



Cosmology with Multiple Galaxies

Chaitanya Chawak¹, Francisco Villaescusa-Navarro^{2,3} , Nicolás Echeverri-Rojas⁴, Yueying Ni^{5,6} , ChangHoon Hahn³ , andDaniel Anglés-Alcázar^{2,7} ¹ Indian Institute of Science Education and Research (IISER) Tirupati, Tirupati-517507, India; chaitanyachawak001@gmail.com² Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA³ Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA⁴ Instituto de Física, Universidad de Antioquia, A.A.1226, Medellín, Colombia⁵ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA⁶ McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA⁷ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269, USA

Received 2023 October 2; revised 2024 April 27; accepted 2024 May 7; published 2024 July 4

Abstract

Recent works have discovered a relatively tight correlation between Ω_m and the properties of individual simulated galaxies. Because of this, it has been shown that constraints on Ω_m can be placed using the properties of individual galaxies while accounting for uncertainties in astrophysical processes such as feedback from supernovae and active galactic nuclei. In this work, we quantify whether using the properties of multiple galaxies simultaneously can tighten those constraints. For this, we train neural networks to perform likelihood-free inference on the value of two cosmological parameters (Ω_m and σ_8) and four astrophysical parameters using the properties of several galaxies from thousands of hydrodynamic simulations of the CAMELS project. We find that using properties of more than one galaxy increases the precision of the Ω_m inference. Furthermore, using multiple galaxies enables the inference of other parameters that were poorly constrained with one single galaxy. We show that the same subset of galaxy properties are responsible for the constraints on Ω_m from one and multiple galaxies. Finally, we quantify the robustness of the model and find that without identifying the model range of validity, the model does not perform well when tested on galaxies from other galaxy formation models.

Unified Astronomy Thesaurus concepts: Cosmological parameters (339); Galaxy properties (615); Galaxy processes (614); Computational methods (1965); Astronomy data analysis (1858)

1. Introduction

Some of the most fundamental questions we can ask in cosmology are: What are the components that make up the Universe? How much does each component contribute? We now know that the Universe should be made up of at least three main components: (1) baryons, representing all the substances and materials we know, (2) dark matter, some fundamental particle that interacts with baryons mostly (perhaps uniquely) through gravity, and (3) dark energy, a mysterious substance (perhaps a property of the vacuum) responsible for the recent acceleration of the Universe. From cosmological data, we believe these three components represent roughly 5%, 25%, and 70% of the current energy content of the Universe, respectively.

Parameters such as Ω_b and Ω_m represent the fraction of the Universe's energy content in terms of baryons and baryons plus dark matter, respectively. Determining them is important to learn about the nature and properties of dark matter and also to learn about the growth rate of the Universe (Huterer 2023). There are many different methods to infer these parameters, from studying the properties of the cosmic microwave background anisotropies to the spatial distribution of galaxies. Recently, Villaescusa-Navarro et al. (2022) claimed that a tight relation between Ω_m and the properties of individual galaxies is present in galaxies from state-of-the-art hydrodynamic simulations. The relationship is present even when varying the value

of astrophysical parameters controlling the efficiency of supernovae and active galactic nuclei (AGNs) feedback. Echeverri et al. (2023) reached the same conclusion when using galaxies generated with a different hydrodynamic and subgrid physics model.

Villaescusa-Navarro et al. (2022) discussed that such a relation might be due to the existence of a low-dimensional manifold where galaxy properties reside. In this view, changing Ω_m modifies the location of the galaxies in that manifold differently than changing the efficiency of astrophysical processes. For instance, increasing the value of Ω_m while keeping Ω_b fixed will increase the overall dark matter content of the Universe. That excess will enhance the dark matter content of galaxies, affecting their density, star formation rate, metallicity, etc. On the other hand, feedback can also affect some of these properties, but it is unlikely that it will significantly affect the dark matter content of most galaxies.

Villaescusa-Navarro et al. (2022) argued that knowing the location of one point in the manifold is enough to characterize it, and therefore, with one single galaxy, it is possible to infer the value of Ω_m . We note that Villaescusa-Navarro et al. (2022) and Echeverri et al. (2023) showed that Ω_m can be inferred with a $\sim 10\%$ precision based on the properties of a single galaxy, perhaps indicating that the manifold should have some intrinsic width associated with it. However, by using multiple galaxies, it should also be possible to infer the value of cosmological and astrophysical parameters by characterizing the impact on galaxy statistics like the stellar mass function. Recently, Busillo et al. (2023) have shown that galaxy scaling relations are sensitive to both cosmology and astrophysics and derived constraints on those from real data (see also Jo et al. 2023 for



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.

the impact on the star formation rate history and the stellar mass function). In this work, we thus ask ourselves how well we can infer cosmological parameters if we only have a few galaxies. Due to computational constraints, we limit our analysis to fewer than 10 galaxies. The training time in the case of two galaxies is approximately 40 hr when using a single NVIDIA A100 GPU. This increases as we increase the number of galaxies. We note that using properties from galaxies directly (instead of summary statistics) enables our models to search through all potential summary statistics and cross-correlations. Put simply, using details from individual galaxies rather than just summary numbers allows our models to explore all possible ways these details relate to each other.

In this paper, we show that using more than one galaxy increases the precision of the models trained to infer Ω_m , but at the same time, allows the models to infer other parameters that were unconstrained when using a single galaxy. To carry out our analysis, we made use of thousands of state-of-the-art hydrodynamic simulations from the Cosmology and Astrophysics with Machine Learning Simulations (CAMELS) project⁸ (Villaescusa-Navarro et al. 2021a, 2022; Ni et al. 2023). To check that our results do not hold just for galaxies generated by a particular code, we perform our analysis using simulations run with three codes that employ different subgrid physics models: (1) AREPO+IllustrisTNG, (2) GIZMO+SIMBA, and (3) MP-Gadget+Astrid.

This paper is organized as follows. We present the data we use and the machine-learning algorithms we employ in Section 2. In Section 3, we present the main results of our analysis. Finally, we summarize the takeaways and conclude in Section 4.

2. Methods

In this section, we first describe the data we use for this work. We then explain the machine-learning algorithms we employ to analyze the data and outline the metrics we utilize to quantify the accuracy and precision of our models.

2.1. Data

In this paper, we train neural networks to infer the value of cosmological and astrophysical parameters using the internal properties of simulated galaxies. These galaxies come from state-of-the-art hydrodynamic simulations of the CAMELS project (Villaescusa-Navarro et al. 2021a).

All simulations follow the nonlinear evolution of 256^3 dark matter plus 256^3 initial fluid elements from $z = 127$ down to $z = 0$ in a cubic periodic volume of $(25 h^{-1} \text{Mpc})^3$. All simulations share the value of these cosmological parameters: $\Omega_b = 0.049$, $h = 0.6711$, $n_s = 0.9624$, $w = -1$, $\Omega_K = 0$, and $\Sigma m_\nu = 0 \text{ eV}$.

The simulations have been run with three different codes and, therefore, can be classified into three different suites:

1. *IllustrisTNG*. The simulations in this suite have been run with the AREPO code (Springel 2010; Weinberger et al. 2019) and they employ the IllustrisTNG subgrid physics model (Pillepich et al. 2018; Nelson et al. 2019).
2. *SIMBA*. The simulations in this suite have been run with the GIZMO code (Hopkins 2015), and they employ the SIMBA subgrid physics model (Davé et al. 2019).

3. *Astrid*. The simulations in this suite have been run with the MP-Gadget code (Feng et al. 2018), and they employ a slightly modified version of the Astrid subgrid physics model (Bird et al. 2022; Ni et al. 2022).

Each suite contains 1000 simulations (from the Latin hypercube set of CAMELS). Each of those simulations has a different value of Ω_m , σ_8 , and four astrophysical parameters that control the efficiency of supernova and AGN feedback: A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2} . We refer the reader to Villaescusa-Navarro et al. (2021a) and Ni et al. (2023) for further details on the specifics of the astrophysical parameters. We emphasize that the astrophysical parameters have different meanings in each suite due to the different subgrid implementations, and they represent variations relative to the corresponding fiducial models of IllustrisTNG, SIMBA, and Astrid. Table 1 briefly describes the astrophysical and cosmological parameters involved in this study.

The value of these six parameters are arranged in a Latin hypercube with boundaries defined by

$$0.1 \leq \Omega_m \leq 0.5 \quad (1)$$

$$0.6 \leq \sigma_8 \leq 1.0 \quad (2)$$

$$0.25 \leq A_{\text{SN1}}, A_{\text{AGN1}} \leq 4.0 \quad (3)$$

$$0.5 \leq A_{\text{SN2}}, A_{\text{AGN2}} \leq 2.0 \quad (4)$$

We note that in the case of Astrid, the A_{AGN2} parameter ranges from 0.25 to 4. We also emphasize that all simulations have different values of the initial random seed. In this work, we focus our attention on the $z = 0$ snapshots of these simulations.

2.2. Galaxy Properties

Halos and subhalos are identified in the simulations using the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009). In this work, we define a galaxy as a subhalo with a stellar mass larger than zero. We follow Echeverri et al. (2023) and only consider galaxies with stellar masses above $5 \times 10^8 h^{-1} M_\odot$ to avoid working with small, likely spurious objects. SUBFIND computes many properties for each galaxy, but in this work, we focus our attention on the following 14:

1. M_g : the subhalo gas mass content, including the circumgalactic medium's contribution.
2. M_{BH} : the black hole (BH) mass of the galaxy.
3. M_* : the stellar mass of the galaxy.
4. M_t : the total mass of the subhalo hosting the galaxy.
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 velocity dispersion of all particles 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 (SFR).
10. J : the galaxy's subhalo spin vector modulus.
11. V : the modulus of the galaxy's subhalo peculiar velocity.
12. R_* : the radius containing half the galaxy's 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 $\sqrt{GM(<R_{\text{max}})}/R_{\text{max}} = V_{\text{max}}$.

For IllustrisTNG simulations, we also consider the following three properties:

1. U : the galaxy absolute magnitude yin the U band.

⁸ <https://www.camel-simulations.org/>

Table 1
Brief Description of the Cosmological and Astrophysical Parameters, and Their Ranges of Variation

Suite	Ω_m	σ_8	A_{SN1}	A_{SN2}	A_{AGN1}	A_{AGN2}
IllustrisTNG	Fraction of energy density in matter (dark matter + baryons)	Variance of the linear field on $8 h^{-1}\text{Mpc}$ at $z = 0$	Galactic winds: Energy per unit SFR	Galactic winds: wind speed	Kinetic mode BH feedback: energy per unit BH accretion rate	Kinetic mode BH feedback: ejection speed/burstiness
SIMBA	Fraction of energy density in matter (dark matter + baryons)	Variance of the linear field on $8 h^{-1}\text{Mpc}$ at $z = 0$	Galactic winds: Mass loading	Galactic winds: wind speed	Quasi-stellar object and jet-mode BH feedback: momentum flux	Jet-mode BH Black Hole feedback: jet speed
ω Astrid	Fraction of energy density in matter (dark matter + baryons)	Variance of the linear field on $8 h^{-1}\text{Mpc}$ at $z = 0$	Galactic winds: Energy per unit Star Formation Rate (SFR)	Galactic winds: wind speed	Kinetic mode Black Hole feedback: energy per unit BH accretion rate	Thermal mode Black Hole feedback: energy per unit BH accretion rate
	$\in [0.1, 0.5]$	$\in [0.6, 1.0]$	$\in [0.25, 4.0]$	$\in [0.5, 2.0]$	$\in [0.25, 4.0]$	IllustrisTNG and SIMBA $\in [0.5, 2.0]$, Astrid $\in [0.25, 4.0]$

2. K : the galaxy absolute magnitude in the K band.
3. g : the galaxy absolute magnitude in the g band.

We note that the above three magnitudes are not present in simulations of the SIMBA and Astrid suites because SUBFIND needs some particular properties not stored in those simulations to estimate the magnitudes. We refer the reader to Villaescusa-Navarro et al. (2022) for further details about these properties.

2.3. Input Data

The input to our models is a 1D vector containing the properties of n galaxies, where $n \in [1, 10]$. For instance, if we use galaxies from the Astrid suite and set $n = 5$, the input vector will contain $5 \times 14 = 70$ values. We remind the reader that for IllustrisTNG galaxies, we take 17 properties for each galaxy, while for SIMBA and Astrid, only 14 are available. Once the simulation suite and the value of n are chosen, we construct 1500 1D arrays with the properties of n unique galaxies (i.e., we enforce that the same galaxy cannot appear twice in the same set. It can, however, appear again in a different set). The 1500 1D arrays are constructed from the same simulation. The reason why we take 1500 arrays is that we have performed several convergence tests, and we find that increasing the number of 1D arrays during training does not yield noticeable improvements in our results.

2.4. Machine-learning Techniques

In this work, we train neural networks to perform likelihood-free inference on the value of two cosmological (Ω_m and σ_8) and four astrophysical (A_{SN1} , A_{SN2} , A_{AGN1} , A_{AGN2}) parameters.

Our models take as input a 1D vector containing the properties of n galaxies and return $2N_{\text{params}}$ numbers, where N_{params} is the number of parameters considered (e.g., $N_{\text{params}} = 1$ if only inferring one parameter). For each parameter i , our models output its marginal posterior mean (μ_i) and standard deviation (σ_i). This is achieved by minimizing the following loss function:

$$\mathcal{L} = \log \left[\sum_{j \in \text{batch}} (\theta_{i,j} - \mu_{i,j})^2 \right] + \log \left\{ \sum_{j \in \text{batch}} [(\theta_{i,j} - \mu_{i,j})^2 - \sigma_{i,j}^2]^2 \right\}. \quad (5)$$

This loss function guarantees that $\mu_{i,j}$ and $\sigma_{i,j}$ represent the parameter's posterior mean and standard deviation i (Jeffrey & Wandelt 2020; Villaescusa-Navarro et al. 2021b).

Our models use several blocks, each containing a fully connected layer, a LeakyReLU nonlinear activation function, and a dropout layer. After the last block, a fully connected layer predicts the network's output. We write our model in PyTorch.⁹ The number of blocks, the number of neurons in the fully connected layers, the learning rate, the weight decay, and the dropout rate are considered hyperparameters.

The value of the hyperparameters is tuned using Optuna¹⁰ (Akiba et al. 2019), which searches the hyperparameter space for the optimal values of the hyperparameters, which minimizes the value of the validation loss. We use at least 100 trials and

the optimization is done by searching the hyperparameter values that minimize the validation loss. We emphasize that we run Optuna for each different configuration; for instance, when changing the simulation suite or the number of galaxies, we retrain using Optuna to find the best hyperparameters for that case.

To train the models, we first split the simulations into training (850), validation (100), and testing (50) sets. We then construct the input 1D arrays by combining the properties of galaxies from the same simulation. We note that it is important to (1) avoid mixing galaxies from different simulations when combining galaxy properties into the input arrays since different simulations sample different parameter values and (2) avoid having galaxies from the same simulation in different sets (e.g., training and testing) since there could be leakage of information if galaxies from the same simulation are somehow correlated.

2.5. Performance Metrics

In this work, we use four metrics to quantify the accuracy and precision of our models. To use these metrics, we need to consider that for a given input 1D vector i , θ_i represents the value of the considered parameter, while μ_i and σ_i represent the posterior mean and standard deviation predicted by the network for that parameter. The four statistics we consider are:

1. Root mean squared error:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (\theta_i - \mu_i)^2}, \quad (6)$$

where the sum runs over all 1D arrays in the considered test set. Smaller values of the RMSE indicate the model is more accurate.

2. Mean relative error, ϵ :

$$\epsilon_i = \frac{1}{N} \sum_{i=1}^N \frac{\sigma_i}{\mu_i}, \quad (7)$$

the mean relative error tells us about the model's precision, with lower values representing, in general, more precise models. The mean relative error does not know anything about the true values. Thus, one can have a very precise but not accurate model.

3. Coefficient of determination, R^2 :

$$R^2(\theta_i, \mu_i) = 1 - \frac{\sum_{i=1}^n (\theta_i - \mu_i)^2}{\sum_{i=1}^n (\theta_i - \bar{\theta})^2}, \quad (8)$$

where, $\bar{\theta} = \sum_{i=1}^n \theta_i$. The R^2 quantifies the model's accuracy, with values close to 1 being accurate, and values close to 0 being poor.

4. Reduced chi-squared, χ^2 :

$$\chi^2 = \frac{1}{N} \sum_{i=1}^N \frac{(\theta_i - \mu_i)^2}{\sigma_i^2}. \quad (9)$$

We made use of these statistics to quantify the precision of the model error bars (posterior standard deviation). Values close to 1 indicate the size of the errors is appropriate, while values below/above 1 indicate the errors are over/underpredicted.

⁹ <https://pytorch.org/>

¹⁰ <http://optuna.org/>

3. Results

We now present the results of our analysis. We first show the results when training the models using the properties of two galaxies, and then we show the results when considering multiple galaxies. We note that Villaescusa-Navarro et al. (2022) and Echeverri et al. (2023) showed that Ω_m can be inferred with a $\sim 10\%$ precision based on the properties of a single galaxy, and the models could not constrain the values of other cosmological and astrophysical parameters within a meaningful margin of error.

3.1. Two Galaxies

We train models using 1D arrays that contain the properties of two galaxies. We then test those models on 1D arrays that contain the properties of two galaxies from the test set. We show the results in Figures 1 (IllustrisTNG and SIMBA), and 2 (Astrid). We find that all models can constrain the value of Ω_m accurately with $\{\text{RMSE}, \epsilon, R^2\}$ equal to $\{0.022, 0.077, 0.966\}$ (IllustrisTNG), $\{0.023, 0.090, 0.956\}$ (SIMBA), and $\{0.028, 0.094, 0.919\}$ (Astrid). We note that these numbers are better than the ones obtained for a single galaxy; for instance, these metrics are $\{0.0365, 0.11, 0.842\}$ when considering one single galaxy from Astrid (Echeverri et al. 2023).

On the other hand, σ_8 remains mostly unconstrained with two galaxies, irrespective of the simulation suite employed, in the same way as our findings for one galaxy (Villaescusa-Navarro et al. 2022; Echeverri et al. 2023). We reach similar conclusions for the AGN parameters of IllustrisTNG and SIMBA simulations. For Astrid, A_{AGN1} remains unconstrained, while A_{AGN2} can be inferred with a $\sim 16\%$ precision; a significant improvement from the $\sim 24\%$ obtained using one single galaxy (Echeverri et al. 2023). Finally, all models can infer the supernova feedback parameters with different precisions. We note that in the case of the supernova parameters, we have discarded a very small fraction of galaxies (0.41% for IllustrisTNG and 0.071% for SIMBA) since they have unreasonably small widths of the posterior, and therefore their χ^2 was really large and affected significantly the reported mean values.

These results show that better constraints on the value of the parameters can be achieved by using two galaxies instead of one.

3.2. Multiple Galaxies

Similar to the case of two galaxies, we also carried out the analysis for up to 10 galaxies considered simultaneously. We trained neural networks to perform likelihood-free inferences to estimate the values of the cosmological (Ω_m and σ_8) and astrophysical (A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2}) parameters using data from N galaxies (N goes from 1 to 10) from the IllustrisTNG, SIMBA, and the Astrid suites. Once trained, the model is tested using the galaxies from the test set for each case. Figure 3 shows how the prediction RMSE and R^2 of the cosmological and astrophysical parameters change as we increase the number of galaxies considered simultaneously. In this case, the quantity reported is the mean value of all galaxies in the test set. In Figure 7 in the Appendix, we show the results of training and testing using 10 galaxies (this figure is the equivalent to Figures 1 and 2).

We find that, as we consider more galaxies simultaneously, the predicted values of the cosmological parameters Ω_m and σ_8

become increasingly more accurate. In the case of the astrophysical parameters (A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2}), their predicted values can either improve or remain the same. The trend in the astrophysical parameters is not the same for all the suites because of the difference in the physical meaning of these parameters in each suite. For instance, the prediction of A_{AGN2} significantly improves when increasing from one galaxy to more than seven for the Astrid model, but remains poorly constrained regardless of the number of galaxies in the IllustrisTNG and SIMBA models. As we discuss below this may be related to the fact that with several galaxies, one can create a proxy for astrophysical quantities that are sensitive to feedback.

3.3. Only Ω_m

It is evident from the results discussed up to now that the model does an excellent job of predicting the value of Ω_m . So, we proceed to train the neural network to predict the posterior mean and standard deviation for only Ω_m instead of all six cosmological and astrophysical parameters. We do this so that the models can focus entirely on minimizing the loss for this parameter, avoiding situations where degeneracies with other parameters can yield suboptimal results for the parameter of interest.

In this case, these models do a slightly better job at inferring the value of Ω_m compared to results obtained when trained to predict all six cosmological and astrophysical parameters.

From Figure 4 we see that the neural network becomes increasingly more precise at inferring the value of Ω_m as we increase the number of galaxies considered simultaneously. For SIMBA and IllustrisTNG suites, the RMSE improves by about 55%, and in Astrid's case, it improves by 37% as we go from 1 galaxy to 10 galaxies. The right panel of Figure 4 shows the results when considering the R^2 statistics instead.

3.4. Most Important Features

Villaescusa-Navarro et al. (2022) carried out a feature importance study that showed that standard feature ranking methods (like computing saliency maps, using SHAP values, or using the inbuilt “feature importance” from scikit-learn) did not yield the important features that the model used to make inferences. This is due to strong internal correlations between galaxy properties, which makes it very difficult for the model to pinpoint the top properties. For that reason, Villaescusa-Navarro et al. (2022) trained a series of gradient-boosted trees models where one feature was discarded at a time. That way, the features could be ranked according to importance, and the results were sensitive. Echeverri et al. (2023) used the same procedure to rank the properties of the Astrid galaxies. Villaescusa-Navarro et al. (2022) and Echeverri et al. (2023) found that the five most important properties, according to their order of importance, for each of the suites are:

1. IllustrisTNG: $\{V_{\text{max}}, M_*, Z_*, R_*, K\}$
2. SIMBA: $\{V_{\text{max}}, M_*, R_{\text{max}}, Z_*, R_*\}$
3. Astrid: $\{M_*, Z_*, V_{\text{max}}, M_g, M_*\}$

With this information on hand, we now ask ourselves whether the constraints we obtain for Ω_m are mostly due to those variables or whether when considering multiple galaxies there may be information coming from other features.




Figure 1. We train neural networks to infer the value of the cosmological (Ω_m and σ_8) and astrophysical (A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2}) parameters from the internal properties of two random galaxies (without using galaxy positions). Next, from each simulation of the test set, we randomly select two galaxies and test the model on them. We show the results as points with error bars representing the posterior mean and the standard deviation (without making assumptions about the shape of the posterior). As can be seen, the models can precisely infer the value of Ω_m for both IllustrisTNG and SIMBA galaxies and, in some cases, the supernova feedback parameters. The value of σ_8 and the AGN parameters is poorly predicted in all cases.

2 galaxies : Astrid




Figure 2. Same as Figure 1 but for Astrid galaxies.

To answer this question, we train models that only use the above galaxy properties but consider multiple galaxies. We show the results in Figure 5 for the case of 10 galaxies. As we can see, when predicting only Ω_m , the model performs only $\sim 6\%$ worse than in the case of SIMBA, and $\sim 21\%$ worse than in cases of Astrid and IllustrisTNG, when compared to training with all properties. This is even after removing 12 galaxy properties in the case of IllustrisTNG (9 in the case of SIMBA and Astrid). We thus conclude that most of the information is contained in the most important variables for individual galaxies. We emphasize that this does not mean that the model uses information from individual galaxies and somehow stacks the results. Even using this subset of variables, one can construct noisy estimates of properties, like the stellar mass function, expected to be affected by cosmology (Jo et al. 2023). Therefore, the source of information may arise from both individual galaxies and collective properties.

3.5. Robustness

One of the most important aspects to consider when working with numerical simulations is the robustness of the results. In other words, how well the model behaves when training on galaxies from one galaxy formation model and testing it on galaxies from another galaxy formation model. This aspect has been investigated before in Villaescusa-Navarro et al. (2022), where it was found that even with a single galaxy, inference of Ω_m from galaxy properties was not robust. This claim was later revisited by Echeverri et al. (2023), who found that the lack of robustness was at least partially due to the presence of a small fraction of outliers. Echeverri et al. (2023) also found that removing these outliers during testing will make the model robust across simulations. In our case, as we train the models

using more than one galaxy, the chances of including the outliers increase, thus making the model more precise but less robust.

In order to verify the robustness of our model, we have considered the case where we train models that use 10 random galaxies produced with a given code and test them using 10 galaxies from another galaxy formation model. We show the results in Figure 6. As can be seen, the results are not robust and training models on galaxies from one simulation suite does not yield accurate results when testing on galaxies from another suite. To some extent, this is expected since constraints with multiple galaxies are tighter than with a single one and we have not removed, a priori, outliers from the cross-distributions. It is interesting to note that the worst case happens when training on Astrid and testing on IllustrisTNG. While we do not have a clear explanation for this, it may be related to the model focusing on aspects that are different among these two simulations. We leave it to future work to explore strategies designed to increase the robustness of the results following the findings of Echeverri et al. (2023). For instance, trying to identify outliers and remove them from the test set to improve the reliability of the predictions or perhaps using a generative model to be able to compute the likelihood of the galaxies directly.

4. Summary and Discussion

Previous works by Villaescusa-Navarro et al. (2022) and Echeverri et al. (2023) have pointed out the existence of a tight relation between the properties of individual simulated galaxies and Ω_m . The authors interpreted these results as a consequence of the existence of a manifold containing galaxy properties. Under that interpretation, properties of the manifold may




Figure 3. We train neural networks to infer the posterior mean and posterior standard deviation of all six parameters as a function of the number of galaxies. The top panels show the results for the RMSE, while the bottom panel displays the results for the R^2 statistics. In all cases, we show the average results, i.e., for a given simulation; we take 1500 different combinations and report the mean values. In general, the more galaxies we consider, the tighter the constraints on the parameters. However, there are some cases where constraints saturate, and adding more galaxies does not yield tighter constraints.

change in distinct manners when varying different parameters, allowing the inference of Ω_m from the properties of a single galaxy.

In this work, we have studied whether using the properties of several galaxies can help better constrain the value of the cosmological and astrophysical parameters. To investigate this,




Figure 4. We train the models to infer the value of Ω_m alone. The left and right panels show the mean values of the RMSE and R^2 as a function of the number of considered galaxies. We find our models perform slightly better when trained to infer Ω_m alone instead of inferring all six properties simultaneously. As can be seen, results improve when considering more galaxies, but in the case of Astrid, constraints tend to saturate when using more than ~ 5 galaxies.




Figure 5. We train models to infer the value of Ω_m using the properties of 10 galaxies. However, instead of using all galaxy properties, we made use of the five most important properties found when using a single galaxy. The panels show the average results for SIMBA (left), IllustrisTNG (middle), and Astrid (right). We found that constraints using just five galaxy properties are similar to the ones obtained using all galaxy properties. This may indicate that the models still extract information from the manifold containing galaxy properties and that information from noisy global quantities (e.g., stellar mass function) is subdominant.

we have trained neural networks to perform likelihood-free inferences on the values of the cosmological (Ω_m and σ_8) and astrophysical (A_{SN1} , A_{SN2} , A_{AGN1} , and A_{AGN2}) parameters by using the internal properties of multiple galaxies. We have made use of the properties of galaxies at redshift $z = 0$ from the IllustrisTNG, SIMBA, and the Astrid simulation suites from the CAMELS project (Villaescusa-Navarro et al. 2021a; Ni et al. 2023). We emphasize that our model only uses information from the galaxy properties, not their positions; in other words, the constraints do not incorporate any information from clustering.

We find that the precision of the predictions improves as we increase the number of galaxies. In the case of IllustrisTNG, SIMBA, and Astrid, the RMSE of Ω_m improves by factors of 2.5, 2.5, and 1.6, respectively. In the case of Astrid, we observe a plateau on the constraints when going beyond ~ 5 galaxies. For IllustrisTNG and SIMBA, the trend indicates that better constraints can be achieved using more than 10 galaxies. When considering the R^2 statistics, we find that the results tend to saturate when using more than ~ 5 galaxies. We have trained models to infer the value of Ω_m alone, i.e., without predicting the value of the other parameters. In this case, we find slightly

more precise results than when training the models to predict all parameters. The results, shown in Figure 4, do not change our conclusions.

For σ_8 , we find a steady improvement in the precision of the predictions (both RMSE and R^2) as we increase the number of galaxies. We emphasize that our models cannot determine the value of σ_8 with a single galaxy. The origin of these constraints may arise not from the properties of individual galaxies but from statistics that can be constructed when using multiple galaxies. For instance, the stellar mass function may be sensitive to the value of σ_8 , and a noisy version of it can be constructed when considering multiple galaxies. We thus speculate that the origin of this information may not be related to the manifold hosting the galaxy properties. We note that Busillo et al. (2023) have obtained cosmological and astrophysical constraints from the properties of a relatively small number of local, star-forming, galaxies.

For the supernova feedback parameters, we also find a consistent improvement in the constraints as we increase the number of galaxies for all suites with the exception of A_{SN2} for Astrid where constraints seem to saturate for more than ~ 5 galaxies. We believe the explanation may be related to the

10 galaxies




Figure 6. Robustness test. We have trained models to infer the value of Ω_m using the properties of 10 galaxies from SIMBA (top row), IllustrisTNG (middle row), and Astrid (bottom row). We then test the models on properties from 10 galaxies from SIMBA (left column), IllustrisTNG (middle column), and Astrid (right column). As can be seen, the models are not robust, and they fail when tested on galaxies from models different from the ones used for training.

previous argument, i.e., with multiple galaxies, one can construct noisy estimates of global statistical properties that may be sensitive to these parameters, like the stellar mass function or the stellar metallicity relation. However, differently to σ_8 , even with a single galaxy we find some constraining power on the value of these parameters, so it can just be that results are just exploiting that to determine the shape of the manifold better. Both factors likely came into play in this setup.

Finally, we find no constraining power for the AGN parameters for A_{AGN1} for the galaxies in IllustrisTNG and Astrid and a modest improvement of $\sim 20\%$ for SIMBA. On the other hand, for A_{AGN2} , the constraints for IllustrisTNG and SIMBA do not improve up to 10 galaxies, while for Astrid there is a significant $\sim 2.5 \times$ improvement in the RMSE value. We note that AGN feedback is expected to have a larger effect on massive galaxies, so the fact that we choose galaxies

randomly (making it more likely to choose small galaxies) can be the reason behind this behavior.

Our results indicate that the models may still be exploiting the information contained on the most important variables used when inferring Ω_m with a single galaxy. In this case, the improvement may be due to a better determination of the galaxy manifold (stacking results for individual galaxies) but also to the impact of cosmology and astrophysics on quantities such as the stellar mass function, where noisy versions of it can be constructed from a set of galaxies. It is however interesting to see that galaxy properties not important for constraints on individual galaxies do not seem to have an impact also when using catalogs.

As expected, our models become more precise but less accurate as we increase the number of galaxies. The reason is that the models are not robust even when considering a single

galaxy (Villaescusa-Navarro et al. 2022). However, we note that Echeverri et al. (2023) found that the models fail, on average, due to the presence of outliers. We thus leave it to future work to tackle the robustness of the models for one and multiple galaxies.

Finally, it is important to compare the results of this work versus those in Hahn et al. (2023), which are based on the core idea of utilizing the impact of cosmology of galaxy properties. That work provides the first constraints on Ω_m and σ_8 obtained from the photometry alone of thousands of NASA-Sloan Atlas¹¹ (NSA) galaxies. The NSA is a catalog of images and parameters of local galaxies, from surveys in the ultraviolet, optical, and near-infrared bands. The NSA provides photometry of $z < 0.05$ galaxies observed by the Sloan Digital Sky Survey. In that work, it is found that adding more galaxies improves the constraints, while here, we find that constraints tend to saturate when considering multiple galaxies. However, in Hahn et al. (2023), the information is not extracted from noiseless galaxy properties but from noisy and dust-attenuated photometry. In that case, even at the level of a single galaxy constraints are poorer than the ones reported here. This is because some information is lost when using photometry instead of galaxy properties. Thus, it is not surprising that stacking thousands of galaxies yields better constraints when using photometry than the ones obtained when using a few galaxies but knowing their properties without errors.

We conclude that better constraints on the value of the cosmological and astrophysical parameters can be obtained by using the properties of multiple galaxies instead of one. In this case, a combination of better knowing the underlying manifold hosting the data and the possibility of constructing noisy

estimates of global quantities is behind the performance of our results. It would be interesting to investigate whether some particular combinations of galaxies yield tighter constraints and, therefore, maximize the information content. That selection should also account for the robustness of the model. We leave all this for future work.

Acknowledgments

The neural networks have been trained using the GPUs from the Rusty cluster at the Flatiron Institute. We thank the anonymous referee for the constructive report. We thank Arnab Lahiry, Natali de Santi, and Helen Shao for useful conversations. C.C. thanks IISER Tirupati for the opportunity to carry out his Master's thesis remotely. The work of F.V.N. is supported by the Simons Foundation. The CAMELS project is supported by the Simons Foundation and the NSF grant AST-2108078. D.A.A. acknowledges support by NSF grants AST-2009687 and AST-2108944, CXO grant TM2-23006X, Simons Foundation Award CCA-1018464, and Cottrell Scholar Award CS-CSA-2023-028 by the Research Corporation for Science Advancement. Details about the CAMELS simulations can be found at <https://www.camel-simulations.org>.

Appendix Ten Galaxies

Figure 7 shows the results we obtained when training networks using the properties of 10 galaxies. As can be seen, these results are systematically better than the ones we obtained using two galaxies (see Figures 1 and 2).

¹¹ <http://www.nsatlas.org/>




Figure 7. Same as Figures 1 and 2 but using 10 galaxies instead of 2.

ORCID iDs

Francisco Villaescusa-Navarro  <https://orcid.org/0000-0002-4816-0455>

Yueying Ni  <https://orcid.org/0000-0001-7899-7195>

ChangHoon Hahn  <https://orcid.org/0000-0003-1197-0902>

Daniel Anglés-Alcázar  <https://orcid.org/0000-0001-5769-4945>

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, in Proc. 25th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (New York, NY: Association for Computing Machinery), 2623
- Bird, S., Ni, Y., Matteo, T. D., et al. 2022, *MNRAS*, **512**, 3703
- Busillo, V., Tortora, C., Napolitano, N. R., et al. 2023, *MNRAS*, **525**, 6191
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, **486**, 2827
- Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, *MNRAS*, **399**, 497
- Echeverri, N., Villaescusa-Navarro, F., Chawak, C., et al. 2023, *ApJ*, **954**, 125
- Feng, Y., Bird, S., Anderson, L., Font-Ribera, A., & Pedersen, C. 2018, MP-Gadget/MP-Gadget: A Tag for Getting a DOI, FirstDOI, Zenodo, doi:10.5281/zenodo.1451799
- Hahn, C., Villaescusa-Navarro, F., Melchior, P., & Teyssier, R. 2023, in Machine Learning for Astrophysics. Workshop at the Fortieth International Conference on Machine Learning (ICML 2023), 14
- Hopkins, P. F. 2015, *MNRAS*, **450**, 53
- Huterer, D. 2023, *A&ARv*, **31**, 2
- Jeffrey, N., & Wandelt, B. D. 2020, arXiv:2011.05991
- Jo, Y., Genel, S., Wandelt, B., et al. 2023, *ApJ*, **944**, 67
- Nelson, D., Springel, V., Pillepich, A., et al. 2019, *ComAC*, **6**, 2
- Ni, Y., Genel, S., Anglés-Alcázar, D., et al. 2023, *ApJ*, **959**, 136
- Ni, Y., Matteo, T. D., Bird, S., et al. 2022, *MNRAS*, **513**, 670
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018, *MNRAS*, **475**, 648
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, **328**, 726
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021a, *ApJ*, **915**, 71
- Villaescusa-Navarro, F., Ding, J., Genel, S., et al. 2022, *ApJ*, **929**, 132
- Villaescusa-Navarro, F., Genel, S., Angles-Alcazar, D., et al. 2021b, *ApJS*, **259**, 61
- Villaescusa-Navarro, F., Genel, S., Anglés-Alcázar, D., et al. 2022, *ApJS*, **265**, 54
- Weinberger, R., Springel, V., & Pakmor, R. 2019, *ApJS*, **248**, 32