



Cosmology with One Galaxy?

Francisco Villaescusa-Navarro^{1,2}, Jupiter Ding², Shy Genel^{1,3}, Stephanie Tonnesen¹, Valentina La Torre⁴, David N. Spergel^{1,2}, Romain Teyssier², Yin Li^{1,5}, Caroline Heneka⁶, Pablo Lemos^{7,8}, Daniel Anglés-Alcázar^{1,9}, Daisuke Nagai¹⁰, and Mark Vogelsberger¹¹

¹ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA; fvillaescusa@flatironinstitute.org

² Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

³ Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

⁴ Department of Physics and Astronomy, Tufts University, Medford, MA 02155, USA

⁵ Center for Computational Mathematics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

⁶ University of Hamburg, Hamburg Observatory, Gojenbergsweg 112, D-21029 Hamburg, Germany

⁷ Department of Physics and Astronomy, University of Sussex, Brighton, BN1 9QH, UK

⁸ University College London, Gower Street, London, UK

⁹ Department of Physics, University of Connecticut, 196 Auditorium Road, Storrs, CT 06269, USA

¹⁰ Department of Physics, Yale University, New Haven, CT 06520, USA

¹¹ Kavli Institute for Astrophysics and Space Research, Department of Physics, MIT, Cambridge, MA 02139, USA

Received 2022 January 14; revised 2022 March 3; accepted 2022 March 11; published 2022 April 20

Abstract

Galaxies can be characterized by many internal properties such as stellar mass, gas metallicity, and star formation rate. We quantify the amount of cosmological and astrophysical information that the internal properties of individual galaxies and their host dark matter halos contain. We train neural networks using hundreds of thousands of galaxies from 2000 state-of-the-art hydrodynamic simulations with different cosmologies and astrophysical models of the CAMELS project to perform likelihood-free inference on the value of the cosmological and astrophysical parameters. We find that knowing the internal properties of a single galaxy allows our models to infer the value of Ω_m , at fixed Ω_b , with a $\sim 10\%$ precision, while no constraint can be placed on σ_8 . Our results hold for any type of galaxy, central or satellite, massive or dwarf, at all considered redshifts, $z \leq 3$, and they incorporate uncertainties in astrophysics as modeled in CAMELS. However, our models are not robust to changes in subgrid physics due to the large intrinsic differences the two considered models imprint on galaxy properties. We find that the stellar mass, stellar metallicity, and maximum circular velocity are among the most important galaxy properties to determine the value of Ω_m . We believe that our results can be explained by considering that changes in the value of Ω_m , or potentially Ω_b/Ω_m , affect the dark matter content of galaxies, which leaves a signature in galaxy properties distinct from the one induced by galactic processes. Our results suggest that the low-dimensional manifold hosting galaxy properties provides a tight direct link between cosmology and astrophysics.

Unified Astronomy Thesaurus concepts: Galaxy formation (595); Cosmological models (337); Astrostatistics (1882); Hydrodynamical simulations (767)

1. Introduction

The discovery that the universe is accelerating its expansion (Riess et al. 1998; Perlmutter et al. 1999) marked an inflection point in cosmology. Determining the nature and properties of the substance responsible for this behavior, dark energy, is one of the most important goals of current cosmology.

In order to accomplish this task we need to extract the maximum information from cosmological surveys. We know that for Gaussian density fields the power spectrum (or the correlation function) is the optimal estimator to extract the maximum available information. However, we do not know what estimator would allow us to extract the maximum amount of information for non-Gaussian density fields, such as the matter and galaxy distribution on nonlinear scales. Quantifying the information content from different estimators is currently a very active area of research (Coulton et al. 2019; Liu & Madhavacheril 2019; Li et al. 2019; Marques et al. 2019; Ajani et al. 2020; Allys et al. 2020; Banerjee et al. 2020;

Dai et al. 2020; de la Bella et al. 2021; Friedrich et al. 2020; Giri & Smith 2020; Hahn et al. 2020; Lee & Ryu 2020; Ryu & Lee 2020; Villaescusa-Navarro et al. 2020; Uhlemann et al. 2020; Zhang et al. 2020; Banerjee & Abel 2021a, 2021b; Bayer et al. 2021a, 2021b; Cheng and Ménard 2021; Gualdi et al. 2021a, 2021b; Hahn & Villaescusa-Navarro 2021; Harnois-Déraps et al. 2021; Kuruvilla 2021; Kuruvilla & Aghanim 2021; Massara et al. 2021; Naidoo et al. 2021; Porth et al. 2021; Samushia et al. 2021; Valogiannis & Dvorkin 2021; Harnois-Déraps et al. 2022).

Another avenue is to use machine-learning techniques, e.g., neural networks, to find an approximation to the optimal estimator (Schmelzle et al. 2017; Ravanbakhsh et al. 2017; Gupta et al. 2018; Fluri et al. 2019; Ntampaka et al. 2020; Ribli et al. 2019; Hassan et al. 2020; Jeffrey et al. 2021; Zorrilla Matilla et al. 2020; Lu et al. 2022; Villaescusa-Navarro et al. 2021c). Recent works have shown that even for fields that are very contaminated by astrophysical effects, it is possible to extract cosmological information from small scales (Villaescusa-Navarro et al. 2021a).

Either way, both approaches yield the same conclusion: small, nonlinear scales seem to contain a wealth of cosmological information. Down to which scale is this statement true?



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Does the information run out at some point? These are difficult questions that we do not attempt to address in this work. Instead, we investigate whether there is any cosmological information in one of the fundamental building blocks of many cosmological surveys: galaxies. In other words, can we infer the value of the cosmological parameters from a single, generic galaxy?

To address this question we employ machine-learning methods to connect the internal properties of individual galaxies to the value of the cosmological and astrophysical parameters. We made use of galaxies from the state-of-the-art hydrodynamic simulations of the CAMELS project (Villaescusa-Navarro et al. 2021b). The internal galaxy properties considered in this work include the stellar mass, the star formation rate, the total mass in the galaxy’s subhalo, and the stellar radius among others. The CAMELS simulations contain around 1 million galaxies at fixed redshift for 2000 different cosmological and astrophysical models from two distinct suites of hydrodynamic simulations based on the IllustrisTNG (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018a; Springel et al. 2018) and SIMBA (Davé et al. 2019) frameworks. This allows us to quantify the dependence of our results on uncertainties in astrophysics processes and on the subgrid physics model.

As we shall see below, we find that we can infer the value of Ω_m with a $\simeq 10\%$ precision just using the internal properties of an individual, generic galaxy; no constraint can be placed on the value of σ_8 . These constraints account for uncertainties in astrophysics, as implemented in each of the suites comprising the CAMELS simulations. However, they are not robust to changes in subgrid physics, due to the intrinsic differences in galaxy properties between the different frameworks. If our interpretation of these results is correct, it would imply that galaxy properties live in manifolds that change with the value of Ω_m , providing a link between cosmology and astrophysics. To our knowledge, this is the first time that this idea has been explored, and it may open new exciting possibilities to connect cosmology with astrophysics through galaxy properties.

To enable the community to reproduce our results, we release all data used in this work together with the codes, databases, and network weights obtained after training. We refer the reader to <https://github.com/franciscovillaescusa/Cosmo1gal> for further details.

This paper is organized as follows. In Section 2 we describe the data we use together with the machine-learning methods employed. We present our results in Section 3 and attempt a physical interpretation for them in Section 4. Finally, we summarize and discuss the main results of this work in Section 5.

2. Methods

In this section we describe the data and the machine-learning models we use to find the mapping between the properties of individual galaxies and the value of the cosmological and astrophysical parameters.

2.1. Simulations

In this work we use galaxies from the simulations of the CAMELS project (Villaescusa-Navarro et al. 2021b). CAMELS contains two suites of state-of-the-art hydrodynamic simulations:

1. *IllustrisTNG*. The simulations in this suite have been run with the AREPO code (Springel 2010; Weinberger et al. 2020) and employ the same subgrid physics model as the original IllustrisTNG simulations (Weinberger et al. 2017; Pillepich et al. 2018b).
2. *SIMBA*. The simulations in this suite have been run with the GIZMO code (Hopkins 2015) and employ the same subgrid physics model as the original SIMBA simulation (Davé et al. 2019), building on its precursor MUFASA (Davé et al. 2016) with the addition of supermassive black hole growth and feedback (Anglés-Alcázar et al. 2017).

All simulations follow the evolution of 2×256^3 dark matter plus fluid elements in a periodic comoving volume of $(25 h^{-1} \text{Mpc})^3$ from $z = 127$ down to $z = 0$. All simulations share the value of the following cosmological parameters: $\Omega_b = 0.049$, $h = 0.6711$, $n_s = 0.9624$, $\sum m_\nu = 0.0 \text{ eV}$, $w = -1$. However, each simulation has a unique value of Ω_m and σ_8 . The hydrodynamic simulations also vary the values of four astrophysical parameters that control the efficiency of supernova and active galactic nucleus (AGN) feedback: A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2} .

In this work we use the LH sets of the IllustrisTNG and SIMBA suites. Each set contains 1000 simulations, where the values of Ω_m , σ_8 , A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2} are arranged in a Latin hypercube defined by

$$\Omega_m \in [0.1, 0.5] \quad (1)$$

$$\sigma_8 \in [0.6, 1.0] \quad (2)$$

$$A_{\text{SN1}}, A_{\text{AGN1}} \in [0.25, 4.0] \quad (3)$$

$$A_{\text{SN2}}, A_{\text{AGN2}} \in [0.5, 2.0], \quad (4)$$

and each simulation has a different value of the random seed that sets the initial conditions. We note that the Latin hypercubes of the IllustrisTNG and SIMBA simulations are different, i.e., there is no correspondence between simulations from the two sets. We emphasize that the astrophysics parameters have different meanings in the IllustrisTNG versus SIMBA suites due to the distinct internal workings of these two models.

We refer the reader to Villaescusa-Navarro et al. (2021b) for further details on the simulations of the CAMELS project.

2.2. Galaxy Properties

We have run SUBFIND (Springel et al. 2001) to identify halos and subhalos in the simulations. In this work we define galaxies as subhalos that contain more than 20 star particles. All galaxies from all simulations are characterized by 14 different properties:

1. M_g . The gas mass content of the subhalo hosting the galaxy, including the contribution from the circumgalactic medium.
2. M_{BH} . The total mass of black holes in the galaxy.
3. M_* . The stellar mass of the galaxy.
4. M_t . The total mass of the subhalo hosting the galaxy, i.e., the sum of the mass in dark matter, gas, stars, and black holes in the subhalo.
5. V_{max} . The maximum circular velocity of the subhalo hosting the galaxy: $V_{\text{max}} = \max(\sqrt{GM(<R)/R})$.

6. σ_v . The mass-weighted velocity dispersion of all particles contained in the galaxy's subhalo.
7. Z_g . The mass-weighted gas metallicity of the galaxy.
8. Z_* . The mass-weighted stellar metallicity of the galaxy.
9. SFR . The galaxy star formation rate.
10. J . The modulus of the galaxy's subhalo spin vector.
11. V . The modulus of the galaxy's subhalo peculiar velocity.
12. R_* . The radius containing half of the galaxy stellar mass.
13. R_t . The radius containing half of the total mass of the galaxy's subhalo.
14. R_{\max} . The radius at which $\sqrt{GM(<R)/R} = V_{\max}$.

For galaxies of the IllustrisTNG simulations we also consider three additional properties:

15. U . The galaxy magnitude in the U band.
16. K . The galaxy magnitude in the K band.
17. g . The galaxy magnitude in the g band.

These three properties are only present in the IllustrisTNG simulations because they are calculated on the fly and saved to the simulation output files by the AREPO code but not by the GIZMO code.

We have employed the above galaxy properties because these are computed by SUBFIND and therefore easily accessible to us. Using other properties would require postprocessing the snapshots, which we leave for future work. We emphasize that while most of the considered properties can be associated with galaxies themselves, there are others that should be seen as properties of the subhalos hosting the galaxies, like V_{\max} , M_t , and σ_v . We note that in this work we are not splitting galaxies according to some property, e.g., large or small, central or satellite.

2.3. Machine-learning Algorithms

We made use of several machine-learning algorithms for three different reasons: first, to assert that our conclusions hold independently of the method used; second, because some tasks (e.g., feature ranking) require a significant amount of computation, and it is difficult to perform them if not using fast methods; and third, to provide some interpretability to our results.

1. *Gradient Boosting Trees*: This method is based on decision trees and therefore computationally efficient. We made use of the XGB package¹² to estimate the value of Ω_m from the 17 galaxy properties. For each task, we tune the value of the following hyperparameters: (1) the learning rate, (2) the maximum depth, (3) the minimum child weight, (4) the value of gamma, (5) the value of `colsample_bytree`, and (6) the number of estimators. The loss function we optimize is the mean squared error. XGB accounts for L2 regularization internally. We note that in this case we perform parameter regression, while with neural networks we do likelihood-free inference.
2. *Neural networks*: We made use of fully connected layers because they are appropriate for the task we consider in this work. Our architecture consists of several fully connected layer blocks. These blocks contain a fully connected layer that is followed by a LeakyReLU activation layer with a slope of 0.2 and a dropout layer where the value of the dropout rate is a hyperparameter.

The very last layer of the architecture is just a fully connected layer not followed by an activation or dropout layer. The hyperparameters we consider are (1) the number of fully connected layers, (2) the number of neurons in each layer, (3) the dropout value, (4) the value of the weight decay, and (5) the learning rate. Our networks are trained to perform likelihood-free inference; they estimate the posterior mean and standard deviation for each parameter by minimizing the loss function of moment networks (Jeffrey & Wandelt 2020).

We use the optuna¹³ (Akiba et al. 2019) package to perform the hyperparameter optimization of both gradient boosting trees and neural networks. In both cases, we first sample the hyperparameter space using between 25 and 30 trials¹⁴ and then we perform Bayesian optimization for 75–80 more trials. In all cases we search the hyperparameter space to minimize the value of the validation loss. The value of the hyperparameters for the best models we found can be obtained in <https://github.com/franciscovillaescusa/Cosmo1gal> together with the network weights and scripts used.

For gradient boosting trees and neural networks we split the data into three different sets: training, validation, and testing. Because galaxies in the same simulations may share features in either low- or high-dimensional spaces, we first split the data by simulation. The training set contains 850 simulations with all their galaxies. The validation set has all galaxies from 100 simulations, while the testing set contains 50 simulations with all their galaxies. By splitting the data in this way, we can guarantee that the galaxies in the test set, together with their associated cosmology and astrophysics, have never been seen by the model before. We note that according to our galaxy selection criteria (i.e., subhalos with more than 20 star particles) the number of galaxies will differ in the different simulations. For instance, models with high values of Ω_m and/or σ_8 will typically contain more galaxies than simulations with lower values of these parameters. This effect could potentially introduce a bias because the values of the parameters are effectively not sampled uniformly. However, because the differences in the number of galaxies between the different simulations are relatively small, we think this effect is not very important. On the other hand, correcting this effect may lead to slight improvements in our results.

2.4. Accuracy and Precision

Throughout this paper we will be discussing the accuracy and the precision of a given model. Here we describe what we mean by these.

Neural networks. With this method, we perform likelihood-free inference, and the output of the networks is the posterior mean (μ) and standard deviation (σ) of a given parameter i , i.e.,

$$\mu_i(\mathbf{X}) = \int_{\theta_i} p_i(\theta_i|\mathbf{X})\theta_i d\theta_i, \quad (5)$$

$$\sigma_i(\mathbf{X}) = \int_{\theta_i} p_i(\theta_i|\mathbf{X})(\theta_i - \mu_i)^2 d\theta_i, \quad (6)$$

¹³ <https://optuna.org>

¹⁴ A trial represents the result of training the model with a given value of the hyperparameters.

¹² <https://xgboost.readthedocs.io>

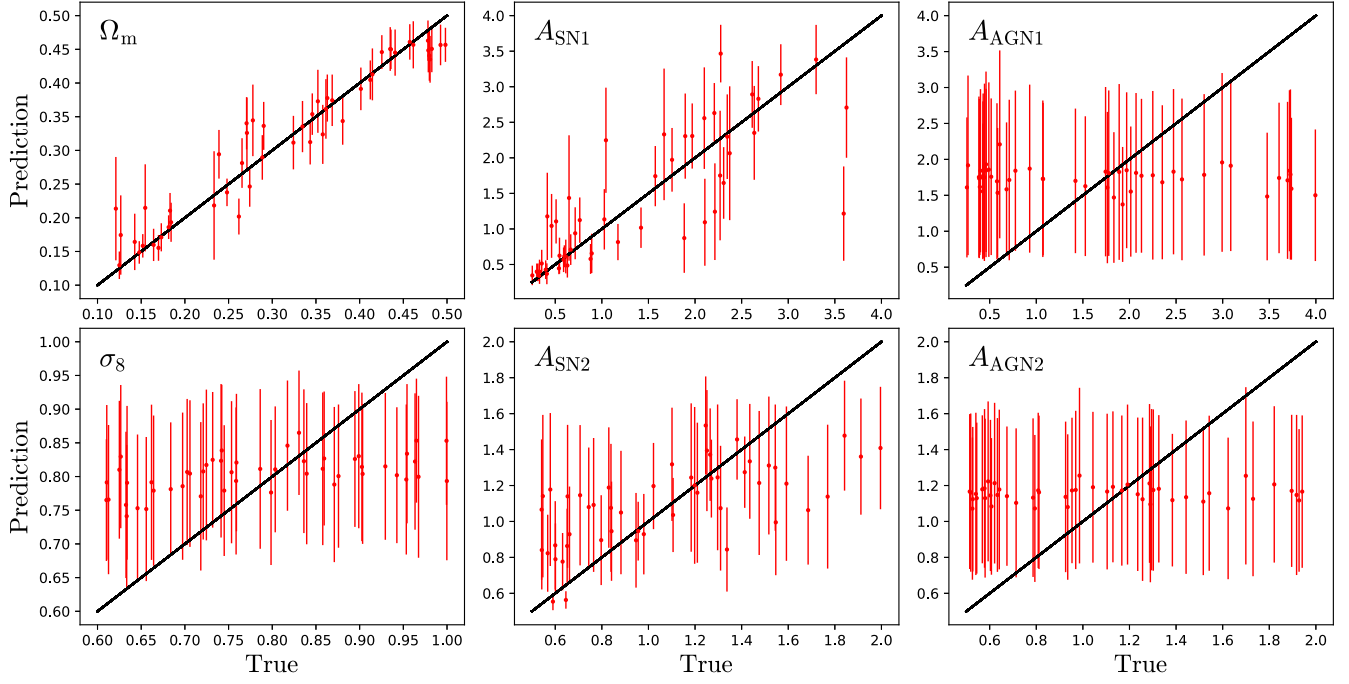


Figure 1. We have trained a neural network to perform likelihood-free inference on the value of the cosmological (Ω_m and σ_8) and astrophysical (A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2}) parameters using as input 17 properties of individual galaxies from the IllustrisTNG simulations at $z=0$. Once the network is trained, we test it using individual galaxies from the test set. The different panels show the posterior mean and standard deviation predicted by the network vs. the true value. Every point with its error bar represents a single galaxy. We find that our model is able to infer the value of Ω_m from the properties of individual galaxies with a $\sim 10\%$ precision.

where \mathbf{X} is the vector containing the galaxy properties and $p_i(\theta_i|\mathbf{X})$ is the marginal posterior over the parameter i ,

$$p_i(\theta_i|\mathbf{X}) = \int_{\theta} p_i(\theta_i|\mathbf{X}) d\theta_1 \dots d\theta_{i-1} d\theta_{i+1} \dots d\theta_n. \quad (7)$$

We define the accuracy of the model for the parameter i as

$$\text{Accuracy}_i = \sqrt{\langle (\theta_i^T - \mu_i)^2 \rangle}, \quad (8)$$

where θ_i^T is the true value of the parameter i and the average is done over all galaxies in the considered set (e.g., the test set). Meanwhile, we define the precision of the model on the parameter i as

$$\text{Precision}_i = \left\langle \frac{\sigma_i}{\mu_i} \right\rangle. \quad (9)$$

Gradient boosting trees. With this method we only perform parameter regression, and the output of the model is the predicted value of the parameter i : $\tilde{\theta}_i$. In this case we only define the model accuracy,

$$\text{Accuracy}_i = \sqrt{\langle (\theta_i^T - \tilde{\theta}_i)^2 \rangle}. \quad (10)$$

We emphasize the differences between our definitions of accuracy and precision. Accuracy quantifies the dispersion around the true values (independently of the size of the error bars for the prediction), while precision estimates the size of the relative errors (independently of whether the values are close to or far from the true values).

Finally, we note that the accuracy and precision as defined above will give more weight to low-mass galaxies, as those are the most abundant in the simulations. When studying how constraints change for different galaxies in a given simulation,

it will prove useful to quantify the dependence of accuracy and precision on stellar mass.

3. Results

We start by training a neural network that takes as input the 17 properties of individual galaxies of the IllustrisTNG simulations at $z=0$ and outputs the posterior mean and standard deviation for each cosmological and astrophysical parameter. Once the network is trained, we test it using the properties of individual galaxies of the test set. In Figure 1 we show the derived posterior means and standard deviations for 50 random galaxies versus their true value.

The network has not found enough information to infer the value of A_{AGN1} , A_{AGN2} , and σ_8 , so it just predicts the mean value with large error bars for these parameters. For the supernova parameters, A_{SN1} and A_{SN2} , the network may be using some information to provide some loose constraints (we provide further details in Appendix D). On the other hand, for Ω_m , the network seems to have found enough information to determine its value for almost all galaxies considered. We emphasize that these constraints are derived for individual galaxies, each having a different cosmology and astrophysics model, that were selected randomly, e.g., independently of their stellar mass and whether they are centrals or satellites.

From Figure 1 we cannot tell whether the network is able to infer the value of Ω_m for any generic galaxy or whether we were lucky with the random selection we carried out in that exercise. To shed light on this question, we compute the averages of the means and standard deviations of the posteriors for all galaxies in each simulation of the test set,

$$\bar{\mu}_{i,s} = \frac{1}{N_s} \sum_{j \in s} \mu_{i,j} \quad \bar{\sigma}_{i,s} = \frac{1}{N_s} \sum_{j \in s} \sigma_{i,j}, \quad (11)$$




Figure 2. We trained neural networks using galaxies from 850 IllustrisTNG simulations and have reserved all the galaxies from 50 additional IllustrisTNG simulations for the test set. For each galaxy of a given simulation of the test set we compute the posterior mean and standard deviation. The bottom panels show the results for 150 individual galaxies of three different simulations with three different values of Ω_m (shown with a horizontal solid line) color-coded according to the value of the stellar mass of the galaxy. Galaxies are organized according to their stellar mass; galaxies on the left are small while the ones on the right are large. We have then computed the posterior mean and standard deviation from all galaxies in a simulation (Equation (11)) and plotted the results in the top panel. The black points in that panel show the results for the simulations in the bottom panels. The numbers inside the top panel show the accuracy and precision of the model. All results are at $z = 0$. As can be seen, our network is able to infer the value of Ω_m for the vast majority of galaxies in a given simulation.

where i denotes the considered parameter (e.g., Ω_m) and j runs over all N_s galaxies of a given simulation s . In the top panel of Figure 2 we show the above values for each of the simulations in the test set. In the bottom-right part of that panel, we quote the accuracy and precision of the model. As can be seen, on average for all galaxies, the network is able to infer the value of Ω_m with an accuracy of 0.034 and a 10.5% precision.

We perform the following exercise to investigate in more detail whether our model works for all galaxies or just a subset of them. First, we select three different simulations of the test set with different values of Ω_m : one low, one high, and one intermediate. From each of those simulations we randomly select 150 galaxies. For each of those galaxies, we compute the posterior mean and standard deviation of Ω_m . In the bottom panels of Figure 2 we show the results. The constraints are color-coded according to the stellar mass of the galaxies. Those plots show that our network not only works for a subset of galaxies but seems to perform well for the majority of the galaxies in a given simulation.

Three features are worth noticing. First, in all cases there seem to be some outliers where the posterior mean is significantly away from the true value. Second, for the models with intermediate and high values of Ω_m , the size of the standard deviation of the posterior is very similar for all galaxies,¹⁵ while for the cosmology with a low value of Ω_m we find that massive galaxies have smaller posterior variances than low-mass galaxies. Third, from the top panel of Figure 2 we can see that in some simulations there seem to be systematic differences between the posterior means and the true value. We will attempt to provide an explanation for these features in Section 4.

From the above results we conclude that there is evidence showing that the value of Ω_m can be inferred from the properties of individual galaxies for the vast majority of the cases. This statement holds for galaxies with very different

¹⁵ For the model with high Ω_m , the minimum and maximum values only vary by a factor of ~ 3 while for the model with low Ω_m the difference is more than a factor of ~ 7 .

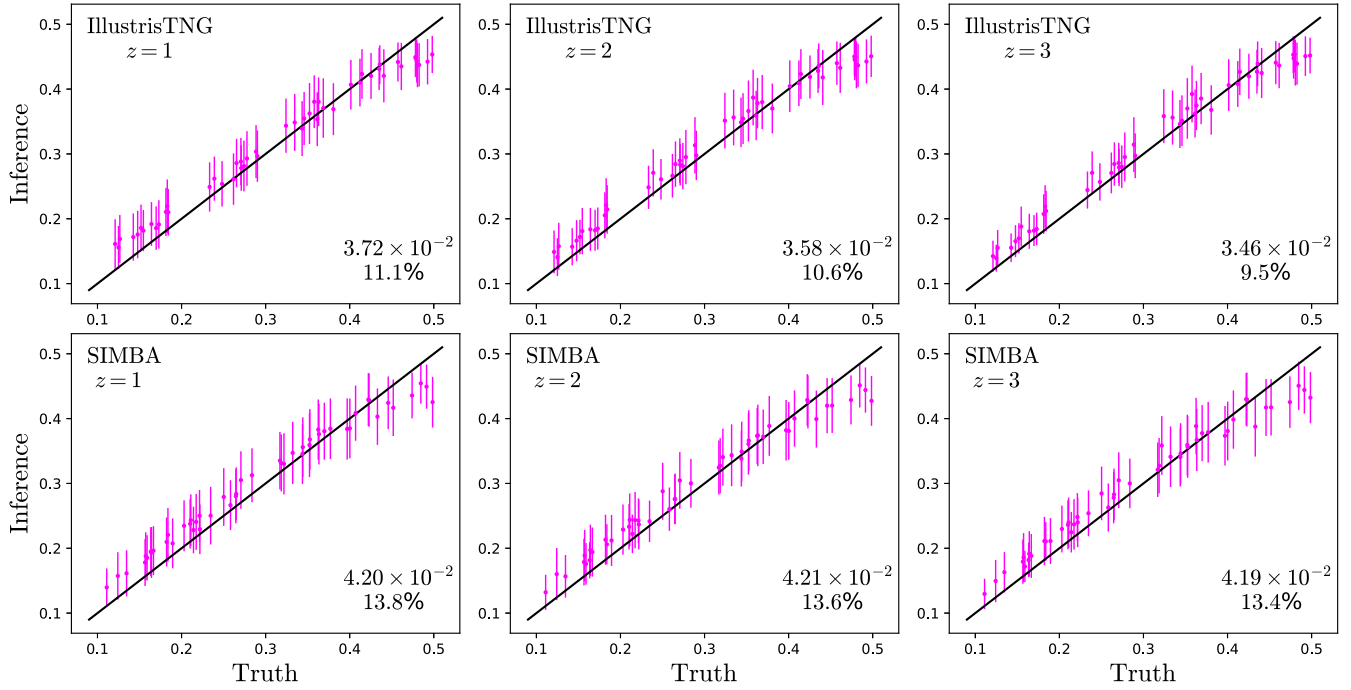


Figure 3. Redshift dependence. We have trained neural networks to infer the value of the cosmological and astrophysical parameters using properties of individual galaxies at different redshifts and for galaxies of the IllustrisTNG and SIMBA simulations. For each galaxy of each simulation of the test set we compute the posterior mean and standard deviation for Ω_m . Next, we compute the mean of those two numbers (Equation (11)) and plot them in the figure for the 50 different simulations in the test set. We show results at redshifts 1, 2, and 3. The numbers in the bottom-right corner show the model accuracy and precision. As can be seen, our networks can infer the value of Ω_m from individual galaxies at redshifts higher than $z = 0$ with an accuracy similar to the one achieved by the models at $z = 0$.

cosmologies, astrophysics, and almost independently on whether the galaxy is massive or dwarf, central or satellite.¹⁶

3.1. Dependence on Method and Data

We have carried out other sanity checks to investigate whether our conclusions hold for different methods and different simulations:

1. We have repeated the above analysis but using galaxies from the CAMELS-SIMBA simulations (using their 14 properties) instead of the IllustrisTNG galaxies, reaching the same conclusions as above. We provide further details of this test and its results in Appendix A.
2. We have repeated the above exercise but performed parameter regression through gradient boosting trees. We have trained these models using both galaxies from the IllustrisTNG and SIMBA simulations (separately). We find that the accuracy of these methods on Ω_m is similar to the one from neural networks.

These tests indicate that our results are robust to the particularities of the method used to perform the mapping between galaxy properties and the value of Ω_m .

3.2. Dependence on Redshift

We now investigate whether our results only hold at $z = 0$, or we can also infer the value of Ω_m from the internal properties of galaxies at higher redshifts. For this, we have trained neural networks to infer the values of the cosmological and

astrophysical parameters from galaxies at redshifts 1, 2, and 3 using both the IllustrisTNG and the SIMBA suites.

Once the models are trained, we test it on individual galaxies from simulations of the test set and compute the average posterior mean and posterior standard deviation from all galaxies in a given simulation (i.e., Equation (11)). We then show these measurements in Figure 3.

As can be seen, results at redshifts higher than zero are qualitatively very similar to the ones at $z = 0$ for both IllustrisTNG and SIMBA galaxies. We computed the accuracy and precision for all models and quote them in the bottom-right part of each panel. We find that both the accuracy and precision of the models are very similar across redshifts, although there is a slight improvement when using galaxies at higher redshifts. The models trained on IllustrisTNG galaxies, however, exhibit a better accuracy and precision than the ones trained on SIMBA galaxies. This is due to the inclusion of the three additional features contained in the IllustrisTNG galaxies (the magnitudes in the U , K , and g bands). This can be seen more clearly in the top-left and bottom-right panels of Figure 4, where models trained on the same variables from IllustrisTNG and SIMBA exhibit similar accuracy and precision. Overall, we conclude that it seems possible to infer the value of Ω_m from the internal properties of galaxies at redshifts $z \leq 3$.

Next, we investigate whether our results are independent of redshift, i.e., whether a model trained on galaxies at a given redshift is able to infer the value of Ω_m from galaxies at a different redshift. We have tried this on different models at different redshifts and found that it does not work. We have also tried a few different things to verify that the reason was not due to the use of comoving versus proper quantities (see, e.g., Shao et al. 2022) but we did not find any improvement. From these tests we conclude that the mapping between the internal

¹⁶ We emphasize that we do not provide the models with information on whether a galaxy is a central or a satellite.

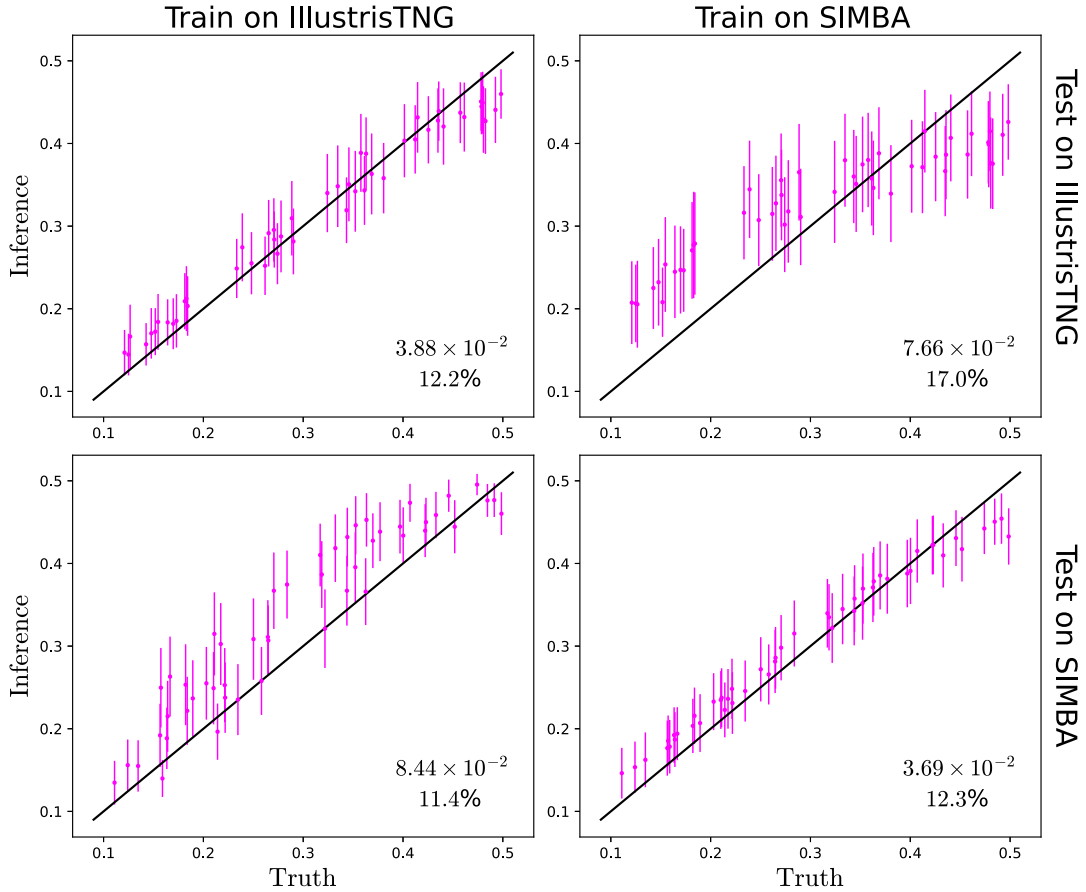


Figure 4. Robustness test. We have trained neural networks to perform likelihood-free inference on the value of the cosmological and astrophysical parameters using the internal properties of individual galaxies at $z = 0$. In this case we made use of the 14 internal properties that are common between the galaxies in the IllustrisTNG and SIMBA simulations. We have trained models using galaxies from either the IllustrisTNG or SIMBA simulations. For each simulation in the test set, we compute the posterior mean and standard deviation for Ω_m for each galaxy on it. We then compute the average value of those two numbers from all galaxies in a given simulation (Equation (11)). These panels show the results for all 50 simulations in the test set when training on galaxies of a given simulation and test it on galaxies of the same simulation or another simulation. In the bottom-right part of each panel we quote the accuracy and precision of the model on the tested galaxies. As can be seen, when the model is tested on galaxies from simulations different to the ones used for training, the model is not able to infer the correct cosmology in most cases. This indicates that the model is not robust and may be using information that is specific to each galaxy formation model.

galaxy properties and Ω_m should have an intrinsic dependence on redshift. In the next section we attempt to provide a physical understanding of this result.

3.3. Robustness

Ideally, we would like to apply this method to the internal properties of real galaxies to derive the value of Ω_m and see whether it agrees with the one derived from standard cosmological measurements (e.g., CMB or galaxy clustering). However, to carry out that task, we need a robust model, i.e., one that works independently of the type of simulations used for training. At its core, CAMELS was designed to test the robustness of models by providing simulations from two distinct suites: IllustrisTNG and SIMBA.

Here we quantify the robustness of our model by testing on one suite the models that were trained on the other suite. We show the results of such exercise in Figure 4. We find that while testing the model on galaxies from the same subgrid model as the one used for training yields precise and accurate results for both the IllustrisTNG and SIMBA models, the model fails when tested on galaxies from different subgrid models. In Appendix B we provide further details on this test.

We have repeated this exercise with the gradient boosting tree method, reaching the same conclusions. We have also tried

with a smaller set of variables, e.g., $\{M_*, V_{\max}, Z_*, \dots\}$, but the models are still not robust. We thus conclude that our models may be learning something particular about each simulation code or that the two suites do not overlap in the parameter space of possible models. In the next section we shall see that one reason behind this behavior is that the two different suites of simulations produce very different galaxies with distinct properties, limiting the range where they both overlap and therefore making the model not robust.

It is interesting to note that the trends we observe in the cross-analysis of Figure 4 go in the same direction. In other words, when the model is trained on IllustrisTNG/SIMBA galaxies, it tends to overpredict the value of Ω_m when tested on galaxies from SIMBA/IllustrisTNG. We would have expected that if the model trained on IllustrisTNG overpredicts when tested on SIMBA, the model trained on SIMBA should underpredict when tested on IllustrisTNG. Investigating this is beyond the scope of this work, but it may be a consequence of the topological properties of the manifold hosting the galaxy properties.

4. Interpretation

In this section we attempt to provide a physical explanation to our findings above. We will focus our attention on Ω_m .

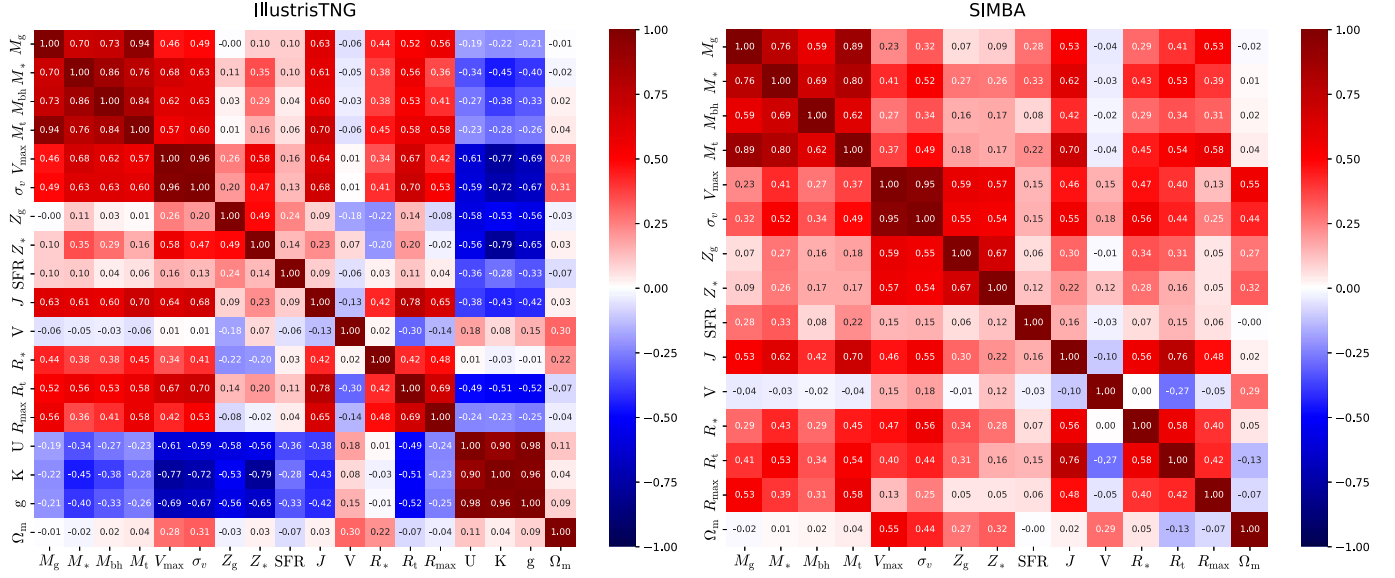


Figure 5. We have computed the correlation matrix (see Equation (12)) of the galaxy properties plus Ω_m for the IllustrisTNG (left) and SIMBA (right) simulations. We find a strong linear correlation among different galaxy properties (e.g., gas mass and total mass), but the correlations between Ω_m and the galaxy properties are relatively mild. This indicates that the value of Ω_m cannot be inferred due to simple, linear correlations between Ω_m and galaxy properties.

4.1. Linear Correlations

We start by investigating whether there are strong linear correlations between the galaxy properties and Ω_m . For that, we plot in Figure 5 the correlation matrix of all galaxy properties plus Ω_m , defined as

$$R_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}}, \quad (12)$$

where

$$C_{ij} = \langle (p_i - \bar{p}_i)(p_j - \bar{p}_j) \rangle, \quad (13)$$

where p_i refers to the i feature of the data vector (galaxy properties plus Ω_m) and $\bar{p}_i = \langle p_i \rangle$. This matrix gives information about the linear correlation between the different variables.

As can be seen, while some galaxy properties seem to be highly correlated (e.g., V_{\max} and σ_v), the linear correlations between Ω_m and the galaxy features are not particularly high. For IllustrisTNG galaxies, the strongest correlated variable with Ω_m is σ_v , while for SIMBA galaxies is V_{\max} .

These tests indicate that our findings are not due to simple linear correlations between Ω_m and galaxy properties. To further reinforce this point, we have explicitly performed linear regression from the galaxy properties to the value of Ω_m , finding that the model performs badly; the model mainly predicts the mean value of Ω_m with very little scatter. We have carried out this exercise for both the IllustrisTNG and SIMBA galaxies and using different numbers of galaxy properties, finding that the results barely change. This reinforces our interpretation that the relation between galaxy properties and Ω_m cannot be described by a simple linear relation.

As a side calculation we have also carried out an analysis with the Principal Component Analysis (PCA) to try to identify the number of components and variables that are responsible for most of the overall data variance (i.e., considering both galaxy properties plus Ω_m). For IllustrisTNG galaxies we find that the first principal component is dominated by Ω_m as well as V , V_{\max} , and σ_v , while for SIMBA the most important features

are Ω_m and V_{\max} , followed by Z_g , Z_* , σ_v , and V . It is interesting to see that Ω_m and V_{\max} seem to form a basis to explain most of the data variance.

4.2. Properties Ranking

Next, we try to identify the most important galaxy features that the network is using in order to carry out the inference. We have used different methods to perform this task, like computing saliency maps and SHapley Additive exPlanation (SHAP) values for the neural networks, and using the feature importance method for random forest and gradient boosting trees regressors. However, we found that these methods did not allow us to identify the most important features, likely because of the strong internal correlations between the different variables. In Appendix C, we provide additional details about our results when using SHAP values.

We therefore tackle this problem as follows. First, we train a model using all galaxy properties and record its accuracy. Next, we remove one of the considered properties and retrain a model using the rest of the properties. We then reincorporate that feature, remove another property, and train another model on those variables. We repeat this procedure until all properties have been independently removed. For instance, we train a model that contains all properties except gas mass, we train another model that contains all properties except stellar mass, we train another model that contains all properties except black hole mass, and so on. For each model we save the accuracy obtained. This method allows us to quantify the worsening of the model accuracy by removing a single feature.

We then continue the exercise by removing the variable that changes the accuracy the least. With the subset of variables left, we repeat the above procedure and train models where we remove one galaxy property at a time and record the model accuracy. In this way we can rank order¹⁷ the features

¹⁷ We note that this method is not guaranteed to give the correct ordering in general. For instance, removing two or more properties at a time may lead to a different ordering.

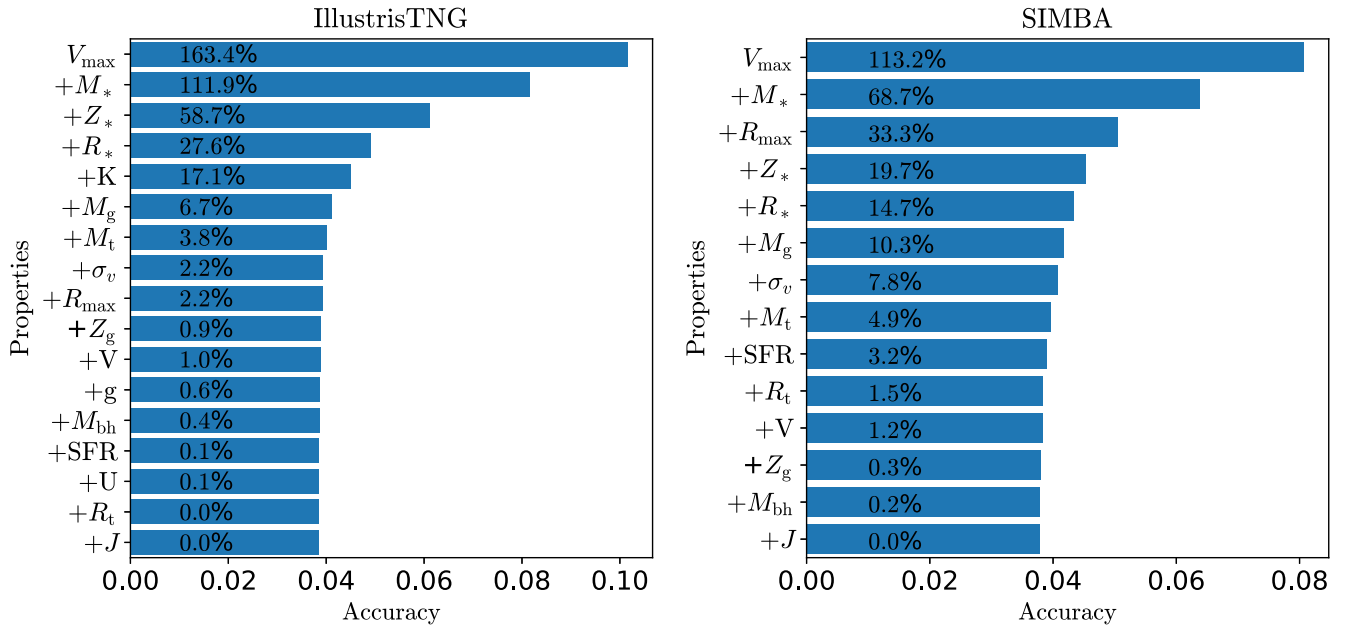


Figure 6. We rank order the galaxy properties for both IllustrisTNG (left) and SIMBA (right) such that the variables contributing the most to the model accuracy are on top while the features contributing the least are on the bottom (see text for details on the procedure used). The horizontal bars indicate the accuracy (in terms of RMSE) achieved by the considered variables, cumulatively from top to bottom, and the black numbers inside them show the loss in accuracy with respect to a model trained using all variables. For instance, for the IllustrisTNG galaxies, a model that only uses V_{\max} achieves an RMSE of ~ 0.1 and performs 163.4% worse than the model trained on all 17 properties (with an RMSE of ~ 0.04). Likewise, a model trained on SIMBA galaxies using $\{V_{\max}, M_*, R_{\max}, Z_*, R_*\}$ achieves an RMSE of ~ 0.04 , which is only 14.7% worse than the model trained on all 14 galaxy properties. We emphasize that this ordering was derived when training gradient boosting trees models to perform regression to the value of Ω_m .

according to their contribution to the model accuracy. Unfortunately, doing this exercise with neural networks while performing hyperparameter optimization is too computationally expensive for this work, so we decided to do it using gradient boosting trees instead of neural networks.

We show the rank-ordered features in Figure 6. We find the two most important features to be V_{\max} and M_* for both IllustrisTNG and SIMBA galaxies. The stellar metallicity and stellar radius are also among the five most important features in both cases. However, for IllustrisTNG galaxies, the K -band magnitude appears as a very relevant property (this property is not present in the SIMBA galaxies) while in the case of SIMBA galaxies the radius associated with the maximum circular velocity, R_{\max} , is selected as an important feature. In Figure 6 we show the accuracy (quantified in terms of the rms error) gained as we add variables cumulatively from top to bottom. For IllustrisTNG galaxies, using the subset of $\{V_{\max}, M_*, Z_*, R_*, K\}$ only degrades results by 17% with respect to the accuracy achieved by training on all 17 properties. Meanwhile, for SIMBA galaxies, using the subset of $\{V_{\max}, M_*, R_{\max}, Z_*, R_*\}$ only degrades results by 15% with respect to training using all 14 features.

Next, with these subsets of five variables we have trained neural networks to perform likelihood-free inference. For IllustrisTNG/SIMBA galaxies we find that the accuracy on predicting Ω_m degrades by 27%/28% when comparing it to the accuracy of model trained using all 17/14 variables. When using the five most important features according to the absolute SHAP values (e.g., $\{M_*, K, M_g, Z_g, V_{\max}\}$ for the IllustrisTNG simulations), we found that the model performs significantly worse: the rms error between the posterior mean and the true value degrades by 47%. These results show that this procedure can find a minimum set of variables that is responsible for most of the model accuracy.

4.3. Visual Inspection

Before attempting a physical explanation of our results with the information gained from the above experiments, we perform a visual inspection of some galaxy features in two and three dimensions to gain intuition. For this, we randomly select 10,000 galaxies from 100 different IllustrisTNG simulations (100 galaxies per simulation); we do the same exercise for the SIMBA simulations. For this exercise we consider three galaxy properties: V_{\max} , M_* , and Z_* . We have chosen these variables because V_{\max} and M_* are the most important ones for both IllustrisTNG and SIMBA galaxies, while Z_* is among the four¹⁸ most important variables in both suites.

In Figure 7 we show the 2D and 3D projections of the data. Each point, representing a galaxy, is color-coded according to its Ω_m value. As can be seen, galaxy properties occupy different regions in the 2D and 3D plots depending on the value of Ω_m . In particular, the dependence of the $V_{\max}-M_*$ relation on Ω_m is particularly pronounced. We will discuss this trend in more detail in the next subsection. We emphasize that galaxies are randomly selected from the simulations, i.e., they not only differ in their value of Ω_m but also in σ_8 and the four astrophysical subgrid parameters.

From Figure 7 we can also see the large, intrinsic differences between the SIMBA and IllustrisTNG galaxies: While they exhibit similar qualitative dependence with Ω_m , they populate the parameter space differently. This is, however, expected, given the large differences between the IllustrisTNG and SIMBA subgrid models. We note that in higher dimensions, the

¹⁸ We note that Z_* is the third and fourth most important variable for the IllustrisTNG and SIMBA galaxies, respectively. The third most important variable for SIMBA is R_{\max} , but it has a rather low importance ranking for IllustrisTNG.

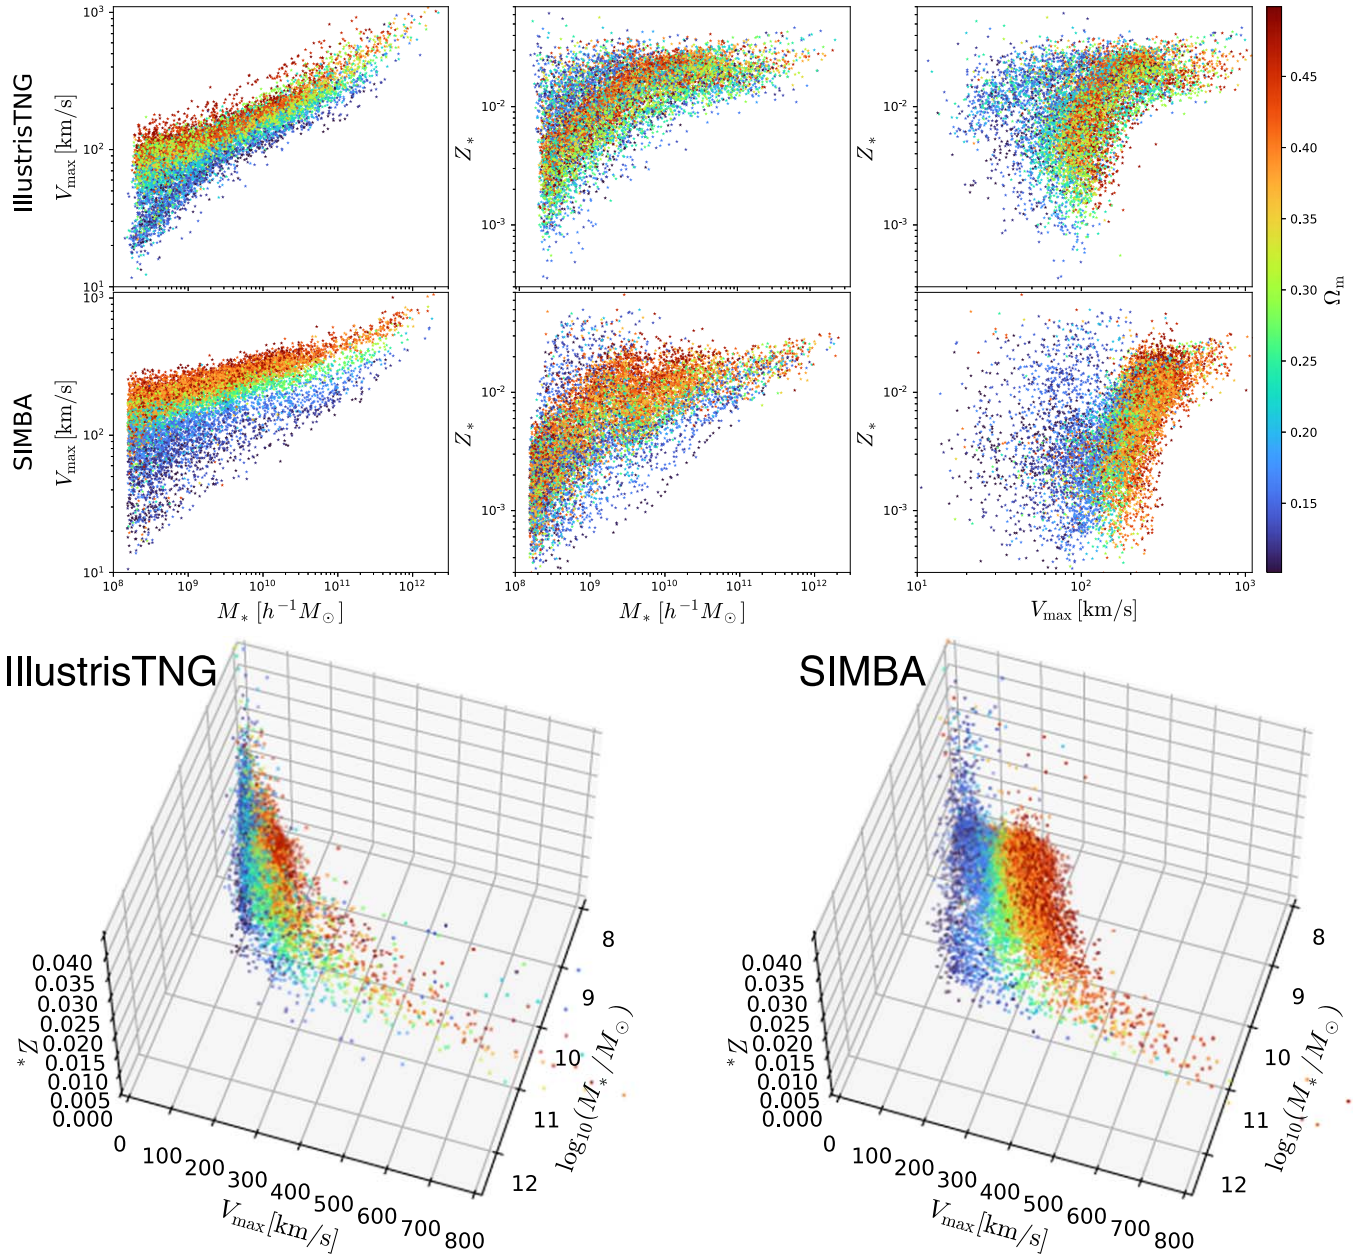


Figure 7. For both the IllustrisTNG and SIMBA suites we have randomly taken 100 simulations. From each simulation we have randomly selected 100 galaxies at $z = 0$, for a total of 10,000 galaxies. Top: For each of those galaxies we show correlations between V_{max} and M_* (left), Z_* and M_* (middle), and Z_* and V_{max} (right) for the IllustrisTNG (top row) and SIMBA (bottom row) galaxies. Each galaxy is color-coded according to the value of Ω_m of its simulation (blue/green/red indicate low/medium/high values of Ω_m). As can be seen, there is a prominent correlation between V_{max} and M_* that changes with Ω_m . We can also observe other more complex trends in the Z_*-M_* and Z_*-V_{max} planes with Ω_m . Bottom: We show the results in 3D. Galaxies occupy different regions in the properties space depending on their value of Ω_m . We believe that in higher dimensions (i.e., considering more galaxy properties), galaxies should occupy even more disconnected regions as a function of Ω_m . We interpret these results as Ω_m changing the manifold where galaxy properties reside in a different way as feedback does. Machine-learning methods can use these patterns to determine the value of Ω_m .

differences between the simulations may be even more pronounced. We believe that this is the reason why our models are not robust, i.e., a model trained on galaxy properties from IllustrisTNG simulations does not work when tested on SIMBA galaxies, and the other way around.

4.4. Physical Interpretation

We now discuss the physics behind our results. As we saw in Figure 7, galaxy properties populate the parameter space differently depending on the value of Ω_m . This indicates that Ω_m induces an effect on galaxy properties, or on a subset of

them, that cannot be mimicked by astrophysical effects. Let us first focus our attention on the two most important properties of both the SIMBA and the IllustrisTNG galaxies: V_{max} and M_* . From the top panels of Figure 7 we can see that at fixed stellar mass, the maximum circular velocity increases monotonically with Ω_m . This may be explained by taking into account that higher values of Ω_m will increase the dark matter density in the universe, and therefore, more dark matter is expected to reside in galaxies, enhancing their gravitational potential well and therefore their V_{max} value. However, feedback from supernovae and AGNs are also expected to affect the stellar mass of the galaxy, introducing some scatter in the M_*-V_{max} relation.

As we shall see below, our tests suggest that Ω_m is not imprinted onto a single property (e.g., V_{\max}) and that even knowing the value of the astrophysical parameters perfectly does not significantly help due to intrinsic scatter among different galaxies within any given model with a fixed set of parameters. Thus, we may think that Ω_m may change the manifold where galaxy properties live and that change is different from the one induced by changes in feedback.

This explanation could shed light on why we cannot determine the value of σ_8 with a single galaxy. In contrast to Ω_m , σ_8 will only change the amplitude of the initial matter fluctuations and that by itself is unlikely to induce systematic differences in galaxy properties. σ_8 may, however, affect the abundance of galaxies, in particular, that of the most massive ones, similarly to what it does for the halo mass function. Thus, while a single galaxy (unless located in the high-mass end) may not be enough to infer σ_8 , a set of galaxies can potentially be used as a probe of σ_8 . We leave this direction of study for future work.

As we saw in the results section, for some simulations, the predictions of the models seem to exhibit an overall bias. This may be due to the following reason. The networks may be learning some function that approximates the galaxy properties' manifold and its dependence on Ω_m . However, due to the limited data we have to train them, it may happen that the learned manifold may be off with respect to the true one. In this case we will expect an overall bias between the prediction of the network and the true value for all galaxies in the considered simulation.

4.5. Breaking Degeneracies with Astrophysics?

We may wonder whether the network is aware of the clear dependence on Ω_m of the V_{\max} versus M_* relation but needs additional information to break the degeneracy between cosmology and astrophysics. Maybe in this case the network is using the other properties (e.g., Z_* , R_* , and K) to first constrain astrophysics (i.e., feedback parameter values) and then determine cosmology. To test this hypothesis we train a network using as input variables $\{M_*, V_{\max}, A_{\text{SN1}}, A_{\text{SN2}}, A_{\text{AGN1}}, A_{\text{AGN2}}\}$. If our hypothesis holds, by providing the network with the true value of the feedback parameters plus M_* and V_{\max} , it would be able to infer Ω_m accurately. However, we find that this model performs very badly when inferring the value of Ω_m : Its accuracy decreases by 91% with respect to the model trained on all variables. This test indicates that the network is not simply extracting information from V_{\max} and M_* and using the other variables to break the degeneracies between cosmology and astrophysics.

Next, we test whether knowing the value of the astrophysical parameters adds additional information to those already contained in the galaxy properties. To quantify this, we train a network using the 17 properties of the galaxies from the IllustrisTNG simulations plus the value of A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2} . We find that results barely improve: The model accuracy and precision increases by 3% and 5%, respectively. This indicates that most of the information the network is extracting is already contained in the internal galaxy properties, and knowing the value of the feedback parameters perfectly does not add any significant additional information. This and the above test indicate that the networks are not trying to infer feedback to break some degeneracies with a particular

observable, but rather that it is the observable itself that is sensitive to Ω_m .

We note that it is well known that galaxy properties change with redshift. This not only explains why the models we train at $z = 0$ do not work at higher redshifts but also why knowing the value of the astrophysical parameters perfectly does not add much information because these values will be the same across redshifts.

4.6. Dark Matter Content

The explanation we formulated above to interpret our findings relies on dark matter playing a crucial role in galaxies. In order to test this hypothesis, we performed the following test. We have trained networks on galaxies from the IllustrisTNG simulations using all properties except V_{\max} , σ_v , M_t , R_t , and R_{\max} . These are quantities that are expected to receive large contributions from the dark matter component of galaxies, and therefore, is a way to quantify how important it is for the network to know the dark matter component or the depth of the gravitational potential well. We find that the network trained with this configuration is still able to infer the value of Ω_m but with much lower accuracy: 96% worse than the model trained on all properties. This test indicates that these variables are very important, although Ω_m leaves some weaker signatures on the other galaxy properties. In a complementary way, we have checked that in the case of the IllustrisTNG galaxies, once we have identified the five most important variables $\{V_{\max}, M_*, Z_*, R_*, K\}$, removing V_{\max} from that set completely cancels the constraining power. In other words, for that subset, V_{\max} is needed to infer Ω_m . From these tests we conclude that the network may be using information either about the dark matter content of the galaxy or about its gravitational potential well.

Next, to reinforce our explanation, we test explicitly whether the dark matter content of galaxies increases with Ω_m . We have taken the 100 galaxies for 100 different models that we discussed above and plot in Figure 8 the V_{\max} versus M_* projection for those galaxies. The top panels show the galaxies color-coded by their value of Ω_m and show the trend we already discussed above (the top panels are identical to the panels in the left column of Figure 7). The panels in the middle row show the results color-coded by the ratio between the dark matter mass¹⁹ and stellar mass in the galaxies. We use that ratio and not the dark matter mass itself because the latter has a strong correlation with stellar mass, making the visualization more challenging. As can be seen, for a fixed value of the stellar mass, the larger the dark matter mass, the higher the value of V_{\max} . This trend is very clear for IllustrisTNG galaxies; meanwhile, for SIMBA it is also clear for low- and high-mass galaxies, while for intermediate galaxies ($9.3 < \log M_*/(h^{-1}M_\odot) < 11$) the dependence is much weaker. This is the same trend we find with Ω_m (top panels), indicating that larger values of Ω_m will tend to increase the dark matter content of galaxies.

We note that increasing the dark matter content of galaxies can also affect other galaxy properties. For instance, changing Ω_m will affect the halo collapse time and concentration, and these may leave an imprint on Z_* , R_{\max} , and R_* . However, the relationship between these variables and the V_{\max} – M_* plane is not clearly visualized in a three-dimensional plot (two plus

¹⁹ The dark matter mass is computed as $M_t - M_g - M_* - M_{\text{BH}}$.




Figure 8. We have randomly taken 100 simulations from the IllustrisTNG (left column) and SIMBA (right column). For each simulation, we take 100 random galaxies at $z = 0$. We then project these galaxies onto the V_{\max} – M_* plane. Each row shows the data color-coded according to the value of Ω_m (top), $\log(M_{\text{dm}}/M_*)$ (middle), and $\log(M_{\text{max}}/M_*)$ (bottom), where $M_{\text{dm}} = M_t - M_g - M_* - M_{\text{BH}}$ is the dark matter mass in the galaxy and $M_{\text{max}} = V_{\max}^2 R_{\text{max}}/G$ is the total matter mass contained within R_{max} . We find that at fixed stellar mass, V_{\max} increases with both Ω_m and M_{dm} , supporting our hypothesis that increasing the value of Ω_m increases the dark matter content of galaxies, making the gravitational potential deeper and therefore enhancing V_{\max} . From the third row we can however see that at fixed stellar mass, V_{\max} is more strongly correlated with M_{max} ; this may explain why the network prefers to extract information from V_{\max} rather than the subhalo total mass or dark matter mass. We note that the reason why we use M_{dm} and M_{max} normalized to the stellar mass is because there is a strong correlation between these quantities and M_* . By taking the ratio we get rid of that dependence, simplifying the visualization of the results.

color) as in Figure 8. We argue that this is due to the high-dimensional manifold on which these features depend on Ω_m .

On the other hand, we may also expect that differences in σ_8 will lead to changes in halo formation time and concentration. Because we cannot infer the value of σ_8 from individual properties of galaxies, we think the effect of Ω_m on galaxy properties should not be primarily driven by the above changes to halo properties, or perhaps be a distinct change to the one induced by σ_8 .

4.7. V_{\max} versus M_t

The above results corroborate our interpretation that changing the value of Ω_m affects the dark matter content of

galaxies, an effect that is physically different from the one from feedback. However, at this point we may wonder why the network prefers to use V_{\max} rather than other properties that are expected to be heavily affected by dark matter, such as the galaxy’s subhalo total mass (M_{tot}) or velocity dispersion (σ_v). To verify that this is indeed the case, we have trained models with galaxies of the IllustrisTNG simulations using as features $\{M_*, M_t, Z_*, R_*, K\}$ and $\{M_*, \sigma_v, Z_*, R_*, K\}$. We find that using these variables the accuracy of the model on Ω_m degrades by 100% and 43%, respectively. This clearly indicates that V_{\max} contains more information than M_t and σ_v . We believe that this may be happening because it is known that V_{\max} correlates more strongly with stellar mass than with subhalo mass

(Conroy et al. 2006). For instance, when halos are accreted into larger halos they may lose a significant fraction of their dark matter content due to tidal forces. That effect will change the dark matter content of galaxies significantly, but the value of V_{\max} may remain rather stable because it mostly probes the mass in the inner regions of the subhalo, which are the least affected by the above processes.

To validate this hypothesis we plot in the bottom row of Figure 8 the galaxies mentioned above but color-coded according to M_{\max}/M_* , where $M_{\max} = V_{\max}^2 R_{\max}/G$. We find for IllustrisTNG galaxies a similar trend to when we used the dark matter mass, while for SIMBA galaxies the trend is now much more evident: For a fixed stellar mass, increasing the value of M_{\max} increases the value of V_{\max} . This indicates that either V_{\max} (or M_{\max}) is a better and more stable proxy for the dark matter content of galaxies than the total subhalo mass or its velocity dispersion.

The above tests may indicate that the network is focusing its attention on the dark matter or total mass content of galaxies in their central region, or maybe directly into the depth of the gravitational potential, rather than in the total dark matter mass in the subhalo's galaxy.

5. Summary and Discussion

In this paper we have shown that it may be possible to infer the value of Ω_m with a precision of $\delta\Omega_m/\Omega_m \simeq 10\%-15\%$ and an accuracy of $\sim 0.035-0.042$ from the internal properties of individual galaxies and their subhalos. This result holds for galaxies of either the CAMELS-IllustrisTNG or CAMELS-SIMBA simulations when using neural networks (to do likelihood-free inference) or gradient boosting trees (to do parameter regression) and at all redshifts considered, $0 \leq z \leq 3$.

We have shown that Ω_m has a large effect on the $V_{\max}-M_*$ relation, although our constraints do not arise from those two variables alone, even if the astrophysical parameters are known perfectly. We believe that the explanation behind our results is that galaxy properties reside in a thin high-dimensional manifold that changes with Ω_m . That change is different from the one induced by astrophysical effects. We think that the physics behind the unique change in the manifold is that Ω_m affects the dark matter content of galaxies. Machine-learning methods can be trained to find these manifolds and therefore to infer the value of Ω_m . We note that Neistein et al. (2012) found that the formation history of a single massive galaxy was enough to calibrate the entire galaxy population in semianalytic models, perhaps indicating the intrinsic correlation among galaxy properties and their location in the high-dimensional manifold.

We note that physically, the effect of Ω_m is very different from the one of σ_8 , which will just change the amplitude of the initial linear matter fluctuations, and therefore, we do not expect it to imprint unique features on galaxy properties. This could explain why our models cannot infer the value of σ_8 from individual galaxy properties.

5.1. Robustness

We caution the reader that our models are not robust; if the models are trained on galaxies from the IllustrisTNG simulations, they cannot infer the value of Ω_m from galaxies of the SIMBA simulations, and vice versa. We believe that this may

be due to the intrinsic differences between the galaxy properties in the two different simulations (see Figure 7).

While this method, in its current form, cannot be used with real data yet due to the lack of robustness, it will be interesting to explore the use of contrastive learning (Le-Khac et al. 2020) to force the network to learn only unique (physical) features that are not simulation/model dependent. Another possible avenue will be to try to develop a theoretical template (e.g., using symbolic regression) and calibrate its parameters directly with real data. We leave these questions for future works.

Once the model is robust, it will be important to quantify how much our constraints degrade by accounting for the observational uncertainties associated with the different galaxy properties. For instance, the value of the relevant galaxy properties such as the stellar mass and the maximum circular velocity can only be inferred with finite precision from observations. We leave for future work exploring the effects of these uncertainties in the constraints of Ω_m . On the other hand, if our interpretation is correct and galaxy properties live in a manifold sensitive to cosmology and astrophysics, one can use that information to reduce uncertainties in galaxy properties by requiring them to be in a manifold. In other words, in the real universe, galaxy properties will reside in a manifold with a fixed cosmology and astrophysics. Thus, there will be high-dimensional correlations that may allow us to determine the value of some galaxy properties with higher accuracy.

5.2. Ω_b

Due to the design of the CAMELS simulations, we can only train models to infer the value of Ω_m and σ_8 , because in all simulations we have kept the values of the other cosmological parameters fixed. It will be important to repeat this work using simulations that vary the value of other cosmological parameters to investigate whether individual galaxies can constrain other parameters but also to study whether degeneracies among parameters will deteriorate the constraining power of this method on Ω_m .

In Section 4 we have seen that our models rely on both galaxy properties and the depth of the gravitational potential well (or the mass in the galaxy core) to infer the value of Ω_m . Thus, galaxy properties may also be sensitive to Ω_b , as varying that parameter will change the abundance of baryons in galaxies. Thus, it will be interesting to investigate whether galaxy properties are also sensitive to Ω_b . On the other hand, it may happen that galaxy properties are sensitive to some particular combination of Ω_b and Ω_m , e.g., to its ratio: Ω_b/Ω_m . While we cannot provide an answer to these questions (as it will require running many simulations with different values of Ω_b), we can however attempt to provide a qualitative indication of what may be happening. For this, we have run six additional IllustrisTNG simulations. In these simulations, the value of the astrophysical parameters is set to the fiducial IllustrisTNG model, while σ_8 is 0.8 and $\{\Omega_m, \Omega_b\}$ is given by $\{0.2, 0.025\}$, $\{0.2, 0.075\}$, $\{0.3, 0.025\}$, $\{0.3, 0.075\}$, $\{0.4, 0.025\}$, $\{0.4, 0.075\}$. For each of these six simulations, we randomly select 100 galaxies.

In Figure 9 we show these galaxies projected on the $V_{\max}-M_*$ plane. Galaxies are color-coded according to the value of Ω_m (left) and Ω_m/Ω_b (right). In the background we show a hexbin plot with the distribution of galaxies from the 1000 IllustrisTNG simulations with a fixed value of Ω_b . As can be seen, for a fixed value of M_* , galaxies do not follow a




Figure 9. In order to explore in a very qualitative manner whether our method is sensitive to Ω_b/Ω_m or to Ω_m and/or Ω_b we have run a set of six simulations with different values of Ω_b (0.025 and 0.075) and Ω_m (0.2, 0.3, and 0.4) using the AREPO and IllustrisTNG model (using the fiducial astrophysical model). For each of those simulations we have randomly taken 100 galaxies. In the two different panels we show V_{\max} vs. M_* of those galaxies color-coded according to their value of Ω_m (left) and Ω_m/Ω_b (right). In the background, we show with a hexbin plot the distribution of galaxies from the IllustrisTNG simulations. From the left panel we can clearly see that galaxies no longer follow a monotonic relation of increasing V_{\max} with Ω_m . On the other hand, from the right panel we can see a much more steady and monotonic relation when using Ω_m/Ω_b . We however note that the colors of the galaxies do not really match the ones from the background simulations with fixed Ω_b .

monotonic relation of higher V_{\max} for larger Ω_m when Ω_b is allowed to vary. It seems that in this case the two different values of Ω_b create a bimodal distribution.

In the right panel of Figure 9 we color code the same galaxies as before but using instead Ω_m/Ω_b . In this case, we find a more monotonic relation between V_{\max} and Ω_m/Ω_b at fixed stellar mass. We note however that the colors of these galaxies are a bit off with respect to the ones from the IllustrisTNG set with fixed Ω_b . Thus, while these results indicate that Ω_b/Ω_m is a more relevant variable than Ω_m when Ω_b is not fixed, we cannot tell whether our method is just sensitive to Ω_m/Ω_b or whether in higher dimensions degeneracies can be broken and we can constrain both Ω_m and Ω_b . We note that the value of Ω_b/Ω_m can be constrained from cosmic microwave background data with high accuracy. Thus, if galaxy properties are indeed sensitive to Ω_b/Ω_m , it will be an interesting way to connect two very different observables and physical quantities of the universe.

We emphasize that we have not provided a full physical interpretation of the results presented in this work, beyond stating that changing Ω_m affects galaxy properties in a way different from the one produced by changing astrophysics parameters. However, we know that the dark matter content/total matter content/depth of the gravitational potential is a very important variable for the network. Besides, our results indicate that the network may be more sensitive to Ω_b/Ω_m rather than to Ω_m . One may wonder if the network is somehow measuring the total mass in the center of the galaxy (e.g., through the depth of the gravitational potential well) and also measuring the mass in baryons in that region. This would allow the network to directly infer Ω_b/Ω_m . In the future, it may be interesting to explore whether there are relations between the total matter and the baryonic content in the inner regions of galaxies that somehow are robust to changes in astrophysics. We note that the idea of measuring Ω_b/Ω_m from individual objects was outlined in White et al. (1993), although there it

only applied to the most massive halos where feedback effects cannot expel baryons out to the intergalactic medium.

5.3. Numerical Effects

Given the surprising results our models have achieved, we should ask ourselves: Where does the information come from? In other words, is the network extracting information from a physical or a numerical effect?

Ω_m is imprinted in the simulations through several different effects; it affects, for example, (1) the amplitude and shape of the linear matter power spectrum used to generate the initial conditions, (2) the mass of the dark matter particles, and (3) the cosmic expansion rate.

If the networks are using some nonphysical feature to get the value of Ω_m from the changes to the power spectrum, it would be expected that they would also be able to infer the value of σ_8 that also affects the linear power spectrum. Because our models are unable to constrain the value of σ_8 , we believe this effect should not be the cause of our results.

The one-to-one correlation between Ω_m and the masses of the dark matter particles (note that in CAMELS, Ω_b is kept fixed at 0.049 in all simulations) is something that can be easily learned by neural networks but is not a physical effect. However, in the considered galaxy properties there is no obvious way where the dark matter particle resolution can show up. The dark matter mass of subhalos obeys the relation $M_{\text{DM}} = N_{\text{dm}} m_{\text{dm}}$, where m_{dm} is the mass of a dark matter particle and N_{dm} is the number of dark matter particles in the subhalo, which should be an integer number. There is an intrinsic degeneracy between N_{dm} and m_{dm} for this to work. Besides, if that were the case, the network would need to estimate the mass in dark matter of the halo by subtracting the gas, stellar, and black hole mass from the total mass of the subhalo. We know from our analysis of the relevant features that none of those properties are important for the networks. Other properties, like V_{\max} , σ_v , V , R_p , and R_{\max} seem more

unlikely to be easily related to the mass of the dark matter particles. Thus, we find this hypothesis not very likely.

Ω_m also changes the expansion rate history in the simulation. However, we cannot think of a situation where the model may be learning a numerical artifact associated with this.

Finally, we note that Ω_b/Ω_m is important for setting the internal structure of galaxies (e.g., how baryon dominated the rotation curve is). Thus, the density of gas in the galaxy is expected to be affected by this, which in turn will affect cooling and feedback. However, these effects are highly nonlinear and it is not obvious whether numerical effects can be imprinted on them.

Thus, while we could not identify a process that will give rise to a numerical artifact that can be learned by the machine-learning models, we cannot completely discard that possibility here.

5.4. Linear Information

On average, our models are able to constrain the value of Ω_m with a $\sim 10\%$ precision and an accuracy of ~ 0.03 for a single, generic galaxy. We may wonder whether there are enough modes in the Lagrangian region of those galaxies to achieve such accuracy. To provide an answer to this question, we consider a volume V and use the Fisher matrix formalism to quantify how much information that volume contains, considering it probes scales from $k_{\min} \sim 2\pi/V^{1/3}$ to $k_{\max} = 64 h \text{ Mpc}^{-1}$. The value of k_{\max} arises from the Nyquist frequency used to generate the initial conditions.

For our setup, the Fisher matrix can be computed as

$$F_{\alpha\beta} = \int_{k_{\min}}^{k_{\max}} V \frac{\partial \log P(k, \theta)}{\partial \theta_\alpha} \frac{\partial \log P(k, \theta)}{\partial \theta_\beta} \frac{k^2 dk}{(2\pi)^2} \quad (14)$$

where in our case $\theta = \{\Omega_m, \sigma_8\}$, $P(k, \theta)$ is the linear matter power spectrum and V the cosmological volume. The integral goes from $k_{\min} \sim 2\pi/V^{1/3}$ to k_{\max} .

In Figure 10, we show with a red solid line the marginalized constraints on Ω_m as a function of k_{\min} . As can be seen, to achieve an error on Ω_m below 0.033 we need a value of $k_{\min} \sim 2 h \text{ Mpc}^{-1}$ or a Lagrangian region of volume $\sim (\pi h^{-1} \text{ Mpc})^3$.

We expect the Lagrangian volume of most galaxies to be smaller than the above estimate (see, e.g., Oñorbe et al. 2014), indicating that the constraints from our models are better than the ones that can be obtained from a linear Gaussian field of the same volume. However, there are several caveats to this calculation. First of all, on scales smaller than $\sim 1 h \text{ Mpc}^{-1}$, the nonlinear matter density field may contain information not contained in the initial Gaussian field (Bayer et al. 2021a). Besides, our models include properties related to velocities (e.g., the galaxy peculiar velocity or the subhalo velocity dispersion) that can also provide additional information to the one contained at linear order. On top of this, on very small scales, cosmological modes are expected to be tightly coupled. Thus, even a relatively small Lagrangian region of a galaxy may be affected by modes larger than it, which could add additional information.

From this test we cannot draw any definitive conclusion on whether the constraints from our models are physical or they just reflect some nonphysical information arising from numerical artifacts.

5.5. Consequences

Our results suggest that galaxy properties would reside in different manifolds for different values of Ω_m . This in turn implies that it should be difficult, if not impossible, to reproduce the galaxy properties from real galaxies for cosmologies with a value of Ω_m far away from the true one. This is a clear prediction of this work that can be tested either using hydrodynamical simulations or semianalytic models.

Regarding hydrodynamic simulations, in CAMELS we vary four astrophysical parameters, while many others are kept fixed. In order to claim that Ω_m induces a distinct effect on galaxy properties, it is important to repeat the analyses carried out in this paper but sampling a much larger volume in parameter space where all astrophysical parameters are varied. This will allow us to investigate whether other astrophysical parameters may mimic the effect of Ω_m on galaxy properties.

On a side note, we note that galaxy properties are known to exhibit some level of intrinsic stochasticity (Genel et al. 2019) in numerical simulations. If our interpretation of the results is correct, this will imply that either the manifold containing the galaxy properties will have some intrinsic width or that galaxies affected by this effect will move along the manifold.

5.6. Future Work

In this work we have focused our attention on individual galaxies. In future work we will investigate the improvement on the parameter constraints when considering several galaxies instead of just one. We think that in this case the manifold where galaxies reside will be much better constrained, and therefore, we expect tighter constraints on all parameters. Furthermore, with many galaxies, it may be possible to extract information from different summary statistics (e.g., stellar-mass function) that may not be contained in the above manifolds.

While in this paper we have focused our attention on inferring the value of Ω_m from individual galaxies, in Appendix D we show that this method can also be used to infer the value of some astrophysical parameters. Given the accurate measurements of the value of the cosmological parameters from other methods, we may consider that this method may be used as a direct probe of astrophysical effects by fixing the value of the cosmological parameters. We will explore this direction in future work.

Most of the properties considered in this work can be measured from surveys. However, some of them, like the maximum circular velocity and the velocity dispersion, cannot be easily measured for large galaxy populations. In future work we plan to substitute those properties with others that can be measured, e.g., the velocity dispersion of the stars or neutral hydrogen, and investigate whether the value of Ω_m can still be inferred with them. It will also be interesting to quantify the accuracy gained by adding other properties not used in this work such as galaxy morphology, stellar age, mass in neutral or molecular gas, individual metal species, and environmental quantities like the overdensity of matter or galaxies in a given scale. Given the similarities we have outlined in this work with subhalo abundance matching,²⁰ it will be worth investigating whether our results improve even for nonobservational quantities, like the peak of the maximum circular velocity,

²⁰ For instance, our results suggest that V_{\max} and M_* are the most important variables. These variables play a crucial role in subhalo abundance matching.




Figure 10. We have used the Fisher matrix formalism to calculate how much information there is in the linear Gaussian density field for a cosmological volume V considering it contains modes from $k_{\min} = 2\pi/V^{1/3}$ to $k_{\max} = 64 h \text{ Mpc}^{-1}$. The solid red line shows the constraints on Ω_m as a function of k_{\min} , while the dashed black line displays the average error on Ω_m from our models. We can see that only volumes larger than $V \sim (3 h^{-1} \text{ Mpc})^3$ will contain enough modes to be able to place a constraint on Ω_m similar to, or better than, the one we obtain. We expect this volume to be larger than the Lagrangian region of most galaxies considered.

which is known to be better correlated with stellar mass than V_{\max} .

Further work is also needed to investigate the dependence of our results on the resolution, i.e., whether the manifold where galaxy properties live also changes with the mass and spatial resolution of the simulations and on the algorithm used to identify halos, subhalos, and galaxies. In similar terms, it is also important to assess the effect of supersample covariance on galaxy properties. While in this work we consider individual galaxies, which are therefore sensitive to local density fluctuations, the small CAMELS boxes do not account for large-scale fluctuations. Although we think this effect will not change our conclusions, it is important to provide a quantitative statement about this. We will also investigate whether our results still hold when using semianalytic models instead of hydrodynamic simulations; we will use CAMELS-SAM (L. Perez et al. 2022, in preparation) for this because their larger boxes will also allow us to investigate the mentioned supersample covariance effects.

Another important avenue to take in this work is the use of more interpretable machine-learning techniques, such as symbolic regression. These techniques are designed to provide analytic expressions between sets of variables, and their functional form may be easier to interpret than neural networks and gradient boosting trees. We note that we have used such techniques in this work but we were not able to obtain expressions accurate enough to capture the underlying relation. We thus leave this research direction for future work.

We believe that this work illustrates the complex interplay between cosmology and astrophysics on different physical scales (from galactic to cosmological) and how cosmological information may still be present within objects shaped by complex astrophysical processes such as galaxies. We also think this work shows how the use of machine-learning

techniques can help us better understand and disentangle complex physical processes and discover new features and techniques to maximize the information we can extract from the data.

To enable the community to reproduce our results, we release all data used in this work together with the codes, databases, and network weights obtained after training. We refer the reader to <https://github.com/franciscovillaescusa/Cosmo1gal> for further details.

We thank the anonymous referee for the constructive report. We thank Tom Abel, Arka Banerjee, Adrian Bayer, Greg Bryan, Neal Dalal, ChangHoon Hahn, Andrew Hearin, Lars Hernquist, Oliver Philcox, Tjitske Starkenburg, Michael Strauss, Masahiro Takada, and Benjamin Wandelt for useful conversations. We thank Uros Seljak for suggesting we perform the calculation in Section 5.4 and Volker Springel for correspondence that gave rise to Section 5.5. We have made use of the XGB²¹ and SHAP²² packages. The neural networks have been trained using GPUs at the Tiger cluster at Princeton University and the Rusty cluster of the Flatiron Institute. The work of F.V.N. is supported by the Simons Foundation. D.A.A. was supported in part by NSF grants AST-2009687 and AST-2108944. C.H. is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2121 Quantum Universe-390833306. All the data and codes used for this work are publicly available in <https://github.com/franciscovillaescusa/Cosmo1gal>. Details on the CAMELS simulations can be found in <https://www.camel-simulations.org>.

Appendix A Results for SIMBA Galaxies

In order to verify that our results hold for both IllustrisTNG and SIMBA galaxies, we have repeated the exercise in Section 3 and trained neural networks on individual properties of SIMBA galaxies to infer the values of the cosmological and astrophysical parameters.

We show the results in Figure 11. We find that, qualitatively, the results for SIMBA galaxies are the same as for IllustrisTNG galaxies. The model is able to infer the value of Ω_m with an accuracy of $\sim 3.7 \times 10^{-2}$ and a precision of 12%. We note that we observe a generic bias for true values of Ω_m below ~ 0.35 . This bias seems to be more severe for SIMBA galaxies than for IllustrisTNG galaxies, even when training on 14 variables (see Figure 4).

From the bottom panels of Figure 11 we can see that the network works for any generic galaxy, not a subset of them. As in the case of IllustrisTNG galaxies, we find a very small fraction of outliers. While the precision of the model when inferring Ω_m is very similar for all galaxies when the true value of Ω_m is intermediate or high, we find that the model is more precise when using massive galaxies for models with low values of Ω_m . This is similar to what we found for IllustrisTNG galaxies, although in that case the differences were even larger.

Overall, we conclude that we can use machine-learning methods to constrain the value of Ω_m independently of the

²¹ <https://xgboost.readthedocs.io>

²² <https://shap.readthedocs.io>




Figure 11. Same as Figure 2 but for SIMBA galaxies at $z = 0$.

simulation suite used to train the model. We emphasize, however, that our models are not robust (see Section 3.3).

Appendix B Robustness Test

In Section 3.3 we investigated the robustness of our models, finding that they are not robust. In other words, training the models on galaxies from one simulation suite does not allow us to infer the correct value of Ω_m from galaxies of the other simulation suite. In this appendix, we show a few more details of this test, investigating whether this is a generic feature for all galaxies or whether the model works in some cases.

We have trained a model using galaxies from the IllustrisTNG simulations at $z=0$ (using all properties except the magnitudes in the U , K , and g bands) and tested it on individual galaxies of the SIMBA simulations. In Figure 12 we show the results of performing the detailed analysis outlined in Section 3. As we already saw in Figure 4 we find that on average, the model is not able to infer the correct value of Ω_m

(top panel). We, however, perform a more detailed analysis of many individual galaxies and show the results in the bottom panels of Figure 12. As can be seen, in general, the model does not work for a generic galaxy. On the other hand, results are not completely off; for instance, see Figure 4 of Villaescusa-Navarro et al. (2021a) for a similar exercise with 2D maps. We find that the true value of Ω_m lies within the model standard deviation in a large fraction of galaxies, although there is obviously a large underlying bias.

We note that the model works better for cosmologies with low and high values of Ω_m and performs worse for intermediate values. This relative success, however, may just be an artifact, namely the network may be using information from priors. For the model with a true value of $\Omega_m \sim 0.27$, there is still a nonnegligible fraction of galaxies where the model seems to be working. This does not look like the fraction of outliers we have seen in all models in the main text. We defer to future work the exploration of the properties of these galaxies and whether they exhibit more similarities to the ones from the IllustrisTNG simulations.




Figure 12. Same as Figure 2 but for a model trained on IllustrisTNG galaxies and tested on SIMBA galaxies.

Appendix C SHAP Values

In order to identify the most important features of our networks, we have computed the SHapley Additive exPlanation (SHAP) value of each galaxy property. This method assigns to each feature of each galaxy a value; larger absolute values for a given property indicate that the feature is having a larger contribution to the final output of the model. In Figure 13 we show the distribution of SHAP values for the different features for the models trained on IllustrisTNG (left) and SIMBA (right) galaxies.

For the IllustrisTNG simulations we find that features such as stellar mass, K -band magnitude, gas mass, gas metallicity, and maximum circular velocity to be among the most important variables. For SIMBA we instead get properties

like total mass, stellar mass, maximum circular velocity, gas mass, and subhalo radius. In order to determine whether these variables are indeed the most important ones, we have retrained neural networks using as input those five variables instead of the 17/14 original ones from IllustrisTNG/SIMBA. However, the performance of the models trained on these variables is relatively poor, much worse than the variables identified in Section 4.2. We think that the reason behind this is that there are multiple variables that are highly correlated, and the model may be extracting information from them in a similar way. Under this condition, the SHAP values, while still reflecting the contribution of each variable to the model prediction, does not inform us of the minimum set of variables we are interested in in order to gain intuition on the physics behind the model.




Figure 13. In order to identify the most important variables used by the model in order to carry out its predictions, we have computed the SHAP (SHapley Additive exPlanation) values for each galaxy in the test set. The panels show the distribution of SHAP values for the galaxies of the IllustrisTNG (left) and SIMBA (right) simulations sorted by the different features. The color indicates the value of the variable from low (blue) to high (red). Larger absolute values indicate that the considered feature has a larger impact on the model’s final prediction.

Appendix D Constraining Astrophysical Parameters

In this paper we have focused on predicting the value of Ω_m . However, we saw in Figure 1 that our models have some constraining power on A_{SN1} and, to a lesser extent, on A_{SN2} . To investigate this more, we repeated the analysis outlined in Section 3 using IllustrisTNG galaxies at $z=0$ and show the results in Figures 14 and 15 for the parameters A_{SN1} and A_{SN2} , respectively.

We find that the model is able to infer the value of A_{SN1} with an accuracy of ~ 0.37 and a precision of $\sim 33\%$, while the constraints on A_{SN2} have an accuracy and precision of ~ 0.29 and $\sim 27\%$, respectively. We note that although the numbers are better for A_{SN2} , the visual inspection of the results reveals

that these are largely affected by priors and the model actually performs better on A_{SN1} .

When inspecting the results from individual galaxies more closely in the bottom panels of Figures 14 and 15, we find that the model performs relatively well for A_{SN1} in general, while for A_{SN2} we can see that in many cases the model is just predicting the mean value with large error bars, independently of galaxy type, cosmology, and astrophysics.

From this exercise we conclude that while the network is capable of using galaxy properties to infer the value of A_{SN1} with large error bars, it can barely say anything beyond predicting the mean value for A_{SN2} . We emphasize that the network cannot infer the value of the other parameters not mentioned in this appendix, i.e., A_{AGN1} , A_{AGN2} , and σ_8 .

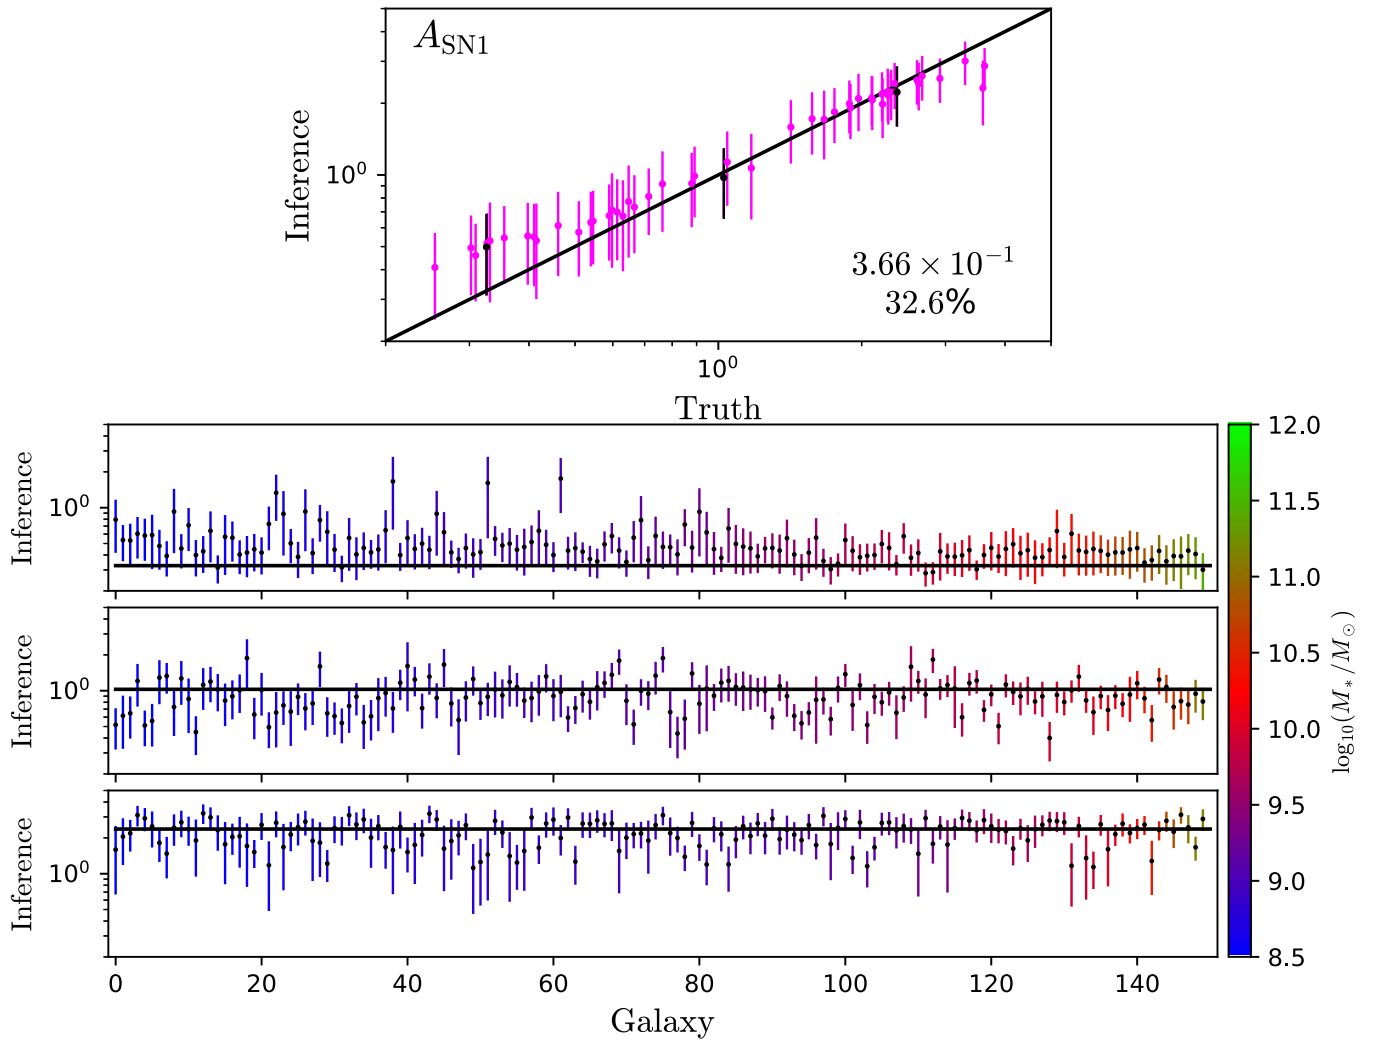


Figure 14. Same as Figure 2 but predicting the astrophysical parameter A_{SNI} from IllustrisTNG.

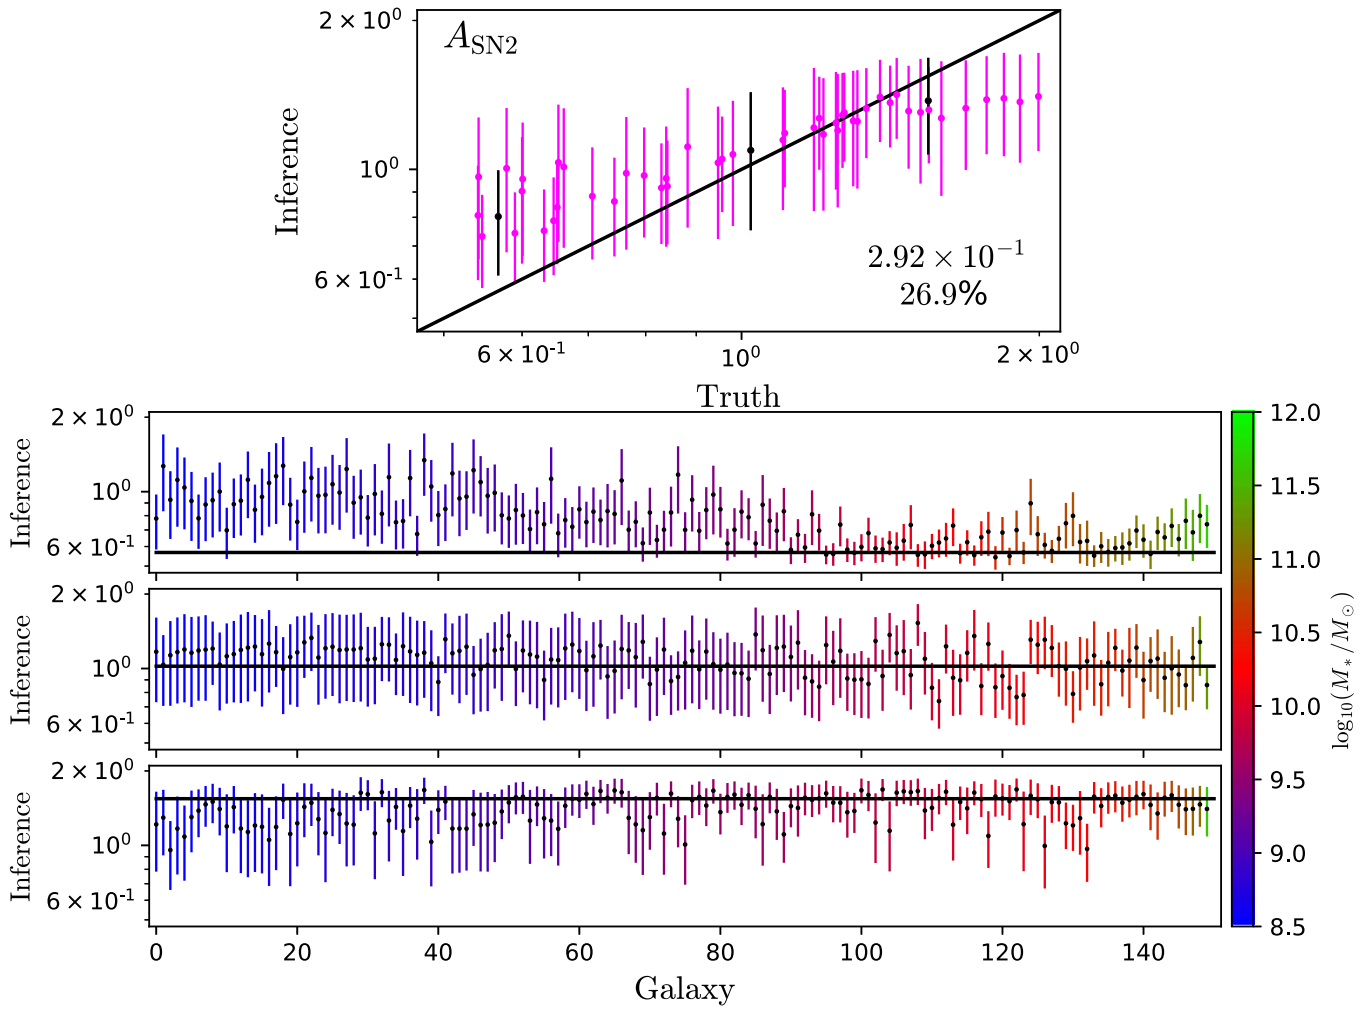


Figure 15. Same as Figure 2 but predicting the astrophysical parameter A_{SN2} from IllustrisTNG.

ORCID iDs

Francisco Villaescusa-Navarro <https://orcid.org/0000-0002-4816-0455>
 Shy Genel <https://orcid.org/0000-0002-3185-1540>
 Stephanie Tonnesen <https://orcid.org/0000-0002-8710-9206>
 David N. Spergel <https://orcid.org/0000-0002-5151-0006>
 Romain Teyssier <https://orcid.org/0000-0001-7689-0933>
 Yin Li <https://orcid.org/0000-0002-0701-1410>
 Caroline Heneka <https://orcid.org/0000-0001-8883-0583>
 Pablo Lemos <https://orcid.org/0000-0002-4728-8473>
 Daniel Anglés-Alcázar <https://orcid.org/0000-0001-5769-4945>
 Daisuke Nagai <https://orcid.org/0000-0002-6766-5942>

References

- Ajani, V., Peel, A., Pettorino, V., et al. 2020, *PhRvD*, **102**, 103531
 Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, arXiv:1907.10902
 Allys, E., Marchand, T., Cardoso, J. F., et al. 2020, *PhRvD*, **102**, 103506
 Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2017, *MNRAS*, **464**, 2840
 Banerjee, A., & Abel, T. 2021a, *MNRAS*, **500**, 5479
 Banerjee, A., & Abel, T. 2021b, *MNRAS*, **504**, 2911
 Banerjee, A., Castorina, E., Villaescusa-Navarro, F., Court, T., & Viel, M. 2020, *JCAP*, **2020**, 032
 Bayer, A. E., Banerjee, A., & Seljak, U. 2021a, arXiv:2108.04215
 Bayer, A. E., Villaescusa-Navarro, F., Massara, E., et al. 2021b, *ApJ*, **919**, 24
 Cheng, S., & Ménard, B. 2021, *MNRAS*, **507**, 1012
 Conroy, C., Wechsler, R. H., & Kravtsov, A. V. 2006, *ApJ*, **647**, 201
 Coulton, W. R., Liu, J., Madhavacheril, M. S., Böhm, V., & Spergel, D. N. 2019, *JCAP*, **2019**, 043
 Dai, J.-P., Verde, L., & Xia, J.-Q. 2020, *JCAP*, **2020**, 007
 Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, **486**, 2827
 Davé, R., Thompson, R., & Hopkins, P. F. 2016, *MNRAS*, **462**, 3265
 de la Bella, L. F., Tessore, N., & Bridle, S. 2021, *JCAP*, **2021**, 001
 Fluri, J., Kacprzak, T., Lucchi, A., et al. 2019, *PhRvD*, **100**, 063514
 Friedrich, O., Uhlemann, C., Villaescusa-Navarro, F., et al. 2020, *MNRAS*, **498**, 464
 Genel, S., Bryan, G. L., Springel, V., et al. 2019, *ApJ*, **871**, 21
 Giri, U., & Smith, K. M. 2020, arXiv:2010.07193
 Gualdi, D., Gil-Marín, H., & Verde, L. 2021a, *JCAP*, **2021**, 008
 Gualdi, D., Novell, S., Gil-Marín, H., & Verde, L. 2021b, *JCAP*, **2021**, 015
 Gupta, A., Matilla, J. M. Z., Hsu, D., & Haiman, Z. 2018, *PhRvD*, **97**, 103515
 Hahn, C., & Villaescusa-Navarro, F. 2021, *JCAP*, **2021**, 029
 Hahn, C., Villaescusa-Navarro, F., Castorina, E., & Scoccimarro, R. 2020, *JCAP*, **2020**, 040
 Harnois-Déraps, J., Martinet, N., Castro, T., et al. 2021, *MNRAS*, **506**, 1623
 Harnois-Déraps, J., Martinet, N., & Reischke, R. 2022, *MNRAS*, **509**, 3868
 Hassan, S., Andrianomena, S., & Doughty, C. 2020, *MNRAS*, **494**, 5761
 Hopkins, P. F. 2015, *MNRAS*, **450**, 53
 Jeffrey, N., Alsing, J., & Lanusse, F. 2021, *MNRAS*, **501**, 954
 Jeffrey, N., & Wandelt, B. D. 2020, arXiv:2011.05991
 Kuruvilla, J. 2021, arXiv:2109.13938
 Kuruvilla, J., & Aghanim, N. 2021, *A&A*, **653**, A130
 Lee, J., & Ryu, S. 2020, arXiv:2006.14477
 Le-Khac, P. H., Healy, G., & Smeaton, A. F. 2020, arXiv:2010.05113

- Li, Z., Liu, J., Zorrilla Matilla, J. M., & Coulton, W. R. 2019, *PhRvD*, **99**, 063527
- Liu, J., & Madhavacheril, M. S. 2019, *PhRvD*, **99**, 083508
- Lu, T., Haiman, Z., & Zorrilla Matilla, J. M. 2022, *MNRAS*, **511**, 1518
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, **480**, 5113
- Marques, G. A., Liu, J., Zorrilla Matilla, J. M., et al. 2019, *JCAP*, **2019**, 019
- Massara, E., Villaescusa-Navarro, F., Ho, S., Dalal, N., & Spergel, D. N. 2021, *PhRvL*, **126**, 011301
- Naidoo, K., Massara, E., & Lahav, O. 2021, arXiv:2111.12088
- Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **477**, 1206
- Neistein, E., Khochfar, S., Dalla Vecchia, C., & Schaye, J. 2012, *MNRAS*, **421**, 3579
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **475**, 624
- Ntampaka, M., Eisenstein, D. J., Yuan, S., & Garrison, L. H. 2020, *ApJ*, **889**, 151
- Oñorbe, J., Garrison-Kimmel, S., Maller, A. H., et al. 2014, *MNRAS*, **437**, 1894
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *ApJ*, **517**, 565
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018a, *MNRAS*, **475**, 648
- Pillepich, A., Springel, V., Nelson, D., et al. 2018b, *MNRAS*, **473**, 4077
- Porth, L., Bernstein, G. M., Smith, R. E., & Lee, A. J. 2021, arXiv:2111.13702
- Ravanbakhsh, S., Oliva, J., Fromenteau, S., et al. 2017, arXiv:1711.02033
- Ribli, D., Pataki, B. Á., Zorrilla Matilla, J. M., et al. 2019, *MNRAS*, **490**, 1843
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, **116**, 1009
- Ryu, S., & Lee, J. 2020, *ApJ*, **894**, 65
- Samushia, L., Slepian, Z., & Villaescusa-Navarro, F. 2021, *MNRAS*, **505**, 628
- Schmelzle, J., Lucchi, A., Kacprzak, T., et al. 2017, arXiv:1707.05167
- Shao, H., Villaescusa-Navarro, F., Genel, S., et al. 2022, *ApJ*, **927**, 85
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, **475**, 676
- Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, **328**, 726
- Uhlemann, C., Friedrich, O., Villaescusa-Navarro, F., Banerjee, A., & Codis, S. 2020, *MNRAS*, **495**, 4006
- Valogiannis, G., & Dvorkin, C. 2021, arXiv:2108.07821
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021a, arXiv:2109.09747
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021b, *ApJ*, **915**, 71
- Villaescusa-Navarro, F., Genel, S., Anglés-Alcázar, D., et al. 2021c, arXiv:2109.10360
- Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, *ApJS*, **250**, 2
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291
- Weinberger, R., Springel, V., & Pakmor, R. 2020, *ApJS*, **248**, 32
- White, S. D. M., Navarro, J. F., Evrard, A. E., & Frenk, C. S. 1993, *Natur*, **366**, 429
- Zhang, G., Li, Z., Liu, J., et al. 2020, *PhRvD*, **102**, 083537
- Zorrilla Matilla, J. M., Sharma, M., Hsu, D., & Haiman, Z. 2020, *PhRvD*, **102**, 123506