle row, each image is again centred on the host galaxy, but in this case only the stellar mass associated with the host galaxy is shown. In the lower row, each image is centred on the companion galaxy, and only the stellar mass associated with the companion galaxy is shown. In the middle and lower rows, white circles are used to denote one and two $R_{1/2}$. All images are labelled with the galaxy SUBFIND IDs. The width of each image corresponds to $7 R_{1/2}$ of the galaxy at the centre of the image, and the white bar at the upper left of each image depicts a physical scale of 10 kpc.

SUBFIND algorithm appears to assigns the stellar mass to the correct galaxy in almost every instance. We therefore set $M_* = M_*^{\text{now}}$ when $r_{\text{sep}} > 2$.

Within the intermediate range of $1 < r_{\text{sep}} < 2$, SUBFIND tends to do a reasonable but suboptimal job at separating the two galaxies. We illustrate this in Fig. 2, where we present stellar composite images³ of four TNG100-1 galaxies that have a close companion at $1 < r_{\text{sep}} < 2$. For each of these galaxy pairs, we show images of the host plus companion (top row), host only (middle row), and companion only (bottom row). These pairs are oriented such that their projected separation is at least 90 per cent of their 3D separation, making the images easier to interpret. In each case, we can clearly see that some stellar mass is incorrectly assigned to the more massive member of the pair. To be conservative, we therefore set $M_* = M_*^{\text{max}}$ for all galaxies with $r_{\text{sep}} < 2$.

In order to assess the impact of this procedure on our sample, we plot the mean $M_*^{\text{now}}/M_*^{\text{max}}$ versus the 3D distance to the closest companion (r) in the lower panel of Fig. 1. For the original ‘Uncorrected’ TNG100-1 sample, $M_*^{\text{now}}/M_*^{\text{max}}$ drops below 90 per cent at $r \lesssim 20$ kpc, presumably driven by numerical stripping. However, after replacing M_*^{now} with M_*^{max} for $r_{\text{sep}} < 2$, the ‘Corrected’ mean $M_*^{\text{now}}/M_*^{\text{max}}$ remains above 95 per cent at all separations, thereby compensating for numerical stripping. This procedure increases the stellar mass for only 1.2 per cent of the galaxies in TNG100-1, confirming that only a small fraction of the sample is significantly

³These images were generated using the IllustrisTNG image visualization tool available at <http://www.tng-project.org>, with the Stellar Composite field. See Nelson et al. (2019) for details.

Figure 5. The mean sSFR of TNG100-1 host galaxies and their controls are plotted versus the distance to the host galaxy’s closest companion (r) in the upper panel. The mean sSFR values have been smoothed using a variable width box kernel. The corresponding sSFR enhancements (Q) are plotted in the lower panel, with $Q = 1$ (dashed line) corresponding to no enhancement. The shaded region surrounding each line shows the 2σ standard error in the mean. The shaded region on the left side of both panels depicts separations at which $f_{\text{crowd}} > 0.2$ (see Section 2.7).

in the mean sSFR of the controls at small r is likely due to the fact that the controls are not strictly isolated; rather, the closest companions of the controls are matched to the separation of the second closest companions of their host galaxies. That is, our methodology is designed to detect the influence of the closest companion, but it does not capture the influence of any additional companions.⁵

In order to more directly compare the sSFRs of host galaxies and their controls, we also estimate the sSFR enhancement (hereafter Q) in a given r bin by dividing the mean sSFR of the host galaxies by the mean sSFR of their controls; i.e.

$$Q(\text{sSFR}) = \frac{\langle \text{sSFR}_{\text{hosts}} \rangle}{\langle \text{sSFR}_{\text{controls}} \rangle}. \quad (2)$$

In the lower panel of Fig. 5, we see values of $Q \sim 1$ at $r \gtrsim 200$ kpc, indicating that there is no statistically significant enhancement of the host galaxy sSFRs at large separations. Conversely, below 200 kpc, sSFR enhancements increase steadily as the separation decreases, reaching a maximum value of $Q \sim 2.0 \pm 0.1$ (a factor of two enhancement in sSFR) at the crowding threshold of this sample ($r_{\text{crowd}} = 15.9$ kpc). Moreover, this steady increase extends well into the region affected by crowding ($r < r_{\text{crowd}}$), suggesting that our choice of $f_{\text{crowd}} = 0.2$ may be quite conservative. Statistically significant (2σ) enhancements in the mean sSFR extend out to 210 kpc, and have very high statistical significance (12σ) at the smallest separations we probe ($r = r_{\text{crowd}}$).

⁵Most of the rise in the mean sSFR of the control sample at small r disappears if we restrict the host galaxy sample to $r_2 > 100$ kpc, thereby creating a relatively isolated sample of galaxy pairs.

3.2 Synthetic images of TNG100-1 host galaxies with close companions

Given the increase in mean sSFR towards small r that we have found, it is natural to speculate that galaxies in these simulations are exhibiting enhancements in their star formation due to close encounters, as has been reported in high resolution idealized merger simulations and cosmological zoom-in simulations. One obvious test of this interpretation is to look for morphological signatures of interactions and disturbances in the synthetic images of these galaxies.

We present stellar composite images of 18 representative TNG100-1 host galaxies with relatively close companions ($r \sim 20$ kpc) in Fig. 6. At these separations, the mean sSFR is enhanced by a factor of ~ 2 for this simulation (see Fig. 5). To cover a representative sample of galaxies in both major and minor pairs, this sample includes images of six host galaxies at low redshift ($z < 0.5$) in each of the following ranges of the companion stellar mass ratio: $0.1 < \mu_c < 0.5$ (upper row), $0.5 < \mu_c < 2$ (middle row), and $2 < \mu_c < 10$ (lower row). In addition, to aid the interpretation of these images, we select pairs which have projected separations (in the image plane) which are within 10 per cent of their 3D separations. Apart from these criteria, these host galaxies were selected at random from the TNG100-1 sample. To be explicit, these galaxies were not selected based on their visual morphologies or other galaxy properties.

Morphological signs of interactions are clearly visible in nearly all of the galaxy pairs depicted in Fig. 6. A diversity of galaxy types and pairs is seen, including pairs of star-forming galaxies (potential wet mergers), pairs of passive galaxies (potential dry mergers), and star-forming+passive pairs (potential mixed mergers). For host galaxies which have controls with non-zero sSFRs, the median Q is 2.1, confirming that our random selection has not generated a subset of galaxies with unusually high (or low) sSFR enhancements. We conclude that the enhanced sSFR we have detected in TNG100-1 does indeed appear to be associated with interactions between galaxies, at least at the relatively small separations ($r \sim 20$ kpc) where the sSFR enhancements are the strongest.

3.3 Enhanced sSFR versus 3D separation in TNG300-1

The TNG300-1 simulation allows us to investigate trends in the enhanced sSFR within a volume that is about 20 times larger than TNG100-1, albeit at lower resolution (see Table 1). This larger volume substantially reduces the statistical uncertainties on the properties measured, and also spans a wider range of cosmological environments.

In Fig. 7, we show the mean sSFR and sSFR enhancement for the TNG300-1 volume. In both panels, we see trends that are qualitatively similar to those of TNG100-1 (Fig. 5), yet yielding much smaller uncertainties. Close inspection of this figure and the underlying data reveal that the mean sSFR is enhanced ($Q > 1$) at the 2σ level for $r < 280$ kpc. While we restrict Fig. 7 to $r < 300$ kpc, we have also extended the analysis out to 1000 kpc (not shown), finding that Q is consistent with 1.0 (within 2σ) throughout the range 300–1000 kpc for TNG300-1. The enhancements increase steadily as r decreases, reaching $Q = 1.72 \pm 0.02$ at the smallest separation for which $f_{\text{crowd}} < 0.2$ (17.4 kpc).⁶

⁶The enhancements reach $Q = 1.98 \pm 0.04$ at the smallest separation shown (9.5 kpc), which corresponds to $f_{\text{crowd}} = 0.5$ (at that separation, half of the host galaxies have been excluded due to crowding).

Figure 6. Synthetic stellar composite images of 18 TNG100-1 host galaxies with close companions. All images are 80 physical kpc on a side, and are centred on the host galaxy. Each image depicts the host galaxy, its closest companion, and any other stellar mass in the volume that is shown. The upper row contains host galaxies with a relatively low mass close companion (companion stellar mass ratio $0.1 < \mu_c < 0.5$), the middle row contains host galaxies with a roughly equal mass close companion ($0.5 < \mu_c < 2$), and the lower row contains host galaxies with a relatively high mass close companion ($2 < \mu_c < 10$).

Figure 7. The mean sSFR of TNG300-1 host galaxies and their controls are shown in the upper panel, while the corresponding sSFR enhancement is shown in the lower panel. The shaded region surrounding each line shows the 2σ standard error in the mean.

3.4 Redshift dependence

As mentioned previously, our sample includes host galaxies in the redshift range $0 \leq z < 1$. The sSFR enhancements we have reported should be representative of galaxies throughout this redshift range, as long as there is no strong redshift evolution in the mean sSFR

enhancement at $z < 1$. We now assess the validity of this approach by dividing our TNG300-1 sample into five narrower redshift bins.

In the upper panel of Fig. 8, we plot the mean sSFR of host galaxies as a function of r for five redshift bins. At all separations, we find that the mean sSFR increases with redshift. The mean sSFR increases by about a factor of 5 from our lowest redshift bin ($0 \leq z < 0.2$) to our highest redshift bin ($0.8 < z < 1$). This trend is consistent with redshift-dependent increases in the SFRs of IllustrisTNG galaxies reported by the IllustrisTNG collaboration (Torrey et al. 2018; Weinberger et al. 2018; Donnari et al. 2019), and is qualitatively consistent with the observed increases in SFRs with redshift (Noeske et al. 2007; Whitaker et al. 2012; Speagle et al. 2014).

In the upper panel of Fig. 8, we see a similar relationship between the mean sSFR and r within each redshift bin, with the mean sSFR clearly rising as r decreases. In the middle panel of Fig. 8, we plot the mean sSFR enhancement (Q) as a function of r for all five redshift bins. We find that the enhancements in all five redshift bins are similar to one another, with statistically significant enhancements extending out to at least 200 kpc in every case. This suggests that interactions cause increased star formation throughout the full redshift range of $0 \leq z < 1$, despite strong evolution in the underlying sSFRs of galaxies. Moreover, the fact that similar trends are seen in all five redshift bins demonstrates that it is reasonable for us to have combined all of the $z < 1$ snapshots together for our analysis.

Nevertheless, close inspection of the middle panel of Fig. 8 reveals that the sSFR enhancements are somewhat smaller at higher redshifts. To better visualize this trend, we define a new quantity ΔQ , which is the difference between Q in a given redshift bin and Q in the lowest redshift bin ($z < 0.2$). In the lower panel of Fig. 8, we plot ΔQ versus r for three representative redshift bins. This plot shows that the sSFR enhancements are significantly lower at high

Figure 8. In the upper panel, the mean sSFR of TNG300-1 host galaxies is plotted versus the 3D distance to their closest companion (r) for five redshift bins spanning $0 \leq z < 1$. For clarity, the trends of the control galaxies are not shown. In the middle panel, the mean sSFR enhancement (Q) is plotted versus r for the same redshift bins. In the lower panel, the offset in Q from the $z < 0.2$ redshift bin is plotted versus r (for clarity, we do not plot the $0.2 < z < 0.4$ and $0.6 < z < 0.8$ redshift bins). In all panels, the shaded coloured regions depict the 2σ standard error in the mean, and the light grey shaded region on the left-hand side depicts the overall $r < r_{\text{crowd}}$ region for TNG300-1. We find similar trends in sSFR enhancements at all redshifts, although the enhancements at small separations decrease with increasing redshift.

redshift ($0.8 < z < 1$) than at low redshift ($z < 0.2$), especially at $r < 50$ kpc. The size of this difference is modest ($\Delta Q \sim 0.3$ at $r = r_{\text{crowd}}$) but statistically significant.

This trend is consistent with predictions from simulations that SFR enhancements are smaller at higher redshift (Perret et al. 2014; Kaviraj et al. 2015; Fensch et al. 2017; Martin et al. 2017). There is also some observational support for this trend (e.g. Kaviraj et al. 2013; Lofthouse et al. 2017). While a detailed study of the underlying factors driving this trend in IllustrisTNG is beyond the scope of this analysis, this topic will be examined in a forthcoming paper.

4 COMPARISON WITH OTHER SIMULATIONS

4.1 Comparison with other IllustrisTNG simulations

In the previous section, we examined the mean sSFR and its enhancement in the TNG100-1 (Fig. 5) and TNG300-1 (Fig. 7) simulations. While similar trends are seen in both, we now carry out a direct comparison between these simulations and investigate how the sSFR enhancements depend on both the resolution and volume of the simulations.

In Fig. 9, we plot the sSFR enhancement versus the distance to the closest companion for the three different resolution runs of the IllustrisTNG 110.7 Mpc cube (see Table 1). In each case, we also display the TNG300-1 results for comparison. In the upper panel of Fig. 9, we see that while there is generally good agreement between the high resolution TNG100-1 and the lower resolution TNG300-1, the sSFR enhancements are consistently higher at $r \lesssim 100$ kpc

Figure 9. sSFR enhancements in TNG300-1 are compared with those in TNG100-1 (upper panel; 8 times higher resolution than TNG300-1), TNG100-2 (middle panel; same resolution as TNG300-1) and TNG100-3 (lower panel; 8 times lower resolution than TNG300-1). In each panel, the shaded region on the left depicts separations at which $f_{\text{crowd}} > 0.2$ for the TNG100 simulation. While we find similar overall trends at all resolutions, the sSFR enhancements at small separations are larger at higher resolution.

for TNG100-1, with differences significant at the 2σ level below 50 kpc. It is therefore possible that the sSFR enhancements have not yet converged at small scales in TNG100-1.

In the middle panel of Fig. 9, we compare two simulations at the same resolution but different volumes (TNG100-2 and TNG300-1). Here, the sSFR enhancements are consistent at all separations, although the 2σ error in the mean is much smaller for the larger volume simulation. This consistency suggests that the TNG100 box is large enough to cover a representative range of environments.

In the lower panel of Fig. 9, we compare TNG300-1 with the lowest resolution 100 Mpc simulation (TNG100-3). We find that, despite its comparatively poor resolution, TNG100-3 nevertheless captures much of the enhanced star formation seen in the larger volume and higher resolution TNG300-1. The most notable exception lies at small separations, where the lower resolution simulation suffers from significant crowding ($f_{\text{crowd}} > 0.2$) out to larger separations (23.4 kpc) than the other simulations, as expected (see Fig. 3). This suggests that higher resolution is particularly important in assessing enhanced star formation at smaller separations.

We have found that sSFR enhancements are captured in a relatively well converged fashion across the full TNG100 resolution suite. The degree of SFR enhancement is, however, impacted at small separations, with higher resolution simulations capturing higher SFR enhancements when galaxies are close. These trends are consistent with the findings of Sparre & Springel (2016), who carried out zoom-in simulations of four merging systems from the Illustris simulation over a range of resolutions. They report stronger SFR enhancements when running at higher resolution, with differences seen both before and after the completion of the merger. Given that their lower resolution runs are comparable in resolution to TNG100-1, this also suggests that the sSFR enhancements we have detected in TNG100-1 are likely to be underestimated. This interpretation has additional support from high resolution merger

Figure 10. The sSFR enhancements in EAGLE, Illustris-1 and TNG100-1 are plotted versus the distance to the closest companion. The shaded region on the left depicts the Illustris-1 crowding limit (18.0 kpc), which is slightly larger than that of TNG100-1 (15.9 kpc) and EAGLE (16.3 kpc). We find similar trends in all three simulations. However, the sSFR enhancements in EAGLE and Illustris-1 are significantly lower than those in TNG100-1.

simulations, which indicate that resolution plays an important role in capturing star formation in interacting systems (Bournaud et al. 2011; Renaud et al. 2014). On the other hand, when comparing high and low resolution runs of the Illustris simulations, Sparre et al. (2015) find similar fractions of star formation occurring in galaxies that reside above the star forming main sequence. This suggests that while lower resolution simulations (such as TNG100-3) may perform moderately well for modest increases in star formation, they may not be as successful at capturing strong starbursts as their higher resolution counterparts (such as TNG100-1).

4.2 Comparison with EAGLE and Illustris-1

Until now, we have investigated the enhancements in sSFR within the IllustrisTNG suite of simulations, exploring how these enhancements depend on resolution, volume, and redshift. However, these simulations were all produced using the same underlying gravity and hydrodynamics solver, and subgrid models. To provide a more independent comparison, we now analyse the EAGLE and Illustris-1 simulations to see if sSFR enhancements are present and, if so, to compare the size and extent of the sSFR enhancements with those in TNG100-1 (which is similar to both EAGLE and Illustris-1 in terms of resolution and volume; see Table 1).

We present measurements of sSFR enhancement as a function of r for the EAGLE, Illustris-1, and TNG100-1 simulations in Fig. 10. We find that the mean sSFR is also enhanced in EAGLE and Illustris-1. The sSFR enhancements in EAGLE rise to $Q \sim 1.6 \pm 0.1$ at the crowding limit of 16.3 kpc, with statistically significant (2σ) enhancements extending out to separations of 162 kpc. For Illustris-1, the enhancements rise to $Q \sim 1.7 \pm 0.1$ at the crowding limit of 18.0 kpc, with significant enhancements extending out to 145 kpc. Remarkably, the enhancements in EAGLE and Illustris-1 are consistent with one another (within the 2σ error bars that are shown) at all separations below 300 kpc.

We also plot the TNG100-1 enhancements (see Section 3.1) in Fig. 10. The sSFR enhancements in EAGLE and Illustris-1 are significantly lower than those in TNG100-1, by approximately a factor of two. These differences are statistically significant (at the 2σ level) below 140 kpc. The differences between TNG100-1 and

Illustris-1 are particularly notable, given that their SFR prescriptions are identical. As such, the differences in their sSFR enhancements can likely be attributed to changes in gas properties. Given that the sizes and gas fractions of IllustrisTNG and Illustris galaxies are known to be different (Pillepich et al. 2018b), it seems likely that galaxy properties play a central role in determining the level of SFR enhancement.

While there are some differences in the behaviours of the sSFR enhancements in TNG100-1, EAGLE, and Illustris-1, what is arguably more striking is the general similarities seen in all three simulation models. In each case, significant enhancements in sSFR are seen, with the strongest enhancements at the smallest separations probed. In addition, the results are consistent with modest enhancements extending out to ~ 150 – 200 kpc in all three simulations. This suggests that enhanced star formation due to the presence of a relatively nearby companion is a generic outcome for cosmological hydrodynamical simulations, despite the fact that the simulations were not tuned in advance to match this relationship.

5 CONNECTING COSMOLOGICAL SIMULATIONS TO OBSERVATIONS

We have presented an analysis of how the mean sSFR and its enhancement (with respect to a control sample) depends on the 3D distance to the closest companion in various cosmological hydrodynamical simulations. In order for us to connect these results with observations, we now investigate how the results change when we quantify proximity in the same way that observers do; namely, using projected separation and relative velocity along the line of sight. We then compare the results with a sample of galaxies from the SDSS.

5.1 Enhanced sSFR versus projected separation

We repeat our analysis from Section 3, now using projected quantities instead of 3D separation to quantify proximity. Rather than simply converting 3D separations into projected separations, we project the sample and then rerun the search for the closest companions, using the projected separation (r_p) and the relative velocity along the line of sight (Δv) to assess proximity. This more closely mirrors the approach used by observers to identify close companions, and accounts for the fact that the closest companion in 3D may not be the closest companion in projection (Patton et al. 2000; Kitzbichler & White 2008; Patton & Atfield 2008; Jian, Lin & Chieuh 2012). Moreover, some of the closest companions in this projected sample will be unrelated foreground/background galaxies, thereby including a source of contamination familiar to observers.

In principle, we could project the simulations from multiple random orientations in order to fully sample the distribution of companion separations. However, given the large volume of the simulations, we elect to simply project the simulations in each of the $x - y$, $y - z$, and $x - z$ planes, and use the variation between these orthogonal projections to assess the need (if any) for additional projections.

Here we describe our methodology in the context of the $x - y$ projection. We take the projected separation r_p to be the distance between the host and companion within the $x - y$ plane. To match the observational Δv criterion along the line of sight, we take the z -axis within the cube to be the line of sight, and compute Δv by summing the contributions from the relative velocity of the pair along the z -axis and the difference in recession velocities due to the Hubble flow.

Figure 11. In the upper panel, the mean sSFR of TNG300-1 host galaxies and their controls is plotted versus projected separation (r_p). We show only the projection on to the $x - y$ plane in this figure. The enhancement in sSFR is plotted versus r_p in the lower panel. For comparison, we also reproduce measurements of Q versus 3D separation (r) from Fig. 7 on the lower panel of this plot. The shaded regions surrounding each line depict the 2σ standard error in the mean. The region which is susceptible to crowding ($r_p < 11.8$ kpc) is depicted in the grey region on the left-hand side of each panel. We find that projection effects significantly dilute the sSFR enhancements that are seen in 3D.

Following the methodology that Patton et al. (2016) used for their observational sample, we compute the distance to each host galaxy’s closest projected companion (r_p) and second closest projected companion (r_2) while counting the number of companions within a projected separation of 2 Mpc (N_2). We again require potential companions to have stellar mass ratios of $\mu_c > 0.1$. Finally, we require potential companions to have a relative velocity along the line of sight (hereafter Δv) of less than 1000 km s^{-1} . This velocity restriction was used by Patton et al. (2016) to avoid obvious projected companions.

Having identified the closest companion for each host galaxy, we now narrow our analysis to host galaxies whose closest companion has $\Delta v < 300 \text{ km s}^{-1}$. This additional restriction on relative velocity increases the likelihood that a given companion is physically associated with the host galaxy (Patton et al. 2000), and corresponds to a maximum line of sight separation of 5 Mpc at $z = 0.5$. Our threshold of 300 km s^{-1} was also used by Patton et al. (2013, 2016), and is similar to the Δv restrictions in many other studies of galaxy pairs (Lin et al. 2004; Sol Alonso et al. 2006; Lambas et al. 2012; Ellison et al. 2013; Davies et al. 2015).

In Fig. 11, we show how the mean sSFR and its enhancement above the control samples depends on r_p in TNG300-1, here only considering the projection on to the $x - y$ plane. In the upper panel, we plot the mean sSFR of host galaxies and their controls versus the projected distance to the host galaxy’s closest companion. The observed trends are qualitatively similar to those seen in 3D (Fig. 7), with a steady increase in the mean sSFR of host galaxies with respect to their controls as the separation decreases. In the lower panel of Fig. 11, we plot the mean sSFR enhancement as a function of r_p . This shows that there are statistically significant en-

hancements (at the 2σ level) in the mean sSFR at $r_p < 260$ kpc, with enhancements rising to $Q \sim 1.7$ at the crowding limit of the sample (11.8 kpc).

In the lower panel of Fig. 11, we also plot the mean sSFR enhancement versus the 3D distance to the closest companion, plotting 3D and projected results on the same axis to facilitate a direct comparison between them. At a given separation, the enhancements are smaller in projected space than in 3D, implying that projection effects are diluting the enhancements that are detected. At small separations ($r_p \lesssim 50$ kpc), projection effects dilute the enhancements by about 20 per cent. At larger separations, the dilution increases to the point that enhancements are present in 3D but no longer detectable in projected space.

We can interpret these trends by noting that, at any given sSFR enhancement, $r_p < r$. This is to be expected, given that a galaxy pair’s projected separation must be smaller than its 3D separation. Moreover, in cases where the closest projected companion and the closest 3D companion are different galaxies, it is rarely the case that $r_p > r$ (for this to happen, the closest 3D companion must be excluded from consideration by failing the Δv criterion).

We also compare the $x - y$ projection results in Fig. 11 with projections in the $y - z$ and $x - z$ planes (not shown), and find that they agree within 1σ at all separations. This indicates that the TNG300-1 simulation is sufficiently large that the results have been averaged over a representative distribution of galaxy pair orientations, thereby mitigating the need to undertake a larger and/or more randomized set of projections.

In summary, we conclude that by measuring sSFR enhancements in projected space (as observers do), one is likely to underestimate the underlying level of sSFR enhancements by about 20 per cent at small separations, while also underestimating the radial extent of the enhancements. Nevertheless, the overall similarity of the trends seen in 3D and projected space clearly shows that the same underlying correlations produce these relationships.

5.2 Comparison with Patton et al. (2013)

As noted in Section 1, there is plenty of observational evidence for enhanced star formation in galaxy pairs, with the greatest enhancements typically seen in the pairs with the closest projected separations. Of particular note is the SDSS study of Patton et al. (2013), who use a similar methodology to ours. They find that the mean SFR is enhanced out to $r_p \sim 150$ kpc, with the enhancements reaching a factor of ~ 3 in the closest pairs. These findings are qualitatively similar to what we have found in TNG300-1 (see Fig. 11), although our enhancements are generally smaller and extend out to larger r_p .

However, there are a number of differences between the Patton et al. (2013) sample and the simulations used in this study. In particular, Patton et al. (2013) computed mean SFR using a low redshift ($z < 0.2$) flux-limited sample that is complete for host galaxies but not companions, whereas we have computed mean sSFR using a $z < 1$ stellar mass-limited sample which is complete for host galaxies and their companions. We therefore carry out a revised analysis of the Patton et al. (2013) SDSS sample, addressing most of these sample differences. We acknowledge at the outset that any such comparison will be imperfect due to some fundamental differences between the available observational data and the quantities reported in the simulations.

5.3 Comparison with SDSS

We begin with the SDSS sample of star-forming galaxies used by Patton et al. (2013), and provide here a brief summary of their data and methodology. The underlying sample from SDSS (York et al. 2000) is derived from SDSS Data Release 7 (Abazajian et al. 2009), using the main galaxy sample described by Strauss et al. (2002). Total stellar mass estimates are taken from Mendel et al. (2014) and are based on the photometry of Simard et al. (2011). The sample is restricted to the redshift range $0.02 < z < 0.2$.

In order to better compare the flux-limited SDSS sample of Patton et al. (2013) with the stellar mass-limited samples of galaxies in the simulations, we apply a minimum stellar mass of $10^{10} M_{\odot}$ to host galaxies in SDSS. This removes a substantial fraction of the host galaxies used in the analysis of Patton et al. (2013), especially at lower redshift. In addition, as dwarf galaxies in SDSS exhibit higher SFR enhancements than massive galaxies (Stierwalt et al. 2015), this restriction might be expected to reduce the overall enhancements in our sample.

Next, in order to maximize the stellar mass completeness of our sample at all redshifts, we impose an additional requirement that all galaxies lie above the ‘red sequence $+3\sigma$ ’ redshift-dependent stellar mass limit of Mendel et al. (2014). This ensures that galaxies of all colours are detectable throughout the sample. We then apply $V_{\max}$ weights to each galaxy, in order to correct for the resulting redshift-dependent mass limits of the sample. To be specific, we apply a statistical weight of w_v for each host galaxy, with

$$w_v = \frac{d(z_{\max})^3 - d(z_{\min})^3}{d(z_i)^3 - d(z_{\min})^3}, \quad (3)$$

where $d(z)$ is the co-moving line of sight distance at redshift z (Hogg 1999), $z_{\min}$ and $z_{\max}$ are the survey redshift limits, and z_i is the redshift at which the sample is complete for galaxies of stellar mass M_i . This weight yields the ratio of the co-moving volume at $z_{\min} < z < z_{\max}$ to the available co-moving volume for a host galaxy of stellar mass M_i .

Following Scudder et al. (2012), we restrict our sample to star-forming galaxies using the emission-line criteria of Kauffmann et al. (2003). We estimate the sSFR of each SDSS galaxy using the fibre SFRs of Brinchmann et al. (2004) and fibre stellar masses from Mendel et al. (2014). Given that the typical stellar mass covering fraction of the SDSS fibres for galaxies in our SDSS sample is approximately 30 per cent, fibre sSFRs are roughly analogous to the sSFRs computed within $R_{1/2}$ of the simulated galaxies used in this study (see Section 2.6). We additionally require galaxies to have $\text{sSFR} > 0.01 \text{ Gyr}^{-1}$, which removes some galaxies which are better described as passive than star-forming (Brinchmann et al. 2004; Bluck et al. 2014). We note that, while the host galaxies are matched in stellar mass with their control galaxies, averaging of sSFRs rather than SFRs will lead to noticeably different results, due to known correlations between SFR and stellar mass (Noeske et al. 2007; Peng et al. 2010).

We then apply the closest companion and control galaxy methodology of Patton et al. (2016) to this revised SDSS sample. This includes a weighting scheme that accounts for overall spectroscopic incompleteness as well as spectroscopic incompleteness that results from fibre collisions (Patton & Atfield 2008; Patton et al. 2011; Simard et al. 2011). We estimate the enhancement in each host galaxy’s sSFR by comparing with at least ten control galaxies, assigning higher statistical weights to the controls that provide the

Figure 12. sSFR enhancement is plotted versus projected separation for $z < 0.2$ galaxies from SDSS and TNG100-1. Both samples are restricted to star-forming galaxies with $\text{sSFR} > 0.01 \text{ Gyr}^{-1}$. The coloured shaded regions denote the 2σ standard error in the mean. The grey region on the left depicts the separations at which the TNG100-1 sample is susceptible to crowding ($r_p < 11.2 \text{ kpc}$).

best simultaneous matches in redshift, stellar mass,⁷ local density, and isolation. We also restrict our sample to host galaxies with closest companions that have $\Delta v < 300 \text{ km s}^{-1}$. This procedure yields a sample of 141 229 host galaxies from SDSS.

Finally, in order to match the IllustrisTNG simulation to the low redshift star-forming sample from SDSS, we apply a maximum redshift of 0.2 to the simulations, and we require a minimum sSFR of 0.01 Gyr^{-1} for host galaxies and their controls.

We plot sSFR enhancement as a function of projected separation for star-forming galaxies from SDSS and TNG100-1 in Fig. 12. We find general agreement between the observations and simulations. In both cases, significant sSFR enhancements are seen at $r_p \lesssim 150 \text{ kpc}$, with the sSFR enhancements increasing to $Q \sim 1.8$ at the smallest separations. The enhancements are somewhat larger in SDSS than in TNG100-1, but given their 2σ error bars, these differences are not statistically significant.

These results imply a general consistency between SDSS and TNG100-1. However, several caveats are in order. First, while our search for companions is complete down to 10 per cent of each host galaxy’s mass in IllustrisTNG, the same is not true for SDSS; in fact, some SDSS host galaxies lie near the $10^{10} M_{\odot}$ limit of the sample. While it is possible to address this incompleteness by imposing stricter stellar mass limits, this removes a substantial fraction of the host galaxies in our SDSS sample, yielding results that only poorly constrain the sSFR enhancements and their dependence on r_p . Secondly, the typical fibre covering fraction of galaxies in our SDSS sample is about 30 per cent, making them more central than the IllustrisTNG sSFRs that are computed within the stellar half-mass radius. Moreover, covering fraction depends on both redshift and galaxy size, leading us to average over a wide range of covering fractions. This means that there are likely to be meaningful differences in the spatial coverage of our SDSS and IllustrisTNG samples.

⁷Simultaneous matching on both redshift and stellar mass has the added benefit of ensuring that the fibre covering fraction distributions of hosts and controls will be similar to one another (Patton et al. 2011).

Nevertheless, the general agreement we have found between SDSS and TNG100-1 suggests that the simulations are successfully capturing star formation that is triggered by galaxy–galaxy interactions. Moreover, this agreement suggests that the approach we have used for detecting enhanced star formation in SDSS is successful at recovering the underlying dependence of sSFR enhancement on the true 3D separations of galaxies as reported earlier in this paper.

6 DISCUSSION

6.1 Are the sSFR enhancements caused by galaxy–galaxy interactions?

In Fig. 5, we discovered statistically significant enhancements in the mean sSFR of host galaxies extending out to 3D closest companion separations of 210 kpc in the highest resolution IllustrisTNG simulation (TNG100-1). The much larger sample in the TNG300-1 simulation allowed us to trace significant enhancements out even further, reaching 280 kpc (Fig. 7). At these large separations, we would not expect galaxies to be undergoing interactions that are strong enough to have any effect on their star formation rates. However, idealized merger simulations indicate that SFR enhancements may be present in galaxies that have recently experienced a close encounter and have subsequently moved to a relatively large orbital separation, as long as the SFR enhancement is relatively long lived (Patton et al. 2013).

In principle, we can test this scenario in IllustrisTNG by tracing out the orbits of interacting galaxies to see if the enhancements we detect at large r are in fact the result of star formation triggered by previous close encounters. However, this would necessitate tracking host galaxies and their companions over multiple snapshots. This is beyond the scope of this paper (which uses the properties of host galaxies and their companions at individual snapshots), but will be investigated in a forthcoming paper.

Here, we instead investigate whether host galaxies that are experiencing enhanced star formation have dark matter haloes that overlap with their closest companions. If so, this would confirm that their dark matter haloes are interacting, and it would increase the likelihood that these galaxy pairs have orbits that are decaying due to dynamical friction and destined to lead to mergers. Conversely, if there is little or no overlap between their dark matter haloes, it would be harder to attribute the enhanced star formation to interactions between these galaxies.

In order to assess the degree to which a host galaxy’s dark matter halo overlaps with its closest companion, we estimate the virial radius of each, and compare the sum of these radii with the separation between the two galaxies. We compute the virial radius of the host galaxy ($R_{\text{vir}}^{\text{host}}$) and its closest companion ($R_{\text{vir}}^{\text{comp}}$) using $R_{\text{vir}} = 120(M_{\text{tot}}/10^{11} M_{\odot})^{1/3}$ kpc (e.g. Dekel & Birnboim 2006), where M_{tot} is the total halo mass of the galaxy. We then compute the relative separation of their dark matter haloes (hereafter $r_{\text{sep}}^{\text{DM}}$) as follows:

$$r_{\text{sep}}^{\text{DM}} = \frac{r}{R_{\text{vir}}^{\text{host}} + R_{\text{vir}}^{\text{comp}}}. \quad (4)$$

Given this definition, a value of $r_{\text{sep}}^{\text{DM}} < 1$ means that the host galaxy and its companion have overlapping virial radii.⁸

⁸With overlapping haloes, SUBFIND has some trouble accurately deblending the dark matter distributions of the two galaxies. However, we mitigate this effect by summing the virial radii of both galaxies in the pair.

Figure 13. Enhanced sSFR is plotted versus the relative separation of the dark matter haloes ($r_{\text{sep}}^{\text{DM}}$) for TNG100-1 and TNG300-1. The shaded regions depict the 2σ standard error in the mean.

We plot the mean sSFR enhancement versus $r_{\text{sep}}^{\text{DM}}$ for the highest resolution (TNG100-1) and the largest volume (TNG300-1) IllustrisTNG simulations in Fig. 13. We find that the enhancements in mean sSFR are found at $r_{\text{sep}}^{\text{DM}} < 1$ for both TNG100-1 and TNG300-1, and nearly all of the enhancements occur at $r_{\text{sep}}^{\text{DM}} < 0.6$. This confirms that the sSFR enhancements we have reported, which occur for closest companions at $r < 280$ kpc (Fig. 7) or $r_p < 260$ kpc (Fig. 11), are likely due to systems which have substantially overlapping dark matter haloes.

This finding is consistent with the picture in which the sSFR enhancements are due to galaxy–galaxy interactions. The fact that we do not see net enhancements at wider separations suggests that post-encounter apocentre separations occur at $r_{\text{sep}}^{\text{DM}} < 1$ and/or that sSFR enhancements dissipate by the time the galaxies reach a separation of $r_{\text{sep}}^{\text{DM}} \sim 1$. However, without tracking sSFR enhancements as a function of orbital properties, we cannot distinguish between these scenarios, nor can we say for certain that the enhancements are triggered by encounters that are close enough to produce strong gravitational interactions between the galaxies.

6.2 What fraction of cosmic star formation is triggered by interacting galaxies?

We have detected clear enhancements in the mean sSFR of IllustrisTNG host galaxies out to relatively large closest companion separations. However, while these enhancements have high statistical significance, that alone does not mean that the enhancements are large enough to have a significant influence on galaxy populations as a whole. To address the larger meaning of these results, we need to consider the size of these enhancements in conjunction with the fraction of galaxies that are involved.

At small 3D separations, we find that the mean sSFR is enhanced by a factor of 2.0 ± 0.1 in the highest resolution IllustrisTNG simulation (TNG100-1; see Section 3.1). The level of enhancement is slightly lower at small projected separations (see Section 5.1). In both cases, these results should be considered lower limits, given that (i) we are unable to estimate the enhancements at even smaller separations (due to crowding) and (ii) the TNG100-1 simulations may not be of high enough resolution to have converged (see Section 4.1).

These levels of enhancement are broadly consistent with those found in simulations and in observations (e.g. Cox et al. 2008; Lambas et al. 2012; Patton et al. 2013; Knapen, Cisternas & Querejeta 2015; Martin et al. 2017; Rodríguez Montero et al. 2019). Given that some close pairs contain one or two passive galaxies (potential mixed mergers or dry mergers), while others may not be influencing one another (e.g. if they are approaching their first close encounter), some of these host galaxies are likely to have sSFR enhancements that are well in excess of the mean sSFR enhancement.

If we consider the full range of closest companion separations at which the mean sSFR of host galaxies is enhanced, we can gain a more complete picture of the importance of interactions on galaxy sSFRs. To this end, we identify the fraction of the host galaxy population that lies within the regime where the mean sSFR is found to be enhanced. We have previously found that the mean enhancement is statistically significant at $r < 280$ kpc in TNG300-1 (Section 3.3). We find that 37 per cent of host galaxies have $r < 280$ kpc. This places an upper limit on the fraction of the sample affected by these net enhancements.

More generally, we can examine the impact of enhanced sSFRs on the host sample as a whole by computing the mean sSFR of all host galaxies (regardless of the distance to their closest companion), and comparing with the mean sSFR of their best controls. This is equivalent to replacing every host galaxy with its best control and assessing how that would change the mean sSFR of the sample. For TNG100-1, we find that the mean sSFR of host galaxies is 0.113 Gyr^{-1} , versus 0.098 Gyr^{-1} for their best controls. This suggests that the presence of closest companions (regardless of separation) boosts the mean sSFR of host galaxies by 14.5 per cent. Approximately one third of this boost (4.7 per cent) comes from host galaxies with close companions ($r < 30$ kpc),⁹ with the remaining two thirds coming almost entirely from host galaxies with closest companions at $30 < r < 200$ kpc. In other words, while only a small fraction of galaxies have a close companion at a given epoch, the combined influence of closest companions at all separations has a notable influence on the IllustrisTNG galaxy population as a whole. Given the consistency between enhancements in IllustrisTNG and SDSS, it is reasonable to infer that the same conclusions may apply to observed galaxy populations at low redshift.

These results are broadly consistent with the findings of Martin et al. (2017), who use the Horizon-AGN cosmological hydrodynamical simulations (Dubois et al. 2014) to estimate the SFR enhancement within a 2 Gyr window surrounding mergers, thereby including both interacting pairs and post-mergers. They report that about 20 per cent of star formation at $z \sim 1$ is triggered by mergers, with 65 per cent of the enhancement occurring before the merger. Conversely, Rodríguez Montero et al. (2019) find that mergers in the SIMBA simulations (Davé et al. 2019) contribute only about 1 per cent to the global SFR budget at $z \sim 2$, with this fraction dropping to about 0.5 per cent at $0 < z < 1$. However, their census refers only to mergers seen shortly after coalescence, and does not include star formation that is triggered in the pre-merger phase.

7 CONCLUSIONS

We have used the IllustrisTNG cosmological hydrodynamical simulations to investigate the relationship between enhanced star

formation and the presence of close companions. The primary goal of this analysis is to bridge the gap between idealized binary merger simulations (which predict enhanced star formation during galaxy–galaxy interactions) and observational studies of galaxy pairs (which report enhanced star formation that is associated with the presence of close companions). We identify the closest companion for massive galaxies ($M_* > 10^{10} M_\odot$) in various IllustrisTNG simulations, taking care to minimize the effects of numerical stripping on the estimated stellar masses and sSFRs of galaxies with close companions. We estimate the enhancement in galaxy sSFRs by comparing with a control sample, using the methodology of Patton et al. (2016) to match on redshift, stellar mass, local density, and isolation. We then analyse how the sSFR enhancement depends on the 3D and projected distance to the closest companion, also comparing our results with the Illustris-1 and EAGLE simulations, and with a sample of massive galaxies from the SDSS.

Our main conclusions are as follows:

- (i) Using the highest resolution IllustrisTNG simulation (TNG100-1; $L_{\text{box}} \sim 110 \text{ Mpc}$), we find that the mean sSFR is enhanced by a factor of 2.0 ± 0.1 at small separations ($r \sim 16$ kpc; see Fig. 5), with clear morphological signs of interactions seen in synthetic stellar composite images of most of these systems (Fig. 6).
- (ii) Using the largest volume IllustrisTNG simulation (TNG300-1; $L_{\text{box}} \sim 300 \text{ Mpc}$), we detect statistically significant enhancements in the mean sSFR out to 3D separations of 280 kpc (Fig. 7).
- (iii) Clear enhancements in the mean sSFR are seen throughout the redshift range $0 \leq z < 1$ for TNG300-1, although there is a gradual decrease in small scale ($r < 50$ kpc) sSFR enhancements as redshift increases (Fig. 8).
- (iv) Similar sSFR enhancements are detected across the full range of TNG100 resolutions, although the highest resolution simulation (TNG100-1) exhibits larger enhancements below 50 kpc (Fig. 9).
- (v) We also detect significant sSFR enhancements in the Illustris-1 and EAGLE simulations, although the enhancements are smaller (by about a factor of two) and extend out to smaller separations than in TNG100-1 (Fig. 10).
- (vi) After redefining our closest companion sample using projected separation, we find that projection effects dilute the TNG300-1 sSFR enhancements by about 20 per cent for close companions, and they narrow the range of separations at which enhancements are detected (Fig. 11).
- (vii) Using low redshift star forming samples of galaxies from TNG100-1 and SDSS, we find general agreement in sSFR enhancements between simulations and observations (Fig. 12).
- (viii) The sSFR enhancements in IllustrisTNG occur in systems where the host galaxy and its closest companion have overlapping virial radii, confirming that these enhancements are associated with interacting galaxies (Fig. 13).
- (ix) By summing the sSFR enhancements at all separations, we estimate that closest companions boost the mean sSFR of massive galaxies in TNG100-1 by 14.5 per cent.

We have shown that cosmological hydrodynamical simulations are able to capture enhanced star formation associated with galaxy–galaxy interactions, exhibiting general consistency with high resolution merger simulations and observations of galaxy pairs. In the second paper in this series, Hani et al. (2020) report that star formation is also enhanced in IllustrisTNG post-merger galaxies. These findings provide motivation for using IllustrisTNG to assess the influence of interactions on additional galaxy properties, such as metallicities, morphologies, and gas content. Moreover, the ability to track the orbital history and future of interacting galaxies in these

⁹Patton et al. (2013) reached a similar conclusion when splitting their SDSS pair sample at a projected separation of 30 kpc.

simulations will enable us to develop a deeper understanding of the physical processes that give rise to these changes in galaxy properties.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a thoughtful and constructive referee's report. We thank all members of the IllustrisTNG, Illustris, EAGLE, and SDSS collaborations for making their data available. DRP and SLE gratefully acknowledge NSERC for Discovery Grants which helped to fund this research. CM, KW, and WB were supported by NSERC Undergraduate Student Research Awards. Support for JM was provided by the NSF (AST Award Number 1516374), and by the Harvard Institute for Theory and Computation, through their Visiting Scholars Program. MHH acknowledges the receipt of a Vanier Canada Graduate Scholarship. We thank Connor Bottrell for helpful discussions which improved this manuscript.

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

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

REFERENCES

Abazajian K. N. et al., 2009, *ApJS*, 182, 543
 Alonso M. S., Lambas D. G., Tissera P., Coldwell G., 2007, *MNRAS*, 375, 1017
 An S.-H., Kim J., Moon J.-S., Yoon S.-J., 2019, *ApJ*, 887, 59
 Barrera-Ballesteros J. K. et al., 2015, *A&A*, 579, A45
 Barton E. J., Geller M. J., Kenyon S. J., 2000, *ApJ*, 530, 660
 Barton E. J., Arnold J. A., Zentner A. R., Bullock J. S., Wechsler R. H., 2007, *ApJ*, 671, 1538
 Bluck A. F. L., Mendel J. T., Ellison S. L., Moreno J., Simard L., Patton D. R., Starkenburg E., 2014, *MNRAS*, 441, 599
 Blumenthal K. A., Barnes J. E., 2018, *MNRAS*, 479, 3952
 Bottrell C., Torrey P., Simard L., Ellison S. L., 2017, *MNRAS*, 467, 2879
 Bournaud F. et al., 2011, *ApJ*, 730, 4
 Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, 351, 1151
 Bustamante S., Sparre M., Springel V., Grand R. J. J., 2018, *MNRAS*, 479, 3381

Bustamante S., Ellison S. L., Patton D. R., Sparre M., 2020, preprint ([arXiv:2004.06121](https://arxiv.org/abs/2004.06121))
 Cao C. et al., 2016, *ApJS*, 222, 16
 Casteels K. R. V. et al., 2014, *MNRAS*, 445, 1157
 Chabrier G., 2003, *PASP*, 115, 763
 Chilingarian I. V., Di Matteo P., Combes F., Melchior A.-L., Semelin B., 2010, *A&A*, 518, A61
 Chown R. et al., 2019, *MNRAS*, 484, 5192
 Cox T. J., Jonsson P., Primack J. R., Somerville R. S., 2006, *MNRAS*, 373, 1013
 Cox T. J., Jonsson P., Somerville R. S., Primack J. R., Dekel A., 2008, *MNRAS*, 384, 386
 Crain R. A. et al., 2015, *MNRAS*, 450, 1937
 Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
 Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
 Davies L. J. M. et al., 2015, *MNRAS*, 452, 616
 De Propriis R., Conselice C. J., Liske J., Driver S. P., Patton D. R., Graham A. W., Allen P. D., 2007, *ApJ*, 666, 212
 Dekel A., Birnboim Y., 2006, *MNRAS*, 368, 2
 Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
 Di Matteo P., Combes F., Melchior A.-L., Semelin B., 2007, *A&A*, 468, 61
 Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
 Donnari M. et al., 2019, *MNRAS*, 485, 4817
 Dubois Y. et al., 2014, *MNRAS*, 444, 1453
 Ellison S. L., Patton D. R., Simard L., McConnachie A. W., 2008, *AJ*, 135, 1877
 Ellison S. L., Patton D. R., Simard L., McConnachie A. W., Baldry I. K., Mendel J. T., 2010, *MNRAS*, 407, 1514
 Ellison S. L., Patton D. R., Mendel J. T., Scudder J. M., 2011, *MNRAS*, 418, 2043
 Ellison S. L., Mendel J. T., Patton D. R., Scudder J. M., 2013, *MNRAS*, 435, 3627
 Ellison S. L., Viswanathan A., Patton D. R., Bottrell C., McConnachie A. W., Gwyn S., Cuillandre J.-C., 2019, *MNRAS*, 487, 2491
 Fensch J. et al., 2017, *MNRAS*, 465, 1934
 Furlong M. et al., 2017, *MNRAS*, 465, 722
 Genel S. et al., 2014, *MNRAS*, 445, 175
 Genel S. et al., 2018, *MNRAS*, 474, 3976
 Hani M. H., Gosain H., Ellison S. L., Patton D. R., Torrey P., 2020, *MNRAS*, 493, 3716
 Hernández-Toledo H. M., Avila-Reese V., Conselice C. J., Puerari I., 2005, *AJ*, 129, 682
 Hogg D. W., 1999, preprint ([astro-ph/9905116](https://arxiv.org/abs/astro-ph/9905116))
 Hopkins P. F., Cox T. J., Hernquist L., Narayanan D., Hayward C. C., Murray N., 2013, *MNRAS*, 430, 1901
 Jian H.-Y., Lin L., Chiueh T., 2012, *ApJ*, 754, 26
 Kauffmann G. et al., 2003, *MNRAS*, 346, 1055
 Kaviraj S. et al., 2013, *MNRAS*, 429, L40
 Kaviraj S., Devriendt J., Dubois Y., Slyz A., Welker C., Pichon C., Peirani S., Le Borgne D., 2015, *MNRAS*, 452, 2845
 Kewley L. J., Rupke D., Zahid H. J., Geller M. J., Barton E. J., 2010, *ApJ*, 721, L48
 Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M.-S., 2015, *MNRAS*, 450, 1349
 Kitzbichler M. G., White S. D. M., 2008, *MNRAS*, 391, 1489
 Knapen J. H., Cisternas M., Querejeta M., 2015, *MNRAS*, 454, 1742
 Lambas D. G., Alonso S., Mesa V., O'Mill A. L., 2012, *A&A*, 539, A45
 Lin L. et al., 2004, *ApJ*, 617, L9
 Lin L. et al., 2010, *ApJ*, 718, 1158
 Lofthouse E. K., Kaviraj S., Conselice C. J., Mortlock A., Hartley W., 2017, *MNRAS*, 465, 2895
 Lotz J. M., Jonsson P., Cox T. J., Primack J. R., 2008, *MNRAS*, 391, 1137
 Marinacci F. et al., 2018, *MNRAS*, 480, 5113
 Martin G., Kaviraj S., Devriendt J. E. G., Dubois Y., Laigle C., Pichon C., 2017, *MNRAS*, 472, L50

- McAlpine S. et al., 2016, *A&C*, 15, 72
- McIntosh D. H., Guo Y., Hertzberg J., Katz N., Mo H. J., van den Bosch F. C., Yang X., 2008, *MNRAS*, 388, 1537
- Mendel J. T., Simard L., Palmer M., Ellison S. L., Patton D. R., 2014, *ApJS*, 210, 3
- Mihos J. C., Hernquist L., 1994, *ApJ*, 425, L13
- Mihos J. C., Hernquist L., 1996, *ApJ*, 464, 641
- Moreno J., Bluck A. F. L., Ellison S. L., Patton D. R., Torrey P., Moster B. P., 2013, *MNRAS*, 436, 1765
- Moreno J., Torrey P., Ellison S. L., Patton D. R., Bluck A. F. L., Bansal G., Hernquist L., 2015, *MNRAS*, 448, 1107
- Moreno J. et al., 2019, *MNRAS*, 485, 1320
- Naiman J. P. et al., 2018, *MNRAS*, 477, 1206
- Nelson D. et al., 2015, *Astron. Comput.*, 13, 12
- Nelson D. et al., 2018, *MNRAS*, 475, 624
- Nelson D. et al., 2019, *Comput. Astrophys. Cosmol.*, 6, 2
- Noeske K. G. et al., 2007, *ApJ*, 660, L43
- Pan H.-A. et al., 2019, *ApJ*, 881, 119
- Patton D. R., Atfield J. E., 2008, *ApJ*, 685, 235
- Patton D. R., Carlberg R. G., Marzke R. O., Pritchett C. J., da Costa L. N., Pellegrini P. S., 2000, *ApJ*, 536, 153
- Patton D. R., Grant J. K., Simard L., Pritchett C. J., Carlberg R. G., Borne K. D., 2005, *AJ*, 130, 2043
- Patton D. R., Ellison S. L., Simard L., McConnachie A. W., Mendel J. T., 2011, *MNRAS*, 412, 591
- Patton D. R., Torrey P., Ellison S. L., Mendel J. T., Scudder J. M., 2013, *MNRAS*, 433, L59
- Patton D. R., Qamar F. D., Ellison S. L., Bluck A. F. L., Simard L., Mendel J. T., Moreno J., Torrey P., 2016, *MNRAS*, 461, 2589
- Pearson W. J., Wang L., Trayford J. W., Petrillo C. E., van der Tak F. F. S., 2019, *A&A*, 626, A49
- Peng Y. et al., 2010, *ApJ*, 721, 193
- Perret V. et al., 2014, *A&A*, 562, A1
- Pillepich A. et al., 2018a, *MNRAS*, 475, 648
- Pillepich A. et al., 2018b, *MNRAS*, 473, 4077
- Renaud F., Bournaud F., Kraljic K., Duc P.-A., 2014, *MNRAS*, 442, L33
- Robotham A. S. G. et al., 2014, *MNRAS*, 444, 3986
- Rodríguez Montero F., Davé R., Wild V., Anglés-Alcázar D., Narayanan D., 2019, *MNRAS*, 490, 2139
- Rodríguez-Gomez V. et al., 2015, *MNRAS*, 449, 49
- Rodríguez-Gomez V. et al., 2019, *MNRAS*, 483, 4140
- Rupke D. S. N., Kewley L. J., Barnes J. E., 2010, *ApJ*, 710, L156
- Rupke D. S. N., Kewley L. J., Chien L.-H., 2010, *ApJ*, 723, 1255
- Satyapal S., Ellison S. L., McAlpine W., Hickox R. C., Patton D. R., Mendel J. T., 2014, *MNRAS*, 441, 1297
- Schaye J., 2004, *ApJ*, 609, 667
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, 383, 1210
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Scudder J. M., Ellison S. L., Torrey P., Patton D. R., Mendel J. T., 2012, *MNRAS*, 426, 549
- Simard L., Mendel J. T., Patton D. R., Ellison S. L., McConnachie A. W., 2011, *ApJS*, 196, 11
- Snyder G. F. et al., 2015, *MNRAS*, 454, 1886
- Sol Alonso M., Lambas D. G., Tissera P., Coldwell G., 2006, *MNRAS*, 367, 1029
- Sparre M. et al., 2015, *MNRAS*, 447, 3548
- Sparre M., Springel V., 2016, *MNRAS*, 462, 2418
- Speagle J. S., Steinhardt C. L., Capak P. L., Silverman J. D., 2014, *ApJS*, 214, 15
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2018, *MNRAS*, 475, 676
- Stierwalt S. et al., 2015, *ApJ*, 805, 2
- Strauss M. A. et al., 2002, *AJ*, 124, 1810
- Thorp M. D., Ellison S. L., Simard L., Sánchez S. F., Antonio B., 2019, *MNRAS*, 482, L55
- Torrey P. et al., 2015, *MNRAS*, 447, 2753
- Torrey P. et al., 2018, *MNRAS*, 477, L16
- Torrey P., Cox T. J., Kewley L., Hernquist L., 2012, *ApJ*, 746, 108
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M. et al., 2014a, *Nature*, 509, 177
- Vogelsberger M. et al., 2014b, *MNRAS*, 444, 1518
- Weinberger R. et al., 2018, *MNRAS*, 479, 4056
- Weston M. E., McIntosh D. H., Brodwin M., Mann J., Cooper A., McConnell A., Nielsen J. L., 2017, *MNRAS*, 464, 3882
- Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, *ApJ*, 754, L29
- Woods D. F., Geller M. J., Kurtz M. J., Westra E., Fabricant D. G., Dell'Antonio I., 2010, *AJ*, 139, 1857
- York D. G. et al., 2000, *AJ*, 120, 1579

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.