



Finding Universal Relations in Subhalo Properties with Artificial Intelligence

Helen Shao¹, Francisco Villaescusa-Navarro^{1,2}, Shy Genel^{2,3}, David N. Spergel^{1,2}, Daniel Anglés-Alcázar^{2,4},
Lars Hernquist⁵, Romeel Dave^{6,7,8}, Desika Narayanan^{9,10}, Gabriella Contardo², and Mark Vogelsberger¹¹

¹ Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA; hshao@princeton.edu, villaescusa.francisco@gmail.com

² Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

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

⁴ Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269, USA

⁵ Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA

⁶ Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh, EH9 3HJ, UK

⁷ Department of Physics & Astronomy, University of the Western Cape, Cape Town 7535, South Africa

⁸ South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa

⁹ Department of Astronomy, University of Florida, Gainesville, FL 32611-2055, USA

¹⁰ University of Florida Informatics Institute, 432 Newell Drive, CISE Bldg E251, Gainesville, FL 32611-5585, USA

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

Received 2021 November 12; revised 2022 January 15; accepted 2022 January 18; published 2022 March 8

Abstract

We use a generic formalism designed to search for relations in high-dimensional spaces to determine if the total mass of a subhalo can be predicted from other internal properties such as velocity dispersion, radius, or star formation rate. We train neural networks using data from the Cosmology and Astrophysics with Machine Learning Simulations project and show that the model can predict the total mass of a subhalo with high accuracy: more than 99% of the subhalos have a predicted mass within 0.2 dex of their true value. The networks exhibit surprising extrapolation properties, being able to accurately predict the total mass of any type of subhalo containing any kind of galaxy at any redshift from simulations with different cosmologies, astrophysics models, subgrid physics, volumes, and resolutions, indicating that the network may have found a universal relation. We then use different methods to find equations that approximate the relation found by the networks and derive new analytic expressions that predict the total mass of a subhalo from its radius, velocity dispersion, and maximum circular velocity. We show that in some regimes, the analytic expressions are more accurate than the neural networks. The relation found by the neural network and approximated by the analytic equation bear similarities to the virial theorem.

Unified Astronomy Thesaurus concepts: Large-scale structure of the universe (902); Astrostatistics (1882); Galactic and extragalactic astronomy (563); Computational methods (1965)

1. Introduction

Galaxies are fascinating objects that host dark matter, gas, stars, and black holes. These elements interact in rich and complex manners whose details we do not yet fully understand (Somerville & Davé 2015). Improving our knowledge on the processes and the physics driving galaxy formation and evolution is not only important for galactic physics, but will also benefit other fields such as cosmology, where the uncertainty in galactic physics represents a major obstacle in extracting fundamental physics information from cosmic surveys (Villaescusa-Navarro et al. 2021).

Galaxies are characterized by many different properties, such as stellar mass, gas metallicity, neutral hydrogen mass, and luminosity in a given band. However, most of these properties are not independent and there are well-known correlations among them. In some cases, the correlations are tight enough to become actual relations such as the Tully–Fisher relation (Tully et al. 1975), the Faber–Jackson relation (Faber & Jackson 1976), and the fundamental plane defined by velocity dispersion, effective radius, and effective surface brightness (Dressler et al. 1987; Bender et al. 1992; Burstein et al. 1997). Some of these correlations and relations are induced by physical mechanisms, and therefore their existence reflects the underlying laws governing a particular process. In that

regard, it is important to search for such relationships in high-dimensional spaces, like the one concerning galaxies, as those relations can help us learn about the underlying physics. Unfortunately, it is not easy to find relations in high-dimensional spaces. On the other hand, machine-learning techniques can perform this task in a relatively straightforward way.

In this work we outline a generic methodology that can be used to search for such relations in high-dimensional spaces and apply it to the case of the properties of subhalos and the galaxies within them. We identify a potentially new and universal relation between the total mass of a subhalo and other internal properties such as radius, velocity dispersion, and gas metallicity by training neural networks on subhalo properties from state-of-the-art hydrodynamic simulations. We then derive a new analytic expression that is able to parameterize the found relationship in an accurate way.

This paper is organized as follows. In Section 2 we describe the data and the machine-learning algorithm we use. We then present our findings in Section 3. Finally, we draw the main conclusions of this work in Section 4.

2. Methods

In this work we search for a relation between the total mass of a subhalo and its other properties. In this section we first outline the generic methodology we employ to search for relationships between different variables in high-dimensional spaces and how to identify analytic expressions that approximate them. We then



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.

describe in detail the ingredients needed to carry out this task: the data used, the neural network architecture and training procedure, and the method used to find the analytic equations.

2.1. Methodology

The generic idea behind this method is to find a relation between a given variable, y , and a set of other variables $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$, i.e.,

$$y = f(\mathbf{x}). \quad (1)$$

In the first step, one can train a neural network to approximate the function f . By testing the accuracy of the network on the test set, one can quantify the accuracy reached by the model. If the accuracy is high, there may be an actual relationship between the variables. One can then use methods such as saliency maps to identify the most important variables that contribute to the relationship. These values can be obtained by computing the average of the gradients of the neural network outputs with respect to the input variable. By training neural networks using only the most important variables, one can check if the model extrapolates better or not than the model trained on all properties; in other words, whether the model may be extracting spurious information from some variables. Next, if the set of variables that contribute the most and achieve a good accuracy is small enough, one can use techniques such as symbolic regression to approximate the found relationship with analytic expressions. If the derived expression is not accurate enough, one can try with simpler models (e.g., power laws) guided by physical principles (if possible). For instance, the symbolic-regression algorithm may find equations that, although not perfect, may capture the main trend, and one can improve on them based on physical principles or by adding additional dependencies not captured by the model. We note that a similar methodology was also used in Wadekar et al. (2021) to improve the connection between halos and their neutral hydrogen mass.

We apply the above scheme to internal subhalo properties from state-of-the-art hydrodynamic simulations, and derive a new relation between the total mass of a subhalo and its other internal properties and also the properties of the galaxy it hosts. We then provide an equation that approximates such a relation and attempt to understand the physics behind it.

2.2. Data

We made use of data from IllustrisTNG100 and IllustrisTNG300 simulations (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018a; Springel et al. 2018; Nelson et al. 2019) together with the CAMELS project (Villaescusa-Navarro et al. 2021), which we briefly describe now.

1. *IllustrisTNG100*. A state-of-the-art magneto-hydrodynamic simulation run with the moving mesh AREPO code (Weinberger et al. 2020). The simulation samples a periodic comoving volume of $(75 h^{-1}\text{Mpc})^3$ using 1820^3 dark matter particles and 1820^3 fluid elements down to $z = 0$. The details of the subgrid galaxy formation model used in this simulation can be found in Pillepich et al. (2018b) and Weinberger et al. (2017), including supernova-driven galactic winds and active galactic nuclei (AGN) feedback. Out of the three simulations run at

different resolutions, we use the one with the highest resolution: IllustrisTNG100-1. This simulation has the highest mass and spatial resolution of the ones considered in this work. We use this simulation to quantify how our networks and analytic expressions behave for subhalos from a simulation with higher resolution and larger volume than the ones they have been trained on.

2. *IllustrisTNG300*. This simulation is identical to IllustrisTNG100 with the only difference being its volume, $(205 h^{-1}\text{Mpc})^3$ and its resolution, with 2500^3 dark matter particles plus 2500^3 fluid elements. Out of the three simulations run at different resolutions, we use the one with the highest resolution: IllustrisTNG300-1. We use this simulation to quantify how our networks and analytic expressions behave for subhalos from a simulation with a much larger volume and slightly different resolution than the ones they have been trained on.
3. *CAMELS-IllustrisTNG*. A suite of 1000 magneto-hydrodynamic simulations run with the AREPO code and using the same subgrid model as the IllustrisTNG simulations above. Each simulation follows the evolution of 256^3 dark matter particles plus 256^3 fluid elements in a periodic comoving volume of $(25 h^{-1}\text{Mpc})^3$ down to $z = 0$. Each of the 1000 simulations has a different cosmology (varying Ω_m and σ_8) but also astrophysics by varying the value of four astrophysical parameters controlling the efficiency of supernova and AGN feedback. These simulations are part of the CAMELS project (Villaescusa-Navarro et al. 2021). We use the subhalos from these simulations to train the neural networks.
4. *CAMELS-SIMBA*. A suite of 1000 hydrodynamic simulations run with the GIZMO code (Hopkins 2015) employing the same subgrid model as the SIMBA simulation (Davé et al. 2019). These simulations are comparable to CAMELS-IllustrisTNG in the sense that each simulation follows the evolution of 256^3 dark matter particles plus 256^3 fluid elements in a periodic comoving volume of $(25 h^{-1}\text{Mpc})^3$ down to $z = 0$. Each of the 1000 simulations has a different cosmology (varying Ω_m and σ_8) but also astrophysics by varying the value of four astrophysical parameters controlling the efficiency of supernova and AGN feedback. We emphasize that the subgrid physics of SIMBA is very different to that of IllustrisTNG, including parameterized galactic winds based on higher resolution FIRE simulations (Muratov et al. 2015; Anglés-Alcázar et al. 2017b) and gravitational torque-driven black hole growth coupled to kinetic outflows (Anglés-Alcázar et al. 2017a). These simulations are part of the CAMELS project (Villaescusa-Navarro et al. 2021). We use the subhalos of these simulations to quantify the accuracy of our networks and analytic expressions when varying the subgrid physics and method used to solve the hydrodynamic equations.
5. *CAMELS-Nbody*. Each hydrodynamic simulation belonging to CAMELS-IllustrisTNG and CAMELS-SIMBA has an associated gravity-only N -body simulation with the same cosmological model and the same initial random phase as its hydrodynamic counterpart. These simulations thus follow the evolution of 256^3 dark matter particles in a periodic comoving volume of $(25 h^{-1}\text{Mpc})^3$ down to $z = 0$. These simulations only account for the

Table 1
The Simulations Used, the Number of Subhalos they Contain, and their Maximum and Minimum Subhalo Masses

Simulation Name	Number of Subhalos	$M_{\text{tot}} (h^{-1} M_{\odot})$	Description
IllustrisTNG100	88,507	$[8.30 \times 10^6 - 2.75 \times 10^{14}]$	A hydrodynamic simulation run with the AREPO code. The simulation has a volume of $(75 h^{-1} \text{Mpc})^3$.
IllustrisTNG300	515,600	$[6.89 \times 10^7 - 1.28 \times 10^{15}]$	A hydrodynamic simulation run with the AREPO code. The simulation has a volume of $(205 h^{-1} \text{Mpc})^3$.
CAMELS-IllustrisTNG	720,548	$[1.18 \times 10^8 - 3.89 \times 10^{14}]$	A set of 1000 simulations run with the AREPO code employing the same subgrid physics as the IllustrisTNG simulations. Each simulation has a different value of the cosmological and astrophysical parameters. Each simulation has a volume of $(25 h^{-1} \text{Mpc})^3$.
CAMELS-SIMBA	1,182,265	$[1.55 \times 10^8 - 5.36 \times 10^{14}]$	A set of 1000 simulations run with the GIZMO code employing the same subgrid physics as the SIMBA simulations. Each simulation has a different value of the cosmological and astrophysical parameters. Each simulation has a volume of $(25 h^{-1} \text{Mpc})^3$.
CAMELS-Nbody	16,965,513	$[5.44 \times 10^8 - 3.93 \times 10^{14}]$	A set of 2000 N -body simulations run with the Gadget-III code. Each simulation has a different value of the cosmological parameters, and they represent the gravity-only counterpart of the CAMELS-IllustrisTNG and CAMELS-SIMBA simulations. Each simulation has a volume of $(25 h^{-1} \text{Mpc})^3$.

Note. For the hydrodynamic simulations, we only consider subhalos that contain more than 20 star particles. For the N -body simulations, we only consider subhalos that contain more than 20 dark matter particles.

gravitational forces and therefore do not model galaxy formation physics. We use these simulations to check if the findings from the above hydrodynamic simulations also apply to subhalos from N -body simulations.

Table 1 summarizes the main characteristics of the different simulations used in this work. Each simulation has a series of halo and subhalo catalogs associated to it that were generated by running SUBFIND (Springel et al. 2001) on the corresponding snapshots at different redshifts. SUBFIND works primarily by identifying local peaks in the three-dimensional density field and separating them by identifying a saddle point between them. In the second step, the overdensities and their surroundings are checked for gravitational self-boundness: those that are self-bound are registered as subhalos, and those that are not are attached to their neighboring overdensities, namely those they share saddle points with. SUBFIND operates on all particle types in the simulations, dark matter and baryonic alike. In this work we focus our attention on subhalos that contain *galaxies*, i.e., SUBFIND subhalos that contain more than 20 star particles. For each subhalo we consider 12 different quantities including both properties of the subhalo itself and of its central galaxy:

1. *Total mass*, M_{tot} . This quantity represents the total mass of the subhalo, i.e., the sum of the mass in gas, dark matter, stars, and black holes.
2. *Black hole mass*, M_{BH} . This quantity is the black hole mass of the subhalo.
3. *Stellar mass*, $M_{\star}$. This quantity is the stellar mass of the subhalo.
4. *Gas mass*, M_{g} . This quantity represents the gas mass of the subhalo.
5. *Spin*, J . This quantity is the modulus of the subhalo three-dimensional spin vector.

6. *Velocity*, V . The quantity is the modulus of the peculiar velocity vector of the subhalo.
7. *Gas metallicity*, Z_{g} . This quantity represents the average metallicity of all gas particles within twice the radius containing half of the subhalo stellar mass.
8. *Stars metallicity*, $Z_{\star}$. This quantity is the average metallicity of all star particles within twice the radius containing half of the subhalo stellar mass.
9. *Radius*, R . This quantity is the comoving radius containing half of the subhalo total mass.
10. *Star formation rate*, SFR. This quantity is the total star formation rate of all gas particles in the subhalo.
11. *Velocity dispersion*, σ . This quantity represents the one-dimensional velocity dispersion of all particles in the subhalo.
12. *Maximum velocity curve*, V_{max} . This quantity is the maximum of the spherically averaged rotation curve, defined as $\sqrt{GM_{\text{tot}}(<r)}/r$.

We note that some of the above properties can be associated with the galaxy inside the subhalo, like the black hole mass, stellar mass, and star formation rate. Others, like the gas mass, can be associated with the galaxy but also its circumgalactic medium. We emphasize that although we generically call all these objects subhalos, some of them are may be very close to actual halos. For instance, the most massive subhalo of a halo tends to represent the main body of the dark matter halo without including other subhalos and their galaxies. Thus, while we do not distinguish between central and satellites in our nomenclature (but see Appendix A), it is important to keep this in mind when associating our subhalos to the physical systems behind them. Moreover, most of these properties are not observable, including the total mass and velocity dispersion of the subhalos. The only property that can be approximately estimated is the maximum circular velocity.

2.3. Neural Networks

The main goal of this paper is to find a relation between the total mass of a subhalo, M_{tot} , and its other properties, i.e.,

$$M_{\text{tot}} = f(\vec{\theta}), \quad (2)$$

where $\vec{\theta}$ is the vector with the other subhalo properties; $\vec{\theta}$ can be all other subhalo properties or a subset of it. Neural networks can be used to approximate the function f . When training a model on a given data set, we first split its data into training (80%), validation (10%), and testing (10%). Some quantities exhibit a very large dynamical range. This can be a problem when training the networks. Thus, we first redefine these variables such as

1. $M_{\text{tot}} \rightarrow \log_{10}(1 + M_{\text{tot}})$,
2. $M_* \rightarrow \log_{10}(1 + M_*)$,
3. $M_{\text{BH}} \rightarrow \log_{10}(1 + M_{\text{BH}})$,
4. $M_g \rightarrow \log_{10}(1 + M_g)$.

Before taking the logarithm we have added 1 to the variables to avoid problems when the value of those variables is equal to 0. Note that we are implicitly assuming that each mass term is divided by $h^{-1} M_{\odot}$. We then standardize all variables using

$$\tilde{x} = \frac{x - \mu}{\delta}, \quad (3)$$

where μ is the mean and δ is the standard deviation of each variable, x .

The architecture of our model consists of a set of fully connected layers whose input is the vector with the considered normalized variables and the output is a single number with the predicted value of the normalized subhalo total mass, $\tilde{M}_{\text{tot}}$. Between each fully connected layer we include a LeakyReLU nonlinear activation function with a negative slope value of 0.2.

The loss function we optimize via gradient descent is the standard mean squared error:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (\tilde{M}_{\text{tot}}^{\text{true}} - \tilde{M}_{\text{tot}}^{\text{pred}})^2. \quad (4)$$

We use the AdamW optimizer (Loshchilov & Hutter 2017) with beta values equal to 0.9 and 0.999. We train the network using a batch size of 256 for 500 epochs. We use PYTORCH¹² to train, validate, and test the networks.

The hyperparameters of our model are (1) the number of fully connected layers, (2) the number of neurons per layer, (3) the value of the learning rate, and (4) the weight decay. We use the OPTUNA code (Akiba et al. 2019) to perform Bayesian optimization and find the best value of these hyperparameters for each case we consider (e.g., when considering all subhalo properties or when considering only a subset of them). For each case we run 100 trials, where each trial consists of training the model using selected values of the hyperparameters. We perform optimization of the hyperparameters with the requirement to achieve the lowest validation loss possible. The model we choose in the end is the one with the lowest validation loss found by OPTUNA. The optimal model that OPTUNA found contains five hidden layers and each layer consists of 171, 236, 884, 266, and 32 neurons, respectively. The ultimate value of the weight decay is 1.1×10^{-5} and the learning rate is

1.1×10^{-4} . Moreover, we did not tune batch size as a hyperparameter and chose its value to be 256.

2.4. Symbolic Regression

While neural networks can approximate very complex relations hidden in the data, their interpretation may be challenging. In general, it is desirable to obtain an expression that characterizes, or approximates, a given relation, because the physics behind it is much easier to understand and interpret in that form. For this purpose we made use of techniques designed to approximate functions with analytic expressions.

We made use of two different methods:

1. *Genetic programming*. The idea behind this method is to start with a series of operators (e.g., +, −, ×, /, exp) and create combinations with the operators and the variables in a so-called generation. Those expressions are then evaluated and the most accurate ones *survive* to the next generation, where *mutations* and *crossovers* can take place to explore different equations and find an optimal one. We made use of the PYSR (Cranmer 2020) and EUREQA¹³ packages for this.
2. *Template fitting*. Differently from genetic programming, here we start with a functional form that has some free parameters and perform a nonlinear least-squares fit to the data to obtain the best value of the parameters. This is done using the curve-fit function from SCIPY-OPTIMIZE.¹⁴ The idea is to use physically motivated expressions or to improve the results obtained from genetic programming.

When searching for analytic expressions we typically combine both methods. For instance, we use genetic programming to obtain the main trend and then we improve on that using template fitting.

3. Results

In this section, we present the results of our analysis with the neural network and symbolic-regression models.

3.1. Neural Networks

We start by training a neural network to predict the total subhalo mass from the other 11 properties using the subhalos of the CAMELS-IllustrisTNG simulations. In the top-left panel of Figure 1 we show the predicted total subhalo mass versus its true value from the subhalos of the CAMELS-IllustrisTNG test set (that the network has not seen before). As can be seen, the network is able to predict the value of the subhalo total mass very accurately: the rms error is 2.10×10^{-2} and 99.99% of the subhalos have a predicted total mass that is within 0.2 dex of its true value.

One may wonder if the tight relation found by the network arises from a strong correlation with a single variable. For instance, the network may have found a trivial relation between the mass and the radius of a subhalo. To check this, we make scatter plots between M_{tot} and R , M_{tot} and σ , and M_{tot} and V_{max} , as these three are the most important ones used by the network (see below), and fit a generic function that goes through the mean of the relation between the variables. Given that function, one can predict the total subhalo mass given the value of the

¹² <https://pytorch.org>

¹³ <https://www.creativemachineslab.com/eureqa.html>

¹⁴ <https://docs.scipy.org/doc/scipy/reference/optimize.html>




Figure 1. We train a neural network to predict the total mass of a subhalo from other 11 properties of the subhalo and the galaxy that it hosts. The training is carried out using subhalos from the CAMELS-IllustrisTNG simulations at redshift $z = 0$. The different panels show the predicted mass vs. the true value for subhalos from the CAMELS-IllustrisTNG (top left), CAMELS-SIMBA (top right), IllustrisTNG100 (bottom left), and IllustrisTNG300 (bottom right) simulations. The plots are two-dimensional histograms with the color bars indicating the number of galaxies in each bin of predicted and true mass. The rms error is quoted on the bottom right of each plot and the percentage of the predictions that lie within 0.2 dex of the actual value is quoted on the top left corners. The trained model exhibits surprising extrapolation properties, indicating that a universal relation may have been found.

considered variable. The rms that these functions return are 3.48×10^{-1} (for R), 3.79×10^{-1} (for σ), and 4.48×10^{-1} (for $V_{\max}$). These rms are much higher than the one obtained by the network, indicating that the relation cannot be attributed to a strong correlation with a single variable.

We then test the model on the subhalos from the CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300 simulations and show the results in the other panels of Figure 1. We find that even if these subhalos come from simulations with different subhalo mass ranges, resolutions, hydrodynamic solvers, and subgrid models than the ones used for training, the model can still accurately predict the total mass of those subhalos. The rms for the IllustrisTNG300 subhalos is similar to the one from the CAMELS-IllustrisTNG, while for CAMELS-SIMBA and IllustrisTNG100 the rms is a factor of $\simeq 3$ larger. We note that out of these three, it is expected that the lowest rms will be obtained by the subhalos from the

IllustrisTNG300 simulation, as those subhalos are similar to those from the IllustrisTNG-CAMELS simulations both in terms of resolution and simulation code. The fraction of subhalos whose predicted total mass is within 0.2 dex of the truth is very high in all cases: 99.77% (CAMELS-SIMBA), 97.90% (IllustrisTNG100), and 99.99% (IllustrisTNG300).

We note, however, a few interesting points. First, the scatter in the predicted total mass is significantly larger for CAMELS-SIMBA subhalos than for subhalos of the other simulations for almost all masses. Second, the network is not able to accurately predict the total mass of the most massive subhalos. Third, the network predicts the total subhalo mass with large scatter for very small subhalos, e.g., those with masses below $10^8 h^{-1} M_{\odot}$ from the IllustrisTNG100 simulation. All these cases can be seen as the network not extrapolating properly. In the case of CAMELS-SIMBA, their subhalo's galaxies are generated with a code that solves the hydrodynamic equations in a different

way and utilizes a distinct subgrid model. In the case of IllustrisTNG100 and IllustrisTNG300, their higher resolution and larger volume allow these simulations to contain subhalos with lower and higher total mass than those in the CAMELS-IllustrisTNG simulations used for training (see Table 1). Thus, it is important to acknowledge that while the percentages of the predictions that fall within 0.2 dex of the truth are calculated to be over 99% for both the CAMELS-SIMBA and IllustrisTNG300 simulations, these values do not appropriately weigh the most massive subhalos that are larger than $10^{14} h^{-1} M_{\odot}$, which exhibit significantly greater errors. Nevertheless, the model overall exhibits superb accuracy and appears to utilize a robust relation between different internal subhalo properties to predict the total mass.

We also examine how well the neural network is able to predict the total mass for central versus satellite subhalos, as these two types of systems exhibit different physical characteristics. We find that for all four simulations, the neural network predictions for the total mass of central subhalos demonstrate slightly higher accuracy than those for the satellite subhalos. The rms for the centrals (satellites) are as follows: for CAMELS-IllustrisTNG, 1.94×10^{-2} (2.26×10^{-2}); for CAMELS-SIMBA, 6.69×10^{-2} (7.11×10^{-2}); for IllustrisTNG100, 3.97×10^{-2} (7.91×10^{-2}); and for IllustrisTNG300 1.90×10^{-2} (2.25×10^{-2}). In general, the neural network effectively captures the trend for the total mass for both types of subhalos. We provide more details in Appendix A.

Our model not only works for subhalos from simulations with different volumes, resolutions, and subgrid models to the ones used for training, but also at redshifts other than the ones used for training. We illustrate this in Appendix B, where we display in Figure 8 the results of testing the network on redshifts other than $z = 0$ (the one used for training). Our model is able to extrapolate very well, and the rms and fraction of subhalos within 0.2 dex is similar to those at $z = 0$ and higher redshifts for subhalos from all simulation types. The largest differences we find are for subhalos of the CAMELS-SIMBA simulations.

One key aspect needed for the network to work at higher redshifts was to use the proper radius, R_p , instead of the comoving radius, $R = R_p/a$, where $a = 1/(1+z)$ is the scale factor. Without this change of variable, the network was unable to predict the correct total mass of the subhalos. This indicates that the relation learned by the network depends on the physical radius rather than the comoving radius. We provide further details on this issue in Appendix C.

Next, we examine which subhalo properties are the most important for the neural network's predictions by computing the saliency value of each input variable as described in Section 2.1. This is shown in Figure 2; we find that the top five most important properties, in descending order, are as follows: R , σ , $V_{\max}$, M_g , and SFR. We note that this exercise is done for subhalos with different masses, and the saliency values represent the average importance over all different subhalo masses; subhalos with different masses may exhibit different saliency values for the different variables.

To explore the effect of these variables on the accuracy of the neural network, we train four separate models to predict the total subhalo mass using different combinations of the five most important properties and compare the results with that of the original model that was trained on all 11 properties. We list the rms of the predictions from each of the models when tested




Figure 2. The saliency values of the input subhalo properties. Larger saliency values indicate that the network output is more affected by changes in the value of that variable and it can be seen as a proxy for feature importance. As shown, the five most important properties are R , σ , $V_{\max}$, M_g , and SFR.

on subhalos from the four different simulations in Table 2. We find that for subhalos of the CAMELS-IllustrisTNG, training on more variables helps in achieving higher accuracy. This is expected since adding more variables should only increase the information content, not decrease it. We observe the same behavior when the network is tested on subhalos from the IllustrisTNG300 simulation.

On the other hand, the model trained on the five subhalo properties achieves a higher accuracy than the model trained on all properties when tested using subhalos of the CAMELS-SIMBA and IllustrisTNG100 simulations. This may be due to the fact that the network, when trained on all properties, may be extracting information from properties that are unique to one particular simulation. In other words, the network may be learning artificial correlations that are only present in simulations with a given physics implementation or resolution. The model trained on five properties seems to learn more robust relationships that are able to better extrapolate when applied to subhalos from other simulations.

We also investigate if the relation found by the network is unique to subhalos of hydrodynamic simulations or whether it may be a more generic property that applies to subhalos of gravity-only N -body simulations. We carry out this task by testing the network trained on the R , σ , and $V_{\max}$ properties of the subhalos of the CAMELS-IllustrisTNG simulations on subhalos from N -body simulations of CAMELS. We show the results in Appendix D (Figure 11). We find that our model is also able to accurately predict the total mass of these subhalos that originate from simulations with different cosmological models. We provide further details of this test in Appendix D.

Finally, we study if our conclusions change if we train on subhalos from simulations other than those of the CAMELS-IllustrisTNG simulations. For this, we train a network using all 11 properties from subhalos of the CAMELS-SIMBA simulations, and test the model on CAMELS-IllustrisTNG, IllustrisTNG100, and IllustrisTNG300. We find that that model is able to capture a similar trend for all simulations but predictions are less accurate than those from the model trained on CAMELS-IllustrisTNG subhalos. Similarly, we also train a neural network using all 11 properties of subhalos from the IllustrisTNG300 simulation and test the model on subhalos from other simulations. In this case we also find that the predictions have an order of magnitude higher rms than the predictions from the neural network trained on CAMELS-

Table 2

Neural Networks Trained to Predict the Total Subhalo Mass from Different Combinations of Input Properties using the Subhalos from the CAMELS-IllustrisTNG Simulations

	All Properties	$\sigma, V_{\max}, R, M_g, \text{SFR}$	$\sigma, V_{\max}, R, M_g$	$\sigma, V_{\max}, R$	σ, R
CAMELS-IllustrisTNG	2.1×10^{-2}	2.3×10^{-2}	2.6×10^{-2}	3.0×10^{-2}	5.3×10^{-2}
CAMELS-SIMBA	6.9×10^{-2}	5.2×10^{-2}	5.6×10^{-2}	6.6×10^{-2}	1.0×10^{-1}
IllustrisTNG100	6.2×10^{-2}	5.3×10^{-2}	6.8×10^{-2}	9.0×10^{-2}	1.3×10^{-1}
IllustrisTNG300	2.1×10^{-2}	2.2×10^{-2}	2.5×10^{-2}	3.0×10^{-2}	5.6×10^{-2}

Note. Each column indicates which properties were used for training, and each row shows the rms errors of the predicted total mass when the model was tested on the respective simulation sets. As seen, using more variables yields more accurate results for subhalos of the CAMELS-IllustrisTNG simulations, but in some cases it makes the extrapolation worse for subhalos of other simulations.

IllustrisTNG. We provide further details on these tests in Appendix E.

One reason that can explain these results is that the subhalos from the CAMELS-IllustrisTNG simulations exhibit a very large variety (given that they come from simulations with different cosmological and astrophysical parameters) that helps the network to find more robust relations. When the network is trained on subhalos from a single simulation with fixed cosmology and astrophysics (e.g., IllustrisTNG300), the model may be learning artificial relations (such as numerical artifacts) that may not extrapolate very well. On the other hand, the subhalos from the IllustrisTNG100, IllustrisTNG300, and CAMELS-IllustrisTNG may share many similarities, since they are generated with the same code. Thus, the network trained on CAMELS-SIMBA subhalos may not extrapolate as well to these three simulations, in the same way that the model trained on CAMELS-IllustrisTNG does not extrapolate to subhalos from the CAMELS-SIMBA simulations as well as for the other IllustrisTNG-based simulations. Overall, although the model trained on CAMELS-IllustrisTNG seems to perform the best when testing on subhalos from other simulations, the models trained on CAMELS-SIMBA and IllustrisTNG300 are also very accurate and possess powerful extrapolation properties. We hence conclude that our results are not largely affected by the type of simulations used for training.

3.2. Analytic Expressions

The above results indicate that the function $M_{\text{tot}} = f(\vec{\theta})$ found by the neural network may be a universal one, as it holds accurately when tested on subhalos from simulations with different cosmologies, astrophysical models, subgrid models and hydrodynamic solvers, different volumes, different resolutions, and at different redshifts than those used when training the model. If that is the case, it would be very interesting to obtain an analytic expression that reproduces, or approximates, f in order to improve our knowledge on the physics behind such an relation.

In order to search for such an equation we first train a genetic programming model on a subset of inputs ($\vec{\theta}$) and outputs (M_{tot}) from subhalos from the CAMELS-IllustrisTNG simulations that cover the whole mass range. We did this using the PYSR package with the binary operators “ADD”, “SUB”, “MULT”, “DIV”, and the unary operator “NEG”. The model was trained for 1000 trials using a batch size of 2048. The output contains a list of equations found by the model, and associated with each equation were its complexity, mean square error (MSE), and score. The complexity is a metric that accounts for the number

of operators, constants, and input variables used; the score of an equation is higher if it is more simple and accurate. To select the best equation, we take into consideration all three values. Among the different equations found, the one that stands out for its simplicity and accuracy is

$$M_{\text{tot}} = A\sigma^\alpha R^\beta, \quad (5)$$

where $A = 10^{6.6}$, $\alpha = 1.9$, and $\beta = 0.9$. We note that A has units of $h^{-1} M_\odot$, while σ and R are assumed to have units of km s^{-1} and $h^{-1} \text{kpc}$, respectively. This equation is able to capture the main trend of the function f but it also exhibits significant biases and large scatter on the high- and low-mass ends, respectively (see Figure 3). The rms values of this equation are as follows: 7.60×10^{-2} (CAMELS-IllustrisTNG), 9.28×10^{-2} (CAMELS-SIMBA), 1.24×10^{-1} (IllustrisTNG100), and 9.32×10^{-2} (IllustrisTNG300). Evidently, this model does not exhibit an accuracy comparable to that of the neural network discussed in Section 3.1.

In order to improve the accuracy of the above expression, we perform template fitting on the CAMELS-IllustrisTNG simulations, considering a relation of the form

$$M_{\text{tot}} = A\sigma^\alpha R^\beta V_{\max}^\gamma, \quad (6)$$

where A , α , β , and γ are the free parameters we fit to the data. We find several interesting things. First, the above equation is more accurate than Equation (5), but its overall performance is still not satisfactory; the rms errors for this model tested on all simulations are listed in Table 3. This is mainly because we find that the free coefficients exhibit some dependence with the total mass of the subhalo. We fix this by using a slightly more complicated equation where the exponents have some implicit dependence on the parameters. Second, the dependence on $V_{\max}$ is crucial to accurately predict M_{tot} for small subhalos ($M_{\text{tot}} < 10^{10} h^{-1} M_\odot$). We note that we find this dependence to be negligible for subhalos above $\simeq 10^{12} h^{-1} M_\odot$. We believe that this is expected since massive galaxies tend to have their maximum circular velocity at small radii deep within the stellar body of the galaxy (e.g., in the bulge), and hence it is expected to be related more to the detailed dynamics at small radii than to the total mass of the subhalo.

Thus, we choose to improve Equation (6) by using a template of the form

$$M_{\text{tot}} = A\sigma^{(\alpha_0 + \alpha_1 \log \sigma)} R^{(\beta_0 + \beta_1 \log R)} V_{\max}^{(\gamma_0 + \gamma_1 \log V_{\max})}, \quad (7)$$




Figure 3. The accuracy of the equation $M_{\text{tot}} = A\sigma^\alpha R^\beta$, where $A = 10^{6.6}$, $\alpha = 1.9$, and $\beta = 0.9$, when used to predict the total mass of subhalos from the different simulations. Although this simple equation is able to capture the main trend of the relation, it can be seen that the predictions exhibit a large scatter in the low-mass end and a significant bias in the high-mass end for all simulations. The rms error is quoted in the bottom right and the percentage of predictions that lie within 0.2 dex of the actual value is quoted on the top left.

where $\log(x) = \log_{10}(x)$. Instead of fitting the above equation over the whole mass range, we fit it over three different mass ranges: (1) $M_{\text{tot}} < 10^{10} h^{-1} M_{\odot}$, (2) $10^{10} h^{-1} M_{\odot} < M_{\text{tot}} < 10^{12} h^{-1} M_{\odot}$, and (3) $M_{\text{tot}} > 10^{12} h^{-1} M_{\odot}$. The equations we derive are as follows:

$$M_{\text{tot}} = \begin{cases} 10^{5.47} R^{(0.96-0.06 \log(R/R_0))} \sigma^{(0.19-0.10 \log(\sigma/\sigma_0))} V_{\text{max}}^{1.94}, & \text{if } M_{\text{tot}} \leq 10^{10} h^{-1} M_{\odot} \\ 10^{5.32} R^{(0.87+0.06 \log(R/R_0))} \sigma^{0.72+0.40 \log(\sigma/\sigma_0)} V_{\text{max}}^{(1.58-0.54 \log(V_{\text{max}}/V_{\text{max}0}))}, & \text{if } 10^{10} h^{-1} M_{\odot} < M_{\text{tot}} < 10^{12} h^{-1} M_{\odot} \\ 10^{7.37} R^{0.96} \sigma^{(1.46+0.24 \log(\sigma/\sigma_0))}, & \text{if } M_{\text{tot}} \geq 10^{12} h^{-1} M_{\odot}. \end{cases} \quad (8)$$

In the above equations, for $M_{\text{tot}} < 10^{10} h^{-1} M_{\odot}$, $R_0 = 2.41 h^{-1} \text{ kpc}$ and $\sigma_0 = 25.6 \text{ km s}^{-1}$. For $10^{10} h^{-1} M_{\odot} <$

$M_{\text{tot}} < 10^{12} h^{-1} M_{\odot}$, $R_0 = 30.84 h^{-1} \text{ kpc}$, $\sigma_0 = 63.25 \text{ km s}^{-1}$, and $V_{\text{max}0} = 123.33 \text{ km s}^{-1}$. For $M_{\text{tot}} > 10^{12} h^{-1} M_{\odot}$, $\sigma_0 = 148.96 \text{ km s}^{-1}$. We also note that units of R , σ , and V_{max} are $h^{-1} \text{ kpc}$, km s^{-1} , and km s^{-1} , respectively.

In Figure 4 we compare the accuracy of these equations versus the one from two networks: one trained only using R , σ , and V_{max} (the variables present in the analytic formula), and the

Table 3

The rms Errors of the Various Models when Tested on the Four Simulations: CAMELS-IllustrisTNG, CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300

Model	CAMELS-IllustrisTNG	CAMELS-SIMBA	IllustrisTNG100	IllustrisTNG300
Equation (5)	7.60×10^{-2}	3.05×10^{-1}	1.24×10^{-1}	9.32×10^{-2}
Equation (6)	5.58×10^{-2}	2.02×10^{-1}	5.48×10^{-2}	4.94×10^{-2}
Equation (8)	3.96×10^{-2}	1.35×10^{-1}	3.08×10^{-2}	3.27×10^{-2}
Equation (G1)	3.32×10^{-2}	1.07×10^{-1}	3.08×10^{-2}	2.74×10^{-2}

Note. Each row indicates the models defined in Equations (5), (6), (8), and (G1).

other trained using all properties. In all cases, the training has been carried out using subhalos of the CAMELS-IllustrisTNG simulations at $z=0$. We find that these equations are able to accurately predict the total mass of a subhalo from its radius, maximum circular velocity, and velocity dispersion, independent of whether the subhalo is from one simulation or another.

There are several interesting features to comment on. First, in general, the accuracy of the equation is worse than the network trained on the same three variables. This indicates that our formula is not capturing the full functional form learned by the network. Second, the model trained using all variables is, in general, more accurate than both the model trained with three variables and the analytic equation. This means that our formula does not account for the additional information on M_{tot} carried out by variables other than R , V_{max} , and σ . Third, our equation is able to predict the total mass of subhalos with very low or very high total mass (beyond the range the model was trained on) more accurately than the neural networks. This may be happening because in that regime the network is in extrapolation mode and may be subject to a number of factors limiting its accuracy, such as learning uninformative priors (Villaescusa-Navarro et al. 2020).

We then test the accuracy of Equation (8) on subhalos from the different simulations at higher redshifts, finding the accuracy of the equation at those redshifts to be very similar as at $z=0$. We note that when using Equation (8) at redshifts other than 0, we need to use the proper radius instead of the comoving radius (see Appendix C).

Finally, we also test Equation (8) on dark matter subhalos from CAMELS-Nbody simulations at redshifts $z=0$ and $z=0.5$. We find that the accuracy of the analytic model is similar to the accuracy of the neural network that was trained using the same three variables (see Figure 11).

Overall, our relatively simple equations are able to capture the main trend of the relation between the total mass of a subhalo and its other internal properties. We note that all above equations have been derived from the PYSR software. In Appendix G we present a set of equations found using the symbolic-regression package EUREQA, which performs with slightly higher accuracy than the above equations but has a more complex form. The rms errors for these equations are also included in the last row of Table 3 for comparison. In Figure 5 we show the rms values for the different simulations when using the neural networks or the different equations.

4. Summary and Discussion

The formation and evolution of galaxies is driven by numerous complex physical mechanisms whose details we do not yet fully understand. One possibility to improve our understanding of these processes is to find relationships

between subhalo and galaxy properties that may reflect the underlying physics governing their formation and evolution. In this work we follow a generic methodology to search for one such relation connecting the total mass of a subhalo (M_{tot}) with other internal properties ($\vec{\theta}$): $M_{\text{tot}} = f(\vec{\theta})$. We have used machine-learning methods to approximate the function f using state-of-the-art hydrodynamic simulations and learned about its properties.

The most important findings of this work are as follows:

1. Neural networks are able to find a relation between the total mass of a subhalo and its other properties, independently of the subhalo type (i.e., a central or a satellite). The model is able to accurately predict the total mass of subhalos from simulations with different cosmologies, different astrophysical models, different subgrid models and hydrodynamic solvers, different volumes, different resolutions, and different redshifts than the ones it has been trained on. The surprising accuracy achieved by the network in these different extrapolations indicates that the network may have found a universal relationship within the subhalo properties.
2. Our networks only work at redshifts $z > 0$ if the proper radius is used instead of the comoving radius. This indicates that the learned relation depends on physical variables and not comoving ones.
3. We find that the models trained on fewer subhalo variables (R , V_{max} , σ , M_g , and SFR) are able to extrapolate slightly better than models trained using all subhalo properties. This could be due to the fact that the network may be leveraging information that is only present in some simulations, and therefore it extrapolates worse when those conditions are no longer met.
4. When the network is trained using R , V_{max} , and σ from subhalos, it is also able to predict the total mass of subhalos from N -body simulations, although with higher scatter for subhalos with low mass.
5. Our results are robust to the type of simulation used for training the networks, including very different subgrid galaxy formation models.
6. A simple relation of the form $M_{\text{tot}} = AR^\alpha \sigma^\beta$ is able to roughly capture the main relation between the total mass of a subhalo and its other properties. However, it exhibits significant biases in the low- and high-mass ends. We find that including V_{max} is crucial to properly capturing the underlying relation in the low-mass end.
7. Equation (8) is a much more accurate expression that works for subhalos from simulations with different cosmologies, astrophysics, subgrid physics, volumes, resolutions, and at different redshifts. Furthermore, this




Figure 4. We fit analytic models to the CAMELS-IllustrisTNG simulations for three different mass ranges and test them on each simulation. This figure compares them to the results of the neural networks. From top to bottom, the rows illustrate the predictions for CAMELS-IllustrisTNG, CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300 simulations. The first column shows the predictions from the analytic model shown in Equation (8); the next column shows the predictions from the neural network trained using only three properties (σ , $V_{\max}$, and R); and the last column shows the predictions from the neural network trained using all 11 subhalo properties. Moreover, the dotted lines at $10^{10} h^{-1} M_{\odot}$ and $10^{12} h^{-1} M_{\odot}$ divide each plot into three sections to indicate the mass ranges that the symbolic models were tested and trained on. The values quoted at the top of each plot for each mass range are the rms errors and the percentage of predictions within 0.2 dex of the truth, respectively. These values are also quoted at the bottom left of each plot for all the mass ranges of that simulation. As shown, in the high- and low-mass ends, the neural networks perform worse than the analytic models when tested on simulations other than CAMELS-IllustrisTNG.




Figure 5. This bar graph compares the rms errors of the various models when tested on the four simulations: CAMELS-IllustrisTNG, CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300. The legend indicates the model associated with each color: Equations (5), (6), (8), and (G1). It can be seen that while the neural network model exhibits very high accuracy in the CAMELS-IllustrisTNG and IllustrisTNG300 simulations, it is unable to extrapolate well to CAMELS-SIMBA and IllustrisTNG100. Moreover, Equation (5), which is similar in form to the virial theorem, performs poorly on all simulations. In general, Equations (8) and (G1) are the best analytic models in terms of accuracy and extrapolation ability, and exhibit similar rms.

equation is able to extrapolate in the low- and high-mass ends more precisely than the neural networks.

8. The neural networks achieve a higher accuracy than Equation (8), indicating that (1) our expression may only be an approximation to the relation found by the network, and (2) that relation should also contain a dependence on other variables, such as gas mass and star formation rate, that our equation lacks.

The rather surprising extrapolation properties of both our neural networks and the analytic equation may indicate that the relation between the total mass of a subhalo and its other properties is set by a relatively simple, universal law.

To our knowledge, the equations derived in this work are new, and, while complex, they are able to capture the main trend of the relation between the mass and other subhalo properties. These equations may help us in understanding the underlying physics behind subhalo/galaxy formation and evolution that, ultimately, is responsible for shaping the distribution of dark matter, gas, and stars in the universe.

We believe, however, that simpler and more accurate equations that capture the apparent universal relationship between the total mass of a subhalo and its other internal properties can be found. As symbolic-regression algorithms improve, the found equations may be more accurate and simple, as well. There are also other adjustments to the pipeline that can be implemented to improve the equations, such as changing the normalization of the data and decreasing the noise from the data by smoothing the data using Gaussian processes. One could also try to use symbolic regression to find a single equation that maps a property, such as radius, to the piecewise equations. This can be performed for all the variables used in the equations to obtain one relation that fits all mass ranges. We leave this for future work.

We also think that although these results are generic, their details may depend on the particular algorithm used to identify subhalos and halos in the simulations. For instance, a different criteria in the definition of radius or velocity dispersion may lead to subhalos for which our network and analytic expression do not work. Moreover, we conjecture that the functional form of the found relations may remain robust for different subhalo-

finding algorithms but the coefficients in the power law may differ. This will need further investigation to ascertain. However, in this case we believe retraining the network and refitting the best values of Equation (8) may fix the issue.

Lastly, we attempt to provide a physical interpretation of our results. From the variables used in Equation 5, as well as its power-law form, we may guess that the network is learning a relation that contains similarities to the virial theorem, which states that in a self-gravitating body the total kinetic (T) energy relates to the gravitational potential energy (U) via $2T + U = 0$, or $M = V^2 R / G$. The fact that the relation learned by the network depends on the proper (physical) size of subhalos, instead of their comoving scale, reinforces this point.

However, it cannot be that same relation because the variables used do not correspond with the ones appearing in the virial theorem, e.g., our radius, R , is defined as the radius containing half of the subhalo total mass, while in the virial theorem R is the radius of the virial system. Besides, our network and analytic expression make use of $V_{\max}$, which does not appear in the virial theorem. On top of this, we know that variables such as the gas mass and the star formation rate are being used by the network to improve its accuracy. However, the network may be using all these variables as a proxy for the ones appearing in the virial theorem.

The universality of our results could also be explained if the underlying relation arises from the virial theorem, which is derived from very generic arguments related to gravity, and therefore is not affected by cosmology, astrophysics, simulation volume, simulation resolution, subgrid physics, and redshift.

Overall, this paper presents a generic formalism to search for relationships in high-dimensional spaces and approximate them with analytic expressions. In future work we will use this formalism to search for relations in galaxy properties from actual observations and derive analytic equations from them.

We thank Rachel Somerville for useful conversations. This work has made use of the Tiger cluster of Princeton University and the Iron and Popeye clusters at the Flatiron Institute, which is supported by the Simons Foundation. F.V.N. acknowledges funding from the WFIRST program through NNG26PJ30C, NNN12AA01C, and from NSF through grant No. AST-2108078. D.A.A. was supported in part by NSF grant Nos. AST-2009687 and AST-2108944. Details on the CAMELS simulations can be found in <https://www.camel-simulations.org>.

Appendix A Central versus Satellite Galaxies

In this section, we compare the neural network's predictions when subhalos are split into central and satellites for all different simulation types. The results for the central subhalos are depicted in Figure 6 while the results for the satellites are shown in Figure 7. As can be seen, while the latter exhibits slightly higher error, the model is still able to accurately predict the total mass for both types of subhalos. This is a surprising result given that the neural network was not given information that directly differentiates the satellites from the centrals, which may indicate that such information is implicitly encoded in relations between the input subhalo properties. Moreover, there are several interesting points to note. First, it is known that satellites with eccentric orbits are being tidally stripped in a time-dependent manner, as well as tidally shocked, as they




Figure 6. We test the neural network trained to predict the total subhalo mass from 11 other subhalo properties on the central subhalos from the CAMELS-IllustrisTNG, CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300 simulations. In this figure, we plot the model predictions against the truth for each simulation on one set of axes to conserve space. For the CAMELS-IllustrisTNG subhalos, we plot the predicted mass against the truth. For the CAMELS-SIMBA subhalos, we plot the predicted mass against the truth plus 3 dex. For the IllustrisTNG100 subhalos, we plot the predicted mass against the truth plus 7 dex. Finally, for the IllustrisTNG300 subhalos, we plot the predicted mass against the truth plus 10 dex. As can be seen, the model is able to perform with very high accuracy for the central subhalos, as more than 99% of its predictions lie within 0.2 dex of the truth for all simulations. However, to clarify, although this percentage is stated as 100.00% for both the CAMELS-TNG and IllustrisTNG300 subhalos, the actual values are 99.999% but have been truncated to two decimal places. This should also be taken into consideration with the fact that the metric does not place more weight on the most massive subhalos, which contain significantly larger prediction errors.




Figure 7. Same as Figure 6 but for the subhalo satellites from the four different simulations.

reach pericenter. Hence, it is not obvious why the relation found by the neural network, one we believe to be rooted in the virial theorem, holds in this case where the system may not be virialized. A possible explanation is that subhalos that are undergoing mass stripping may still be in a semi-equilibrium state, which allows the conditions of the virial theorem to apply. However, further investigation on such satellites is needed to determine their effect on the neural network's predictions. Second, while the neural network does not explicitly account for the tidal radius of the satellites or their location relative to their hosts, which are essential pieces of information for the orbiting bodies, it is still able to accurately predict the total subhalo mass. Here, we believe that the radius used to train the neural network implicitly encodes the information about the tidal radius of the satellites. Given that the radius we use is defined to contain half the total subhalo

mass, the effect of tidal stripping would reduce both mass and radius. Third, it is possible that the larger scatter seen for satellites is a result of both such physical effects and numerical effects related to the specific substructure-finding algorithm, which may sometimes have difficulty identifying the full extent of subhalos when they pass close to the host's center. A more detailed analysis of the relative dynamics between the satellites and their hosts in the context of the found relations is needed, and a more elaborate treatment of this can possibly reduce the scatter in the fittings. We leave this for future investigation.

Appendix B Higher Redshifts

Here we quantify how well our neural network, trained on CAMELS-IllustrisTNG subhalos at $z = 0$, can predict the total mass of subhalos from the different simulations at higher




Figure 8. We test the neural network that was trained on CAMELS-IllustrisTNG subhalos using all 11 properties at $z = 0$ on subhalos from different simulations at different redshifts. The percentage of predictions that fall within 0.2 dex of the true value is quoted at the top left of each plot, along with the rms error at the bottom right. As demonstrated, the model is able to extrapolate in redshift very accurately in all cases. This may indicate that the network has indeed found a universal relationship among subhalo properties.

redshifts. We show in Figure 8 the results of testing the network on subhalos at redshifts $z = 0, z = 0.5, z = 1, z = 1.5$, and $z = 2$. We find that our model can accurately predict the total mass of the subhalos at all considered redshifts, independently of the simulations they come from. We emphasize that CAMELS-SIMBA simulations do not use the same hydrodynamic code as the IllustrisTNG-like simulations, and utilize a very distinct subgrid physics model. Therefore, the high accuracy of the model's predictions at higher redshifts for the SIMBA subhalos is not a trivial result. This reinforces our belief that the neural network may have found a universal relation since it has never seen subhalos at redshifts other than $z = 0$.

It is important to emphasize that in this exercise we have fed the network with the proper radius of a subhalo instead of its

comoving radius. In Appendix C we show what happens when the comoving radius is used at higher redshifts. This indicates that the relation learned by the network depends on the physical radius and not on the comoving radius.

Appendix C Comoving versus Proper Radius

While exploring the neural network's performance at higher redshifts, we found that using different definitions of the subhalo radius led to drastically different results. Namely, using the comoving radius, defined as the proper (physical) radius divided by the scale factor $a = 1/(1+z)$, for subhalos at redshifts higher than $z = 0$ caused the neural network predictions to have a vertical shift, resulting in very low accuracy. This can be seen in the left panel of Figure 9. We




Figure 9. These plots show the large difference in accuracy of the neural network’s predictions when using the comoving radius (left) vs. the proper radius (right) of the subhalos at redshift $z = 2$. Here, the neural network was trained on subhalos of the CAMELS-IllustrisTNG simulations at $z = 0$ and tested on subhalos from the IllustrisTNG300 simulations at $z = 2$. The rms error is quoted in the bottom right and the percentage of predictions that lie within 0.2 dex of the actual value is quoted on the top left. This shows that the relation learned by the network depends on proper lengths and not comoving ones.




Figure 10. We trained a neural network to learn the total subhalo mass from the 11 subhalo properties of the CAMELS-IllustrisTNG simulations at redshift $z = 2$ using the comoving radius as a variable, and tested the model on the subhalos from CAMELS-IllustrisTNG at redshift $z = 0$. At this redshift, the comoving and proper radius definitions are equal. As shown, the model performs poorly when predicting the total mass for the subhalos at a different redshift from the subhalos used during training. This illustrates the importance of using the proper variables when training the network if extrapolation properties are desired. The rms error is quoted in the bottom right and the percentage of predictions that lie within 0.2 dex of the actual value is quoted on the top left.

note that since the network was trained at $z = 0$, where the comoving radius equals the proper radius, it is not obvious which radius definition should be used. Thus, we tested the network using the proper radius of the subhalos and show the results in the right panel of Figure 9. As can be seen, in this case the network accuracy reaches similar values as the model at $z = 0$. We thus conclude that the relation learned by the network, and the one that the network employs to extrapolate to other redshifts and simulations, uses the proper radius of the subhalos and not their comoving one.

This is actually a very important point that shows the importance of training neural networks using the proper variables if extrapolation properties are desired. To illustrate this point, we train a neural network using subhalos of the CAMELS-IllustrisTNG simulations at $z = 2$ using the comoving radius as one of the input variables. The results of testing the network at $z = 2$ are shown on the left panel of Figure 10. It can be seen that the network achieves a very high accuracy when predicting the total mass of subhalos at $z = 2$. The right panel of Figure 10 shows instead what happens when the network is tested on subhalos at $z = 0$. We note that at this

redshift, the proper and comoving definitions are equal. Evidently, the network is not able to predict the correct total mass of the subhalos. This case shows how training the neural network on variables that are not the ones behind a relation may significantly affect the extrapolation properties of the network itself. It is, however, important to emphasize that the proper variables behind a given case may not be known a priori.

Appendix D N-body Simulations

Here we study if the model that is trained on subhalos from hydrodynamic simulations can predict the total mass of subhalos from gravity-only N -body simulations. Since the subhalos of the N -body simulations do not have some properties that are only present in the subhalos of the hydrodynamic simulations, such as SFR, gas mass, and others, we use the model that was trained on the three variables which the subhalos from both simulations share: R , $V_{\max}$, and σ . The left panel of Figure 11 shows the results when the model is




Figure 11. We trained a neural network model to predict the subhalo total mass from R , σ , and $V_{\max}$ using the subhalos from the CAMELS-IllustrisTNG suite (at $z = 0$). The leftmost plot shows the model predictions for a test set of the CAMELS-IllustrisTNG simulations (at $z = 0$). The following two plots depict the model predictions when tested on the N -body subhalos from the CAMELS-IllustrisTNG Dark Matter simulations at redshifts $z = 0$ and 0.5 . The rms error is quoted in the bottom right and the percentage of predictions that lie within 0.2 dex of the actual value is quoted on the top left.

tested on subhalos from the CAMELS-IllustrisTNG simulations at $z = 0$. As can be seen, the model is able to accurately predict the total mass of these subhalos over the whole mass spectrum. The middle and right panels of that figure show the results when that model is tested on subhalos from N -body simulations with different cosmologies at redshifts $z = 0$ and $z = 0.5$. Here, we also find that the model is able to predict the total mass of the N -body subhalos with high accuracy, although the rms is slightly higher than that for subhalos of hydrodynamic simulations and the fraction of N -body subhalos with its total mass within 0.2 dex of their true value is slightly lower than that for the hydrodynamic simulations.

These results show that the relation found by the network is not unique to subhalos from hydrodynamic simulations, but points toward something more fundamental that may just depend on gravity itself, such as a relation similar to the virial theorem.

Appendix E Training on Other Simulations

In order to quantify the dependence of our neural network results on their training set, we train two networks using subhalos from the CAMELS-SIMBA and IllustrisTNG300 simulations using all 11 properties. We then test each model on the subhalos

of the other simulations. We show the results in Figure 12 for the network trained on subhalos from the CAMELS-SIMBA simulations. It can be seen that even though the model is using a robust relation to predict the total mass, it is not as accurate as the model that was trained on subhalos from the CAMELS-IllustrisTNG simulations. The rms values for the neural network trained on CAMELS-SIMBA are: 7.67×10^{-3} (CAMELS-IllustrisTNG), 9.27×10^{-4} (CAMELS-SIMBA), 1.41×10^{-2} (IllustrisTNG100), and 4.70×10^{-3} (IllustrisTNG300).

In Figure 13 we show instead the results of the network trained on subhalos of the IllustrisTNG300 simulation. The accuracies of these predictions are similar to those of the model trained on CAMELS-SIMBA. The rms values for the neural network trained on IllustrisTNG300 are: 3.56×10^{-2} (CAMELS-IllustrisTNG), 9.84×10^{-2} (CAMELS-SIMBA), 1.14×10^{-1} (IllustrisTNG100), and 2.06×10^{-2} (IllustrisTNG300).

These results show that although the data used in the training set can have a relatively minor effect on the accuracy of the network, for which we provide further interpretation in Section 3.1, the main conclusions are unchanged, i.e., the network can accurately predict the total mass of subhalos from simulations with different cosmologies, astrophysics, subgrid physics, resolutions, and volumes than the ones used to train the network.




Figure 12. We train a neural network to learn the total subhalo mass from 11 other subhalo properties on the CAMELS-SIMBA simulations at redshift $z = 0$. We then test the model on a test set of subhalos from the CAMELS-SIMBA simulations, as well as on the subhalos of the CAMELS-IllustrisTNG, IllustrisTNG100, and IllustrisTNG300 simulations. In this figure, we plot the model predictions against the truth for each simulation on one set of axes to conserve space. For the CAMELS-IllustrisTNG subhalos, we plot the predicted mass against the truth. For the CAMELS-SIMBA subhalos, we plot the predicted mass against the truth plus 2 dex. For the IllustrisTNG100 subhalos, we plot the predicted mass against the truth plus 4 dex. Finally, for the IllustrisTNG300 subhalos, we plot the predicted mass against the truth plus 6 dex. It is apparent that the model is able to predict with relatively high accuracy as more than 92% of the predictions lie within 0.2 dex of the truth for all simulations. However, compared to the neural network that was trained on the CAMELS-IllustrisTNG subhalos (see Figure 1), this model is performing worse for all simulations except CAMELS-SIMBA. The rms error and the percentage of the predictions within 0.2 dex of the true value are quoted next to each scatter plot in the figure.




Figure 13. Same as Figure 12 but training on subhalos from the IllustrisTNG300 simulation.

Appendix F Input Properties

Here we investigate the effects of reducing the number of input subhalo properties on the neural network’s ability to predict the total subhalo mass. For that, we train a neural network to learn the total mass using only five subhalo properties, R , σ , $V_{\max}$, SFR, and M_g , on subhalos from the CAMELS-IllustrisTNG simulations. As illustrated in Figure 14, the neural network is still able to predict the total mass with high accuracy (see also Table 2). Notably, the model predictions for subhalos of the CAMELS-SIMBA and IllustrisTNG100 simulations exhibit lower rms errors compared to the predictions from the model trained on all 11 properties.

A potential explanation for this behavior is that when the training is carried out using all variables, the neural network may be exploiting relations that are unique to a particular simulation. These relations may be numerical artifacts or peculiarities of the particular simulations considered. In that case, training with more variables will improve the network precision when the network is tested on simulations of the same kind as the ones used for training, but may actually perform worse when tested on other simulations. By training on a smaller number of variables, the network may learn more robust correlations that can extrapolate better.

We note, however, that the differences between the results obtained by using all variables and five variables is not very large, indicating that the original network is already marginalizing over uninformative variables.




Figure 14. We train a neural network using the CAMELS-IllustrisTNG simulations to learn the total subhalo mass from five selected properties: R , σ , $V_{\max}$, SFR , M_g at $z = 0$. In this figure, we plot the model predictions against the truth for each simulation on one set of axes to conserve space. For the CAMELS-IllustrisTNG subhalos, we plot the predicted against the truth. For the CAMELS-SIMBA subhalos, we plot the predicted against the truth plus 3 dex. For the IllustrisTNG100 subhalos, we plot the predicted against the truth plus 7 dex. Finally, for the IllustrisTNG300 subhalos, we plot the predicted against the truth plus 10 dex. The rms error and the percentage of the predictions within 0.2 dex of the true value are quoted above each scatter plot in the figure. We can see that, even from only five properties, the neural network is able to predict the total mass with very low error. Even though it is performing worse than the original model for the CAMELS-IllustrisTNG simulation, it is doing better for the IllustrisTNG100 and CAMELS-SIMBA subhalos.

Appendix G Eureqa Equations

We train symbolic-regression models using the EUREQA package on the CAMELS-IllustrisTNG simulations to search for analytic expressions that can predict the total subhalo mass from the 11 other subhalo properties. First, we train a model on a subset of subhalos that span the entire mass range from the CAMELS-IllustrisTNG simulations. While the equation that the model discovered is able to capture the main trend of the total subhalo mass, it performs with lower accuracy than the analytic equations discussed in Section 3.2. To improve the results and aid the model’s search, we train symbolic-regression models on subhalos spanning the three different mass ranges used in Section 3.2: (1) $M_{\text{tot}} < 10^{10} h^{-1} M_{\odot}$, (2) $10^{10} h^{-1} M_{\odot} < M_{\text{tot}} < 10^{12} h^{-1} M_{\odot}$, and (3) $M_{\text{tot}} > 10^{12} h^{-1} M_{\odot}$. The equations that were found are as follows:

$$M_{\text{tot}} = \begin{cases} 10^{5.60} R^{2.72+2.00 \log \tilde{\sigma} - 0.26 \log \tilde{R}} \tilde{V}_{\max}^{1.65-0.68 \log \tilde{\sigma} \log \tilde{R}}, & \text{if } M_{\text{tot}} \leq 10^{10} h^{-1} M_{\odot} \\ 10^{5.77} R^{15.48+3.82 \log \tilde{V}_{\max} \log \tilde{\sigma}} \tilde{V}_{\max}^{1.89 \log \tilde{V}_{\max} \log \tilde{R} - 3.93} \tilde{\sigma}^{-5.17-1.93 \log \tilde{\sigma} \log \tilde{R}}, & \text{if } 10^{10} h^{-1} M_{\odot} < M_{\text{tot}} < 10^{12} h^{-1} M_{\odot} \\ 10^{9.74} R^{1.94+0.41 \log^2 \tilde{\sigma}} \tilde{\sigma}^{1.96+\log \tilde{R}-0.02 \log^2 \tilde{\sigma}}, & \text{if } M_{\text{tot}} \geq 10^{12} h^{-1} M_{\odot}. \end{cases} \quad (\text{G1})$$

In the above equations, we define $\tilde{R} = R/R_0$, $\tilde{\sigma} = \sigma/\sigma_0$, and $\tilde{V}_{\max} = V_{\max}/V_{\max 0}$. For $M_{\text{tot}} < 10^{10} h^{-1} M_{\odot}$, $R_0 = 2.41 h^{-1} \text{ kpc}$ and $\sigma_0 = 25.6 \text{ km s}^{-1}$. For $10^{10} h^{-1} M_{\odot} < M_{\text{tot}} < 10^{12} h^{-1} M_{\odot}$, $R_0 = 30.84 h^{-1} \text{ kpc}$, $\sigma_0 = 63.25 \text{ km s}^{-1}$, and $V_{\max 0} = 123.33 \text{ km s}^{-1}$. For $M_{\text{tot}} > 10^{12} h^{-1} M_{\odot}$,

$\sigma_0 = 148.96 \text{ km s}^{-1}$ and $R_0 = 113.24 h^{-1} \text{ kpc}$. We note that the units of R , σ , and $V_{\max}$ are $h^{-1} \text{ kpc}$, km s^{-1} , and km s^{-1} , respectively.

These equations, while more complex, are able to predict the total mass with higher accuracy than the analytic expressions we found using linear regression fitting (Equation (8)), with a rms value of 3.32×10^{-2} for the CAMELS-IllustrisTNG subhalos. When we test this on the CAMELS-SIMBA, IllustrisTNG100, and IllustrisTNG300 simulations we find that the equations are able to effectively extrapolate to the different simulations and to mass ranges different than those used to train the model. The results for this are presented in the panels of the first column of Figure 15. As shown, the rms values for these predictions are lower than those of Equation (8) for all four simulations. Similar to Equation (8), we see that the symbolic-regression models are able to accurately predict the total mass in the high- and low-mass regimes without

generating biases as the neural networks do. However, Equation (G1) still does not attain the overall low rms of the neural networks trained using the same three variables. From this, we conclude that the neural networks are employing a relation that is more complex than the models found by symbolic regression.




Figure 15. Same as Figure 3 but using Equation (G1) derived from EUREQA.

ORCID iDs

Helen Shao  <https://orcid.org/0000-0002-0152-6747>
 Francisco Villaescusa-Navarro  <https://orcid.org/0000-0002-4816-0455>
 Shy Genel  <https://orcid.org/0000-0002-3185-1540>
 David N. Spergel  <https://orcid.org/0000-0002-5151-0006>
 Daniel Anglés-Alcázar  <https://orcid.org/0000-0001-5769-4945>
 Lars Hernquist  <https://orcid.org/0000-0001-6950-1629>
 Romeel Davé  <https://orcid.org/0000-0003-2842-9434>
 Desika Narayanan  <https://orcid.org/0000-0002-7064-4309>
 Mark Vogelsberger  <https://orcid.org/0000-0001-8593-7692>

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, in in Proc. of the 25rd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining
- Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2017a, *MNRAS*, **464**, 2840
- Anglés-Alcázar, D., Faucher-Giguère, C.-A., Kereš, D., et al. 2017b, *MNRAS*, **470**, 4698
- Bender, R., Burstein, D., & Faber, S. M. 1992, *ApJ*, **399**, 462
- Burstein, D., Bender, R., Faber, S., & Nolthenius, R. 1997, *AJ*, **114**, 1365
- Cranmer, M. 2020, PySR: Fast & Parallelized Symbolic Regression in Python/Julia, Zenodo, doi:10.5281/zenodo.4041459
- Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, **486**, 2827
- Dressler, A., Lynden-Bell, D., Burstein, D., et al. 1987, *ApJ*, **313**, 42
- Faber, S. M., & Jackson, R. E. 1976, *ApJ*, **204**, 668
- Hopkins, P. F. 2015, *MNRAS*, **450**, 53
- Loshchilov, I., & Hutter, F. 2017, arXiv:1711.05101
- Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, **480**, 5113
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2015, *MNRAS*, **454**, 2691
- Naiman, J. P., Pillepich, A., Springel Enrico Ramirez-Ruiz, V., et al. 2018, *MNRAS*, **477**, 1206
- Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **475**, 624
- Nelson, D., Springel, V., Pillepich, A., et al. 2019, *ComAC*, **6**, 2
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018a, *MNRAS*, **475**, 648
- Pillepich, A., Springel, V., Nelson, D., et al. 2018b, *MNRAS*, **473**, 4077
- Somerville, R. S., & Davé, R. 2015, *ARA&A*, **53**, 51
- 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
- Tully, R. B., de Marseille, O., & Fisher, J. R. 1975, *BAAS*, **Vol. 7**, 426
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021, *ApJ*, **915**, 71
- Villaescusa-Navarro, F., Wandelt, B. D., Anglés-Alcázar, D., et al. 2020, arXiv:2011.05992
- Wadekar, D., Villaescusa-Navarro, F., Ho, S., & Perreault-Levasseur, L. 2021, *ApJ*, **916**, 42
- Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291
- Weinberger, R., Springel, V., & Pakmor, R. 2020, *ApJS*, **248**, 32