

Physics-informed neural networks in the recreation of hydrodynamic simulations from dark matter

Zhenyu Dai ^{1,2}★ Ben Moews ^{3,4,5,6} Ricardo Vilalta ^{1,2} and Romeel Davé ^{7,8,9}

¹Department of Computer Science, University of Houston, 3551 Cullen Blvd, Houston, TX 77204-3010, USA

²Department of Physics, University of Houston, 3507 Cullen Blvd, Houston, TX 77204-5005, USA

³Business School, University of Edinburgh, 29 Buccleuch Pl, Edinburgh EH8 9JS, UK

⁴Center for Statistics, University of Edinburgh, Peter Guthrie Tait Road, Edinburgh EH9 3FD, UK

⁵McWilliams Center for Cosmology, Carnegie Mellon University, Hamerschlag Dr, Pittsburgh, PA 15213, USA

⁶Pittsburgh Supercomputing Center, 300 South Craig St, Pittsburgh, PA 15213, USA

⁷Institute for Astronomy, University of Edinburgh, Royal Observatory, Edinburgh EH9 3HJ, UK

⁸Department of Physics and Astronomy, University of the Western Cape, Bellville 7535, South Africa

⁹South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa

Accepted 2023 October 31. Received 2023 October 31; in original form 2023 March 24

ABSTRACT

Physics-informed neural networks have emerged as a coherent framework for building predictive models that combine statistical patterns with domain knowledge. The underlying notion is to enrich the optimization loss function with known relationships to constrain the space of possible solutions. Hydrodynamic simulations are a core constituent of modern cosmology, while the required computations are both expensive and time-consuming. At the same time, the comparatively fast simulation of dark matter requires fewer resources, which has led to the emergence of machine learning algorithms for baryon inpainting as an active area of research; here, recreating the scatter found in hydrodynamic simulations is an ongoing challenge. This paper presents the first application of physics-informed neural networks to baryon inpainting by combining advances in neural network architectures with physical constraints, injecting theory on baryon conversion efficiency into the model loss function. We also introduce a punitive prediction comparison based on the Kullback–Leibler divergence, which enforces scatter reproduction. By simultaneously extracting the complete set of baryonic properties for the SIMBA suite of cosmological simulations, our results demonstrate improved accuracy of baryonic predictions based on dark matter halo properties and successful recovery of the fundamental metallicity relation, and retrieve scatter that traces the target simulation’s distribution.

Key words: methods: analytical – methods: statistical – galaxies: evolution – galaxies: haloes.

1 INTRODUCTION

The Lambda cold dark matter model, coined the standard model of cosmology due to its widespread adoption and explanatory power, plays a crucial role in modern cosmology and astrophysics. Galaxy formation and evolution occur within virialized structures resulting from density perturbations in the early Universe, subjected to gravitational collapse (Frenk & White 2012). This leads to large-scale structure in the form of a cosmic web evolving from a somewhat smooth starting point, with dark matter haloes as gravitationally bound overdensities of the postulated main contributor to the matter content of the Universe.

As galaxies form through the condensation of cooling gas within these haloes, their resulting baryonic properties share a natural relationship with the dark matter accumulations they live in (see, for example, Rees & Ostriker 1977; Blumenthal et al. 1984). As dark matter distributions are often sufficient to extract cosmological parameters of interest, N -body simulations restricted to simulating

dark matter particles or grids through gravitational interactions are common in cosmology (Springel et al. 2005; Boylan-Kolchin et al. 2009; Klypin, Trujillo-Gomez & Primack 2011; Riebe et al. 2013; Potter, Stadel & Teyssier 2017). Options in this regard include direct numerical integration and the inclusion of a scale factor to model the expansion of the Universe via general relativistic effects. These simulations are computationally cheap and fast to run compared to more complex alternatives (Efstathiou et al. 1985).

In contrast to the success of large-scale structure, modelling galaxies through N -body simulations has historically been more difficult, as baryonic physics plays a vital role through non-linear and dissipative processes at this level of granularity (Somerville & Davé 2015). At the same time, the relevance of the baryonic properties of galaxies is evident by the reason dark matter bears its name; direct observations in the electromagnetic spectrum rely on baryonic processes emitting such radiation. Methods to include the luminous baryonic matter in our simulations are thus required to enable comparisons to these observations in the first place.

Multiple avenues to circumvent these limitations have been developed. These include abundance matching, which connects the halo mass to a range of baryonic properties through the stellar

* E-mail: zdai3@uh.edu

mass. However, this traditionally requires the assumption of the preservation of rank ordering in mass (see, for example, descriptions in Wetzel & White 2010; Campbell et al. 2018), although this has been superseded by an earlier shift towards the inclusion of intrinsic scatter as well as a preference for ordering by velocity instead of mass (Reddick et al. 2013). Recent models, however, do not solely rely on mass as the halo proxy but instead use a combination of factors and introduce scatter for a non-monotonic relation (Lehmann et al. 2016; Stiskalek et al. 2022). As an alternative to the rank ordering of all galaxies without separating central and satellite galaxies, halo occupancy distribution modelling makes assumptions about the satellite distributions and uses halo mass functions as an input (Berlind & Weinberg 2002).

The penultimate step in this evolution of complexity is semi-analytical models, which incur a heavier computational burden as a trade-off for using a complete physical framework. The issue these models face is the number of free parameters for which constraints have to be found (see, for example, Somerville & Primack 1999; Lu et al. 2014). These include, but are not limited to, GALFORM as presented in Cole et al. (2000) and Baugh et al. (2018), GALICS and GALICS 2.0 (see Hatton et al. 2003; Cattaneo et al. 2017), and SAGE (Croton et al. 2016). Simpler analytical formalisms for the evolution of baryonic galaxy properties exist, for example the bathtub model by Bouché et al. (2010) and the reservoir model by Krumholz & Dekel (2012), as well as the equilibrium model (Davé, Finlator & Oppenheimer 2012; Saintonge et al. 2013; Mitra, Davé & Finlator 2015; Mitra et al. 2017).

Hydrodynamic simulations, which require the most computational resources and time, are generally considered the gold standard of an ab initio approach to galaxy formation and evolution (Vogelsberger et al. 2020). Particle-based and grid-based methods exist based on the foundation that baryonic matter can be modelled by treating gas as an ideal fluid. The former rely on computations of discrete masses, or particles, while the latter split the simulation volume into discrete spaces (Dolag et al. 2008; Somerville & Davé 2015). Influential recent hydrodynamic simulations include, for example, ILLUSTRIS and ILLUSTRISTNG (see Genel et al. 2014; Pillepich et al. 2018), EAGLE as introduced in Schaye et al. (2015), and HORIZONAGN by Dubois et al. (2016), as well as MUFASA and its successor SIMBA, the latter of which is used for the lion’s share of our experiments and is described in more detail in Section 2.1 (Davé, Thompson & Hopkins 2016; Davé et al. 2019).

In recent years, the physical sciences have experienced an exponentially rising interest in applying machine learning (ML) algorithms (Carleo et al. 2019). These developments include the acceleration of hydrodynamic simulations of galaxy formation and evolution, with Kamdar, Turk & Brunner (2016) providing one of the first examples by predicting baryonic properties in ILLUSTRIS using extremely randomized trees, commonly abbreviated as ‘extra trees’, an ensemble model based on decision trees. This heralded the dominance of tree-based ensembles in related research, including Agarwal, Davé & Bassett (2018) for MUFASA and Lovell et al. (2021) for EAGLE, as well as Jo & Kim (2019), McGibbon & Khochfar (2022), and de Santi et al. (2022) for ILLUSTRISTNG.

While these works apply tree-based ensembles, de Santi et al. (2022) also compare other algorithms such as k -nearest neighbours, light gradient-boosting machines, and feed-forward neural networks, and combine them with a linear regressor for improved predictions. Similarly, de Andres et al. (2022) use random forests, feed-forward neural networks, and natural and extreme gradient boosting on THE THREE HUNDRED data described by Cui et al. (2018), and provide an example of favouring boosting over tree-based methods. In this case,

extreme gradient boosting is reported to reach the best accuracy, while natural boosting retrieves the most scatter as a probabilistic regressor.

Jespersen et al. (2022) apply graph neural networks to ILLUSTRISTNG. However, they only use the dark matter version and compute baryonic properties with a semi-analytical model, effectively making the predictor emulate such a model instead of a hydrodynamic simulation. Other works exist that focus on alternative algorithms or hybrid approaches. Moster et al. (2021), for example, apply wide and deep neural networks (WDNNs) and reinforcement learning. However, the baryonic properties are calculated not through a hydrodynamic simulation but EMERGE, an empirical model that statistically links properties from surveys instead of simulating baryonic physics (Moster, Naab & White 2018). The prediction problem is flipped by von Martens et al. (2022), who retrieve dark matter haloes from their baryonic properties in ILLUSTRISTNG instead.

Moews et al. (2021), on the other hand, extend the equilibrium model by incorporating largest-progenitor merger trees,¹ and combine it with extra trees into a hybrid approach. One disadvantage of such analytical models is that they predict a limited set of properties. They can, however, be looped into both the training and prediction steps of machine learning algorithms by first predicting a limited set of properties with the analytical model and then using these predictions together with halo properties to predict the complete set of baryonics. Related work on hybrid models is done by LSST Dark Energy Science Collaboration (2020), who generate mock catalogues by combining empirical and semi-analytical models. This leads to the weighted Monte Carlo sampling of baseline catalogues to improve statistical realism.

Stiskalek et al. (2022), similarly to Moster et al. (2021), also use a WDNN approach for ILLUSTRISTNG and HORIZONAGN, and include a comparison with extra trees, but with a particular focus on reproducing the intrinsic scatter of the galaxy–halo connection. The reason is a well-known failure of commonly used machine learning algorithms when applied to hydrodynamic simulations to retrieve the expected scatter. The prediction of a Gaussian probability distribution for estimating targets does, of course, impose a constraint on the scatter. Still, probabilistic approaches to scatter reproduction are commonly encountered features in the literature (see, for example, Lehmann et al. 2016; Desmond et al. 2017; Mitra et al. 2017; Cao et al. 2020).

Machine learning methods that rely on standard metrics like the mean squared error (MSE) face certain limitations. The dominant drawback for this application area is that these models aim to recreate the provided data without knowledge of the underlying physical theory. Some prior work tries to solve this issue, for example the mentioned reports on weighted sampling through baseline libraries by LSST Dark Energy Science Collaboration (2020) and pre-prediction of baryonic subsets to aid the machine learning model (Moews et al. 2021). Another more direct way is injecting theory directly into the learning process of neural network architectures, with such models coined physics-informed neural networks (PINNs; Raissi, Perdikaris & Karniadakis 2019).

Originally developed for the finding of solutions for partial differential equations, PINNs are rooted in older research on neural networks for ordinary and partial differential equations (see, for example, Dissanayake & Phan-Thien 1994; Lagaris, Likas &

¹This might be a constant source of frustration for some readers of papers on these topics, as decision trees and merger trees are, apart from technically both having a tree-like structure, two entirely different concepts.

Fotiadis 1998). In such models, knowledge represented as suitable equations offers domain constraints. Various approaches exist that include physical equations into the loss function, while replacing a lengthy simulation with a machine learning model has become common practice in many scientific endeavours (McCarn Deiana et al. 2022). The line of thinking behind this approach is simple; instead of evaluating the model solely on prediction accuracy through metrics such as the MSE, additional loss components can enforce compliance with additional parameter relationships. For a more in-depth overview of this rapidly expanding area of physics-driven deep learning, we refer interested readers to the reviews by Karniadakis et al. (2021) and Cuomo et al. (2022).

PINNs are not entirely alien to the field of astrophysics in general. Recent work by Mishra & Molinaro (2021) targets the simulation of radiative transfer by minimizing the residual of the underlying transfer equations, while Martin & Schaub (2022) bypass inefficiencies of gravity models to learn representations of small-body gravity fields directly. Other recent examples include Branca & Pallottini (2022) on solutions for interstellar medium chemistry and the finding of quasi-normal modes of non-rotating black holes (Cornell, Ncube & Harmsen 2022). The development and application of these architectures offer a powerful way to add knowledge to observational evidence to enforce adherence to theoretical models in training the machine learning algorithm.

In this paper, we present the first way to incorporate the concept of PINNs into the active research field of the completion of dark matter-only information in cosmological simulations with baryonic properties. To achieve this goal, we modify the training process of a deep learning model through two different novel extensions of the loss function. The first part is based on the standard approach of injecting physical theory into these models and adds the stellar-to-halo mass relation (SHMR) as described by Moster et al. (2010) into the training process. This double power law is subsequently used by Moster et al. (2018) to parametrize the instantaneous baryon conversion efficiency, meaning the efficiency with which gas is transformed into stars, and provides a theory-oriented constraint.

While this relation can, of course, be argued to be an empirical prescription rather than physical theory, an argument for being physics-informed is given through the baryon conversion efficiency, with higher M_h values leading to the collection of more baryons. In detail, this link is more complicated, and although sufficient for an exploratory study on the benefit of these types of constraints, we discuss the limitations. In a broader sense, the description as physics-informed stems from the name of the model family itself.

The second part, which presents another novel addition, forces the model to recreate the scatter of the underlying hydrodynamic simulation by including the Kullback–Leibler divergence (KLD) as an asymmetric distance measure between approximations of the probability distributions for training targets and associated model predictions (Kullback & Leibler 1951; Ferdosi et al. 2011). This differs in its motivation from the inclusion of the SHMR from Moster et al. (2010), as it targets the recreation of scatter and not an improvement in mean accuracy. In doing so, we show that our extensions of existing machine learning approaches in baryon inpainting from dark matter halo properties are a powerful tool for modern cosmological simulations.

Compared to the baseline model, our results demonstrate improvements in both predictive accuracy and the reproduction of scatter, and show suitable correlations between target variables and model predictions. The contribution pertains to the broader field of machine learning in astrophysics, including adding physical theory into predictive models and using distributional loss components.

The remainder of this paper is structured as follows. In Section 2, we provide an overview of our machine learning approach and data. Section 2.1 describes the SIMBA suite of cosmological simulations and our data set. Section 2.2 covers the functionality and justification of our baseline PINN model; Sections 2.3 and 2.4 introduce extensions of the loss function for the SHMR and a distribution comparison between predictions and target values, respectively. Section 3 presents our experiments and their results. Specifically, Section 3.1 explains the fitting of our theory constraint to SIMBA and the weighting of loss function components. Section 3.2 provides the results for our predictions, and Section 3.3 shows correlations for individual data point accuracy. Section 4 discusses our findings, limitations of our approach, and follow-ups. Lastly, Section 5 provides our conclusions.

2 DATA AND METHODOLOGY

2.1 Simulation data from the SIMBA suite

The SIMBA simulation models the co-evolution of gas and dark matter within an expanding metric using the GIZMO code (see Hopkins 2015), which itself is based on the GADGET-2 (Springel et al. 2005). It employs the meshless finite mass hydrodynamics method, which marries the convenience of a mass-conserving particle-based code with the shock and instability-capturing advantages of a Riemann solver-based scheme.

Many so-called sub-grid processes have been added to GIZMO to model the formation and evolution of galaxies. These include radiative cooling and photoionization heating, chemical enrichment from stellar evolution, the formation of stars and supermassive black holes, the energy release (‘feedback’) from supernovae and black hole accretion discs, and the growth and destruction of dust. The complete model detailing these sub-grid prescriptions is described in Davé et al. (2016, 2019).

SIMBA simulations begin in the linear regime at redshift $z = 249$, and are evolved to $z = 0$, meaning today. 151 snapshot outputs are stored at various redshifts along the way. The main SIMBA run models a random cube of 147 Mpc (comoving) on a side, represented by 1024^3 gas elements and 1024^3 dark matter particles. The minimum (adaptive) spatial resolution is 0.7 kpc. The simulation assumes a *Planck*-concordant cosmology (see Planck Collaboration XIII 2016) of $\Omega_m = 0.3$, $\Lambda = 0.7$, $H_0 = 68 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\sigma_8 = 0.82$, and $n_s = 0.97$. This results in a mass resolution of 1.8×10^7 per gas element and 9.5×10^7 per dark matter particle.

Each snapshot is analysed using the CAESAR galaxy/halo catalogue package. For each halo identified within SIMBA using GIZMO’s native three-dimensional (3D) Friends-of-Friends (FoF) finder, CAESAR identifies galaxies as collections of stars and dense gas via a 6D FoF algorithm. The most massive galaxy within a halo is defined as the central, and the others are satellites. A large range of physical and photometric properties are computed for each halo and galaxy. For this work, the key galaxy quantities are the stellar mass (M_*), star formation rate (SFR), SFR-weighted gas-phase metallicity (Z), neutral hydrogen mass (M_{HI}), molecular hydrogen mass (M_{H_2}), and central supermassive black hole mass (M_{BH}).

These properties are obtained by summing the relevant particles in each galaxy. For haloes, the relevant quantities are the total mass (M_h), the dark matter half-mass radius (r_h), and the dark matter velocity dispersion (σ_h). The catalogues are stored as Hierarchical Data Format version 5 (HDF5) files, and CAESAR provides a simple yet powerful PYTHON-based access interface. The SIMBA snapshots

and catalogues are all publicly available online for use by the scientific community.²

Our work uses central galaxies from the `m100n1024_151` version of the SIMBA main runs. Entries for which $M_{\text{BH}} = 0$ are dropped to avoid zero-mass black holes distorting the distributions of predictions in our experiments, which is in line with related research covered in Section 1. This allows predictive models to focus on the recovery of accurate shapes for exploratory studies, while a sharp cut-off from seeding processes in hydrodynamic simulations complicates this in continuous prediction spaces. These considerations, as well as potential solutions, are further discussed in Section 4.

Similarly, we restrict the range of included halo masses to $11 \leq \log_{10}(M_h) \leq 14$, following the same related research. Apart from these pre-processing steps, we refrain from any further alterations. While suitable data selections could lead to improved predictions, the goal of this paper is to maintain generalizability. For the model training and prediction processes, we apply min–max normalization to input and target variables to scale values within the same interval [0, 1] and then revert predictions to their proper scales.

2.2 Physics-informed neural network model

The concept of PINNs rests on the assumption that data-driven statistical learning can be enhanced with domain knowledge. This integration can be implemented using different frameworks, as covered in the overview of Section 1 and related reviews (Karniadakis et al. 2021; Cuomo et al. 2022). One approach is to generate more data specifically crafted to enforce domain knowledge into the model. While simple, this approach requires large amounts of additional data to cover a broad region of the input-variable space, effectively introducing observational biases into the inductive analysis (Yang & Perdikaris 2019; Kashefi, Rempe & Guibas 2021).

Another approach involves designing specific learning algorithms that embed knowledge into the learning architecture, for example, convolutional neural networks for images and speech (see LeCun & Bengio 1995), graph neural networks (see Zhou et al. 2020; Wu et al. 2021), and networks for Hamiltonian systems, among others (Jin et al. 2020). Implementing this approach is difficult because embedding physical laws within a neural network architecture is limited to simple processes. Complex processes require architectural designs that cannot be easily realized with current learning frameworks; even relatively simple processes need complex and elaborate designs. Regarding practical applications, a second complication of these purpose-designed models is creating a network specific to a given problem, making the transfer to new domain applications time-consuming.

The third approach we adopt here enriches the loss function to explicitly incorporate constraints, usually as partial differential equations (Lagaris et al. 1998; Raissi et al. 2019). The benefit is a decoupling between the machine learning strategy and the embedded knowledge; the corresponding framework has broader applicability since the learning architecture is built independently of the underlying physical laws. Normally, the loss function is extended to include a measure of the accuracy of each prediction and the degree of alignment with a physical constraint.

The ability of PINNs to generalize well goes beyond approximation theorems. The representational power of neural networks is well known; under limited assumptions, any continuous function can be approximated to an arbitrarily close fit using a neural network with a

finite number of hidden nodes and one hidden layer (Cybenko 1989; Yarotsky 2017). It should be noted that representational power is not equivalent to generalization power. Adding domain knowledge to the loss function improves the generalizability of PINNs by reducing the bias component of error while keeping the variance component under control (Hastie, Tibshirani & Friedman 2009).

In traditional PINNs, the final loss, \mathcal{L}_f , is a weighted combination of two losses: one is the data-driven empirical loss, \mathcal{L}_s , and the other is the domain-knowledge constraint, \mathcal{L}_k , with respective weights w_s and w_k , so that

$$\mathcal{L}_f = w_s \mathcal{L}_s + w_k \mathcal{L}_k. \quad (1)$$

The first term is the conventional loss obtained in traditional neural networks through methods such as the squared difference between predictions and target values. Training a network NN_θ that aims to find a near-optimal parameter vector θ , for example, through gradient descent, yields a surrogate for a solution to processes such as complex simulations. The corresponding output $\hat{f}(x)$, where x is a feature vector, serves to adjust the weights during training by looking at a loss function $\mathcal{L}_s(\hat{f}(x), y)$. Here, y is the true response variable, or the target value from the underlying hydrodynamic simulation in our case, and $\hat{f}(x)$ is our model-provided estimate.

The additional term, \mathcal{L}_k , adds domain constraints, usually as differential equations. Specifically, the term captures the partial differential residuals as covered in Section 2.3. In our study, the inputs to the neural network are $x := [M_h, r_h, \sigma_h]$ corresponding to dark matter halo mass at present, dark matter half-mass radius, and dark matter halo velocity dispersion, respectively.

The outputs are $y := [M_*, \text{SFR}, Z, M_{\text{H}_1}, M_{\text{H}_2}, M_{\text{BH}}]$ corresponding to stellar mass, SFR, metallicity, neutral and molecular hydrogen masses, and black hole mass, respectively. During training, we aim to minimize the residual sum of squares,

$$\mathcal{L}_s = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M (y_i^{(j)} - \hat{f}^{(j)}(x_i))^2. \quad (2)$$

Here, N is the number of examples in the training data, while M is the respective target space subject to prediction.

As the subsequent distribution-based loss is calculated over each epoch’s output for the training data, we take the arithmetic mean of the residual sum of squares for each training iteration. Domain knowledge minimizes a different loss function, \mathcal{L}_k , forcing the final model to obey the constraint from physical knowledge. Here, we assume points $\{x_i\}$ sampled across the entire input space. In the next section, we explain the domain loss in detail. With the defined loss functions, the neural network NN_θ is trained to obtain parameters θ using efficient optimization methods, such as gradient descent. The weights w_s and w_k enable different contributions to the final loss function and can be tuned automatically as part of the optimization process.

2.3 Inclusion of baryon conversion efficiency

As the name suggests, the SHMR links the stellar mass of a given galaxy to the dark matter halo mass. Suitable parametrizations have been shown to reflect the galaxy mass function observed in the third data release of the Sloan Digital Sky Survey (see Panter et al. 2007) more closely, as it does not assume a constant SHMR. The instantaneous baryon conversion efficiency ϵ is the rate at which gas is transformed into stars. This efficiency can be described by an SHMR parametrized through a double power-law model as shown by Moster et al. (2010).

²<https://simba.roe.ac.uk>

The latter introduces a parametrization that follows observations by avoiding a surplus of galaxies at low and high masses. For two slopes β and γ that are used to determine the decrease in efficiency at lower and higher masses, respectively, the parametrization takes the form

$$\epsilon(M, z=0) = 2\epsilon_N \left(\left(\frac{M_h}{M_1} \right)^{-\beta} + \left(\frac{M_h}{M_1} \right)^{\gamma} \right)^{-1}. \quad (3)$$

Here, ϵ_N is the normalization, while M_1 denotes the characteristic mass at which the respective efficiency is the same as its normalization. Moster et al. (2018) use this to parametrize the instantaneous baryon conversion efficiency, showing that peak conversion efficiency takes place at halo masses similar to the characteristic mass,

$$M_{\max} = M_1 \left(\frac{\beta}{\gamma} \right)^{(\beta+\gamma)^{-1}}, \quad (4)$$

with the general assumption of $\beta, \gamma > 0$. The integrated baryon conversion efficiency is dependent on redshift (see Moster, Naab & White 2013), and the mentioned work allows for parameters of the instantaneous efficiency to vary, with

$$\begin{aligned} \log_{10} M_1(z) &= M_0 + M_z(1 - (z+1)^{-1}) \\ &= M_0 + M_z \left(\frac{z}{z+1} \right), \end{aligned} \quad (5)$$

and with the normalization and slopes given by

$$\begin{aligned} \epsilon_N(z) &= \epsilon_0 + \epsilon_z(1 - (z+1)^{-1}) = \epsilon_0 + \epsilon_z \left(\frac{z}{z+1} \right), \\ \beta(z) &= \beta_0 + \beta_z(1 - (z+1)^{-1}) = \beta_0 + \beta_z \left(\frac{z}{z+1} \right), \\ \gamma(z) &= \gamma_0. \end{aligned} \quad (6)$$

As we operate at $z = 0$, these considerations are simplified and need only be optimized for single values. In Section 3.1, we will cover this optimization for SIMBA data using maximum likelihood estimation (MLE), with ranges provided by prior research on this parametrization. This view on the SHMR can then be included, for a given data set size of N , into the loss function of equation (1) as

$$\mathcal{L}_k = \sum_{i=1}^N \left(\frac{\hat{M}_*}{M_h} - 2\epsilon_N \left(\left(\frac{M_h}{M_1} \right)^{-\beta} + \left(\frac{M_h}{M_1} \right)^{\gamma} \right)^{-1} \right)^2. \quad (7)$$

This injection of domain knowledge provides an additional constraint for the model, which subsequent experiments show to benefit the model in recovering the mean relation better.

Including the M_*-M_h relationship used in this study is physics-informed through the link to the baryon conversion efficiency, which is given through the fact that bigger dark matter haloes are able to collect more baryons, thus forming more stars. This is, of course, a simplified treatment, but sufficient for exploratory work, and drawbacks are further discussed in Section 4.

2.4 Constraints from predictive distributions

Reproducing the scatter found in hydrodynamic simulations when predicting baryonic properties based on dark matter haloes is a known challenge in the literature (Cui et al. 2018; Stiskalek et al. 2022). As described in Section 1, probabilistic approaches relying on parametrized distribution families are common. In this work, we target the reproduction based on the training data distribution directly

by introducing a second extension to the standard loss function in equation (2) (see Section 2.2).

The KLD is a statistical distance measure to assess the difference between two distributions. Introduced by Kullback & Leibler (1951), it has found various applications in astrophysics in recent years (see, for example, Ben-David, Liu & Jackson 2015; Hee et al. 2016; Moews et al. 2019; Nicola, Amara & Refregier 2019). For a given reference distribution P and proposal distribution Q , the KLD can be written as

$$D_{\text{KL}}(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)}, \quad (8)$$

or, correspondingly, with an integral for absolutely continuous probability distributions. One important point is that the KLD is not a distance metric due to its status as an asymmetric difference measure. This means that

$$D_{\text{KL}}(P||Q) \neq D_{\text{KL}}(Q||P), \quad (9)$$

as it calculates a directional information loss when approximating P via Q . It also does not satisfy the triangle equality,

$$d(a, c) \leq d(a, b) + d(b, c), \quad (10)$$

with points $\{a, b, c\} \in M$ for a given metric space M . Given the above, the KLD is applicable only when a ‘true’ reference distribution is used. Fortunately, this is the case here, as we want to calculate the difference between the distributions of model predictions and their respective targets.

Following Fussell & Moews (2019), who propose the incorporation of the KLD into the loss function in the context of generative modelling – although without implementing the proposal due to conflicting success metrics – we extend equation (1) to

$$\mathcal{L}_f = w_s \mathcal{L}_s + w_k \mathcal{L}_k + w_{\text{KL}} \mathcal{L}_{\text{KL}}. \quad (11)$$

Here, the KLD is part of the overall loss function through a normal assumption placed on the target and prediction distributions, as the loss needs to remain easily differentiable for error backpropagation during training. This means that

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(\mathcal{N}(\mu, \sigma) || \mathcal{N}(\hat{\mu}, \hat{\sigma})), \quad (12)$$

for target mean and variance, μ and σ , and corresponding predictions, $\hat{\mu}$ and $\hat{\sigma}$. In doing so, and as described in the overview in Section 1, we impose an additional constraint on the learning process by forcing the model to recreate the approximate distribution of the underlying training data. This extension focuses on the more accurate reproduction of the scatter found in the underlying simulation, not on the improvement of the mean accuracy. In later sections, we will see how this avoids an underprediction of tails in the scatter of the baryonic properties of interest.

Fig. 1 provides a schematic for the complete model incorporating both the SHMR from Section 2.3 and the KLD, which we dub the ‘Hybrid’ model going forward. The lower part of the figure shows the neural network model, going from input data containing information from a set of dark matter halo properties to baryonic predictions, \hat{y} , as the model output. The components of the loss function are encased by a dashed line, listing the MSE, the KLD introduced in this section, and the SHMR. The respective weights of these components can be selected based on the predictive performance of the neural network, which we cover as part of our experiments in Section 3.1.

The normal assumption made for the respective distributions is, of course, not without fault. In particular, it remains to be seen in the results whether this assumption can lead to a decrease in mean predictive accuracy for sufficiently non-normal distributions.

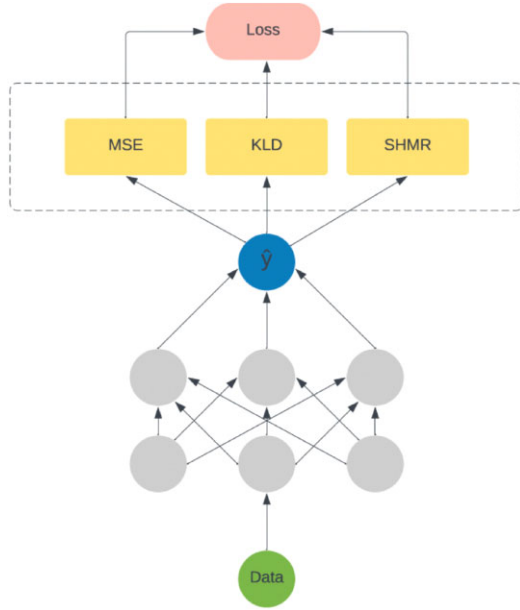


Figure 1. Illustration of the hybrid model. The loss function used in this framework is a combination of the MSE, the KLD, and the SHMR.

However, it provides an approximation that is differentiable for the purpose of backpropagation during training and is computationally cheap enough to result in reasonable training times on small numbers of graphics processing units (GPUs) without requiring access to stacks in large-scale supercomputing solutions. We reserve part of the discussion in Section 4 to list possible alternatives in related follow-up research.

3 EXPERIMENTS AND RESULTS

3.1 The stellar-to-halo mass relation in SIMBA

We first train a simple feed-forward neural network, a multilayer perceptron (MLP), using only the MSE in the loss function, and then use this baseline model to predict all six parameters. In order to efficiently explore and optimize the hyperparameter space, we adopt a sequential model-based optimization (SMBO) approach. The SMBO algorithm iteratively updates the surrogate model, explores the hyperparameter space by sampling new configurations, evaluates their performance using the objective function, and refines the search process based on the observed results, allowing for a reasonably fast and continuous optimization. For this, we use the Optuna hyperparameter optimization framework to determine the optimal number of layers and neurons in the hidden layers (Akiba et al. 2019). The search space was set to be between 3 and 10 layers and 30–100 neurons per layer, following common choices for similarly structured problems.

After invoking the optimization module, the best configuration displayed 9 hidden layers with 80 artificial neurons, each with a rectified linear unit (ReLU) activation function, equally following current practice. In addition, we split off 10 per cent of the data set as a validation set to perform a search on the weights, freely chosen in $[0, 1]$, for the loss function components, $w = [w_s, w_k, w_{KL}]$. These weights are optimized through a combination of grid-search Bayesian optimization and meta-heuristics using Optuna (Akiba et al. 2019). They are optimized to minimize the overall loss on a validation set, averaged over 10 runs. As in most optimization techniques, the

Table 1. Optimization of best-fitting loss function component weights using SMBO. The table lists the MSE and MAPE for various loss function weight combinations in the hybrid model. The values in brackets list the associated standard deviations.

(w_s, w_k, w_{KL})	MSE	MAPE
(1, 0.6, 1.3)	0.025 79 (0.0016)	1.098 (0.0532)

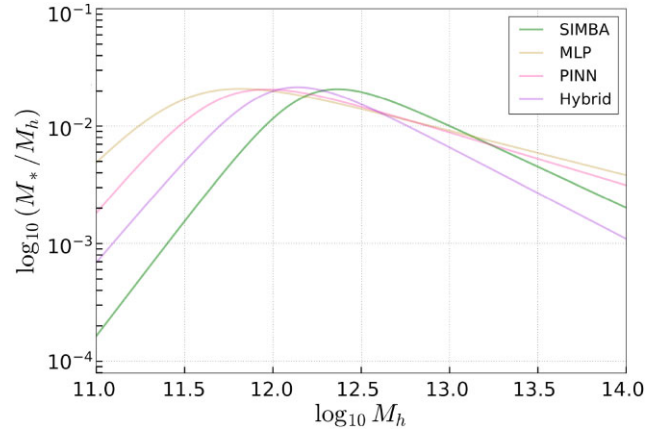


Figure 2. Double power-law fitting of model predictions for the baseline MLP, the PINN, the hybrid PINN + KLD model, and the SIMBA data.

trivial solution of assigning the least possible value to each weight is avoided by imposing a constraint on the weights, meaning the sum of the weights is a constant.

While an MLE approach is infeasible due to the model retraining for each run, this allows us to gauge a suitable combination of weights. The best-fitting combination is tested for 10-fold cross-validation using the MSE and the mean absolute percentage error (MAPE), with the results listed in Table 1. As Fig. 2 and the top panels of Fig. 3 show, this basic model can recover the basic pattern of the SHMR represented in equation (3). However, the pattern exhibits a diminished variance at lower and higher halo masses and is expectedly flat around the characteristic M_{\max} value.

As a result, the predictions show a flatter relation than 1:1 relative to the true distribution from SIMBA, indicating a bias. We note that SIMBA’s galaxy sample is limited in stellar mass, which results in a diagonal completeness limit in the SHMR. This leads to galaxies with a high M_*/M_h ratio for their halo mass being preferentially included in the sample; such asymmetric completeness (known as Malmquist bias) is common in astrophysics, but can be difficult for machine learning algorithms to recover.

Our PINN, on the other hand, is constructed by encoding equation (3) into the loss function as shown in Section 2.3. The parameters of the resulting loss component shown in equation (7) are estimated through MLE. MLE is performed to optimize the parameters of equation (3) and, accordingly, of the PINN part of the loss function in equation (7). Here, the likelihood function is the density function regarded as a function of given parameters θ ,

$$\mathcal{L}(\theta) = \sum_{i=1}^n f(x_i|\theta), \quad \text{with } \theta \in \Theta, \quad (13)$$

where Θ is the corresponding parameter space and $f(\cdot)$ is the probability density function. We can estimate the optimal combination of

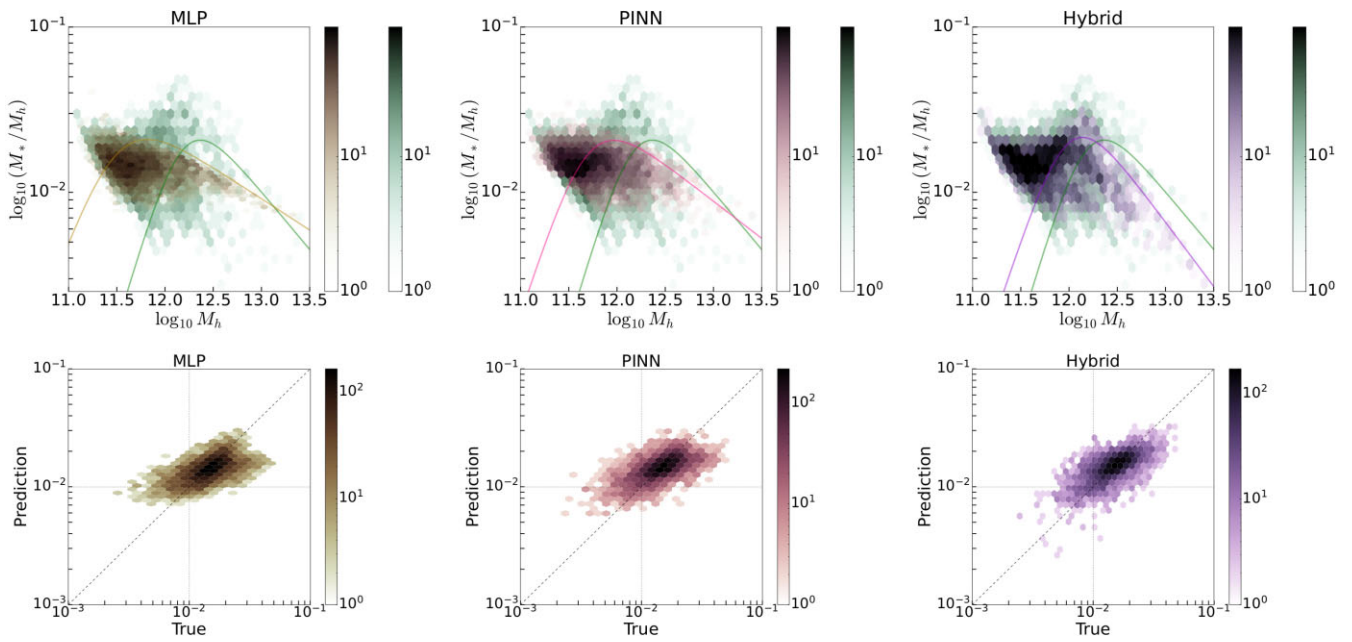


Figure 3. Top: Hexagonal joint histograms of the predicted SHMR from different models and the true SHMR calculated from the SIMBA test set, as well as corresponding double power laws as described in equation (3). From left to right, the panels show the results from the baseline MLP, the PINN, and the hybrid PINN + KLD model. The SIMBA data are shown in green. Bottom panels: Hexagonal joint histograms of the predicted versus target SHMR for different models and the test set, with colour coding as described above. The diagonal line indicates $x = y$ for the resulting correlation plots.

parameter values of interest, $\hat{\theta}$, as

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \mathcal{L}(\theta|x), \quad (14)$$

which results in a choice of parameter values that make the observed data the most probable. Different values in a wide range have been reported on equation (3) (Shankar et al. 2017; Kravtsov, Vikhlinin & Meshcheryakov 2018; Behroozi et al. 2019). Instead of directly applying these numbers, they become the initial guess of our estimator, and the range of values is considered a constraint. The best-fitting parameters for the test set are $\epsilon_N = 0.02$, $M_1 = 12.19$, $\beta = 1.1$, and $\gamma = 0.78$. With this extra constraint, the predicted SHMR pattern better recovers the curvature around M_{\max} in Fig. 3. Equation (3), however, as a double power-law function, this additional constraint cannot introduce the scatter found in the cosmological simulation and still displays a slightly biased recovery of the true SHMR relation.

To retrieve the scatter, we need a way to provide information as part of the loss function, effectively forcing the model to recreate it. For this, we extend the loss function with an additional component using the KLD, which we introduced in Section 2.4. This distributional divergence measurement has a history of being used in various cosmological applications, as discussed in Section 1, including the study of properties and auxiliary observational data on baryonic physics (Yasin et al. 2023). The resulting predictions from this hybrid scheme track the respective target values more faithfully and with less bias.

We calculate the mean and covariance matrix of the training data and construct a multivariate Gaussian distribution in six dimensions, one for each prediction target property. By including the KLD-based loss in equation (8) into the overall loss function, we create a hybrid model combining the PINN with a KLD measurement as written in equation (11). Table 2, listing the MSE as well as the coefficient of determination (R^2) and the Pearson correlation coefficient (ρ), demonstrates the benefit on the predictive power of these models.

Table 2. Performance comparison between different models. The table lists the coefficient of determination (R^2) and Pearson’s correlation coefficient (ρ) for different neural network models corresponding to Figs 2 and 3, as well as the best-fitting relation following Moster et al. (2010). The hybrid model is a combination of PINN with KLD in the loss.

Model	MSE	R^2	ρ
Moster relation	0.545	0.730	0.855
MLP	0.024	0.821	0.906
PINN	0.023	0.827	0.910
Hybrid (PINN + KLD)	0.020	0.848	0.921

We plot the SHMR relative to the halo mass in Fig. 3 to provide a visual overview. This confirms that the hybrid approach between extra physical knowledge and distributional adherence outperforms both the baseline model and the PINN alone. The latter provides information on the mean trends of the SHMR, while the KLD helps the neural network to mimic the substantial scatter, which an exact equation does not provide. Compared to the baseline MLP, the figure shows that the distributional component helps predict the SHMR more accurately at lower halo masses while tracing the downward scatter at higher halo masses.

To better understand the correlation between targets and predictions, the bottom panels of Fig. 3 show the SHMR values from SIMBA on the horizontal axis, with SHMR values from model predictions on the vertical axis. The result would follow the plot’s diagonal line for a perfect model. The visuals, which indicate a better correlation for the PINN compared to the MLP, and for the hybrid model compared to both, are confirmed by the corresponding correlation measurements in Table 2. As discussed in Section 2.4, the normal assumption for the KLD constraint could be argued to enable a decrease in mean accuracy in some cases, and these preliminary results demonstrate the overall ability of the model to use the constraint in a non-detrimental manner. Apart from the SHMR, the KLD also improves the model’s

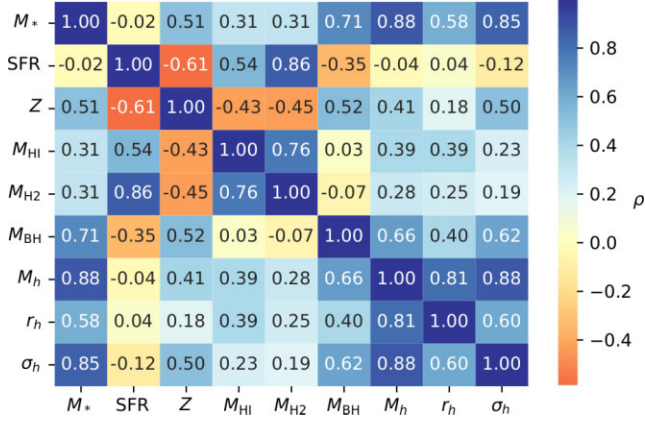


Figure 4. Correlation matrix of input and output parameters. The colour bar indicates the Pearson correlation coefficient values, calculated for variables in the SIMBA data set used in this work.

prediction ability on other baryonic properties, which we will cover, together with potential explanations for an improvement instead of no discernible difference in case of no negative effect of the normal assumption, in Sections 3.2 and 4.

3.2 Predicting baryonic properties from dark matter haloes

After establishing the effect on SHMR prediction, we expand our analysis to the full set of six baryonic properties, $\{M_*, \text{SFR}, Z, M_{\text{HI}}, M_{\text{H2}}, M_{\text{BH}}\}$. To better understand the cause of these results, we calculate ρ values between the entire set of available parameters in Fig. 4.

Fig. 5, where these variables are plotted against the halo mass, shows that each parameter is reasonably well predicted upon visual inspection. In particular, the model excels for M_* and M_{BH} , while the performance on SFR and M_{H2} is subject to a scatter taper at higher halo masses. In the values of Fig. 4, we can see that the correlations between the dark matter halo properties used as inputs, meaning

$\{M_h, r_h, \sigma_h\}$, for these variables are considerably lower compared to the rest of the investigated properties. At the same time, good results are equally reflected in strong correlations for stellar and black hole masses. As covered in Section 3.1, this is aided by the additional SHMR constraint in the loss function.

Following this, we explore secondary correlations between different variables, analogous to a similar analysis performed by Agarwal et al. (2018). Fig. 4 indicates that galaxies with higher SFR exhibit positive correlations with M_{HI} and M_{H2} , and a negative correlation with Z. To test whether our model correctly learns the split in the specific SFR,

$$\text{sSFR} = \frac{\text{SFR}}{M_*}. \quad (15)$$

We plot these properties against the stellar mass and colour data points using the distance in sSFR values from the mean M_* -sSFR relation, $\Delta \log_{10} \text{sSFR}$, in Fig. 6. The colour coding highlights the ML model's recovery of the second-parameter correlation through the sSFR, which is also seen in SIMBA. This means that the data points plotted in the figure are predictions of our model, not data points extracted from the simulation. The resulting mean scaling relations are drawn for model predictions and the corresponding data from the underlying SIMBA simulation, with the former tracing the results of Agarwal et al. (2018) and Davé et al. (2019).

Both mean scaling relations are obtained by fourth-order polynomial fitting, and we provide the 6th to 93rd percentile range for SIMBA data indicated by grey shading. In doing so, the figure provides the simulation distribution as a reference in addition to the colour-coded model predictions. We use a bin size of 0.2 for the $\log_{10} M_*$ values in solar masses along the horizontal axis for the latter. The divergence between confidence intervals and the mean scaling relation for metallicity in SIMBA at lower stellar masses is an artefact of a small number of data points available at this range, but we include this left-hand interval to exemplify potential peculiarities encountered in such analyses.

As our approach is a prediction problem using approximations, it is not perfectly consistent, just like other works on machine learning for baryonic inpainting into dark matter haloes. Here, the

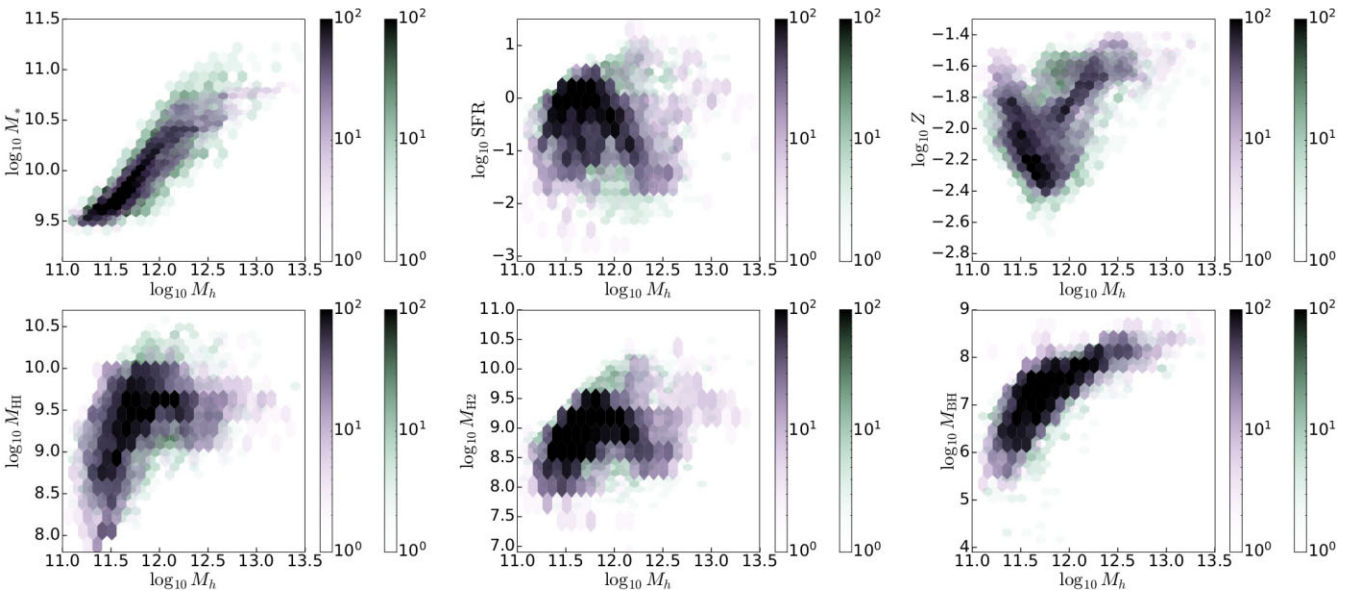


Figure 5. Hexagonal joint histograms of target values and model predictions. The panels show, for the set of six baryonic properties of interest, plots against the dark matter halo mass. Model predictions are in purple, while the SIMBA data are in green.

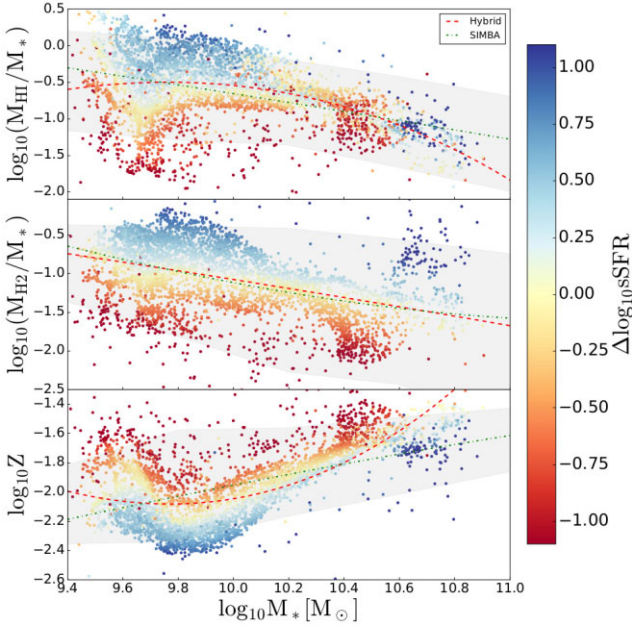


Figure 6. Secondary correlations at $z = 0$. The panels show, from top to bottom, neutral hydrogen, molecular hydrogen, and metallicity as a function of the stellar mass. Mean scaling relations are drawn with green dot-dashed lines for SIMBA and red dashed lines for our hybrid model. Galaxies are coloured by the distance from the mean M_* -sSFR relation.

underprediction of scatter at higher stellar masses is notable for M_{H_2} compared to SIMBA. At the same time, the relationships of Z , M_{H_1} , and M_{H_2} to the sSFR are preserved around the mean scaling relations, although in a cleaner split than is the case in the target values provided by the SIMBA cosmological simulation suite.

While similar patterns can be found in SIMBA, the discrepancies in polynomial fits can be attributed to the wave-like nature common to neural network predictions. This can depend on the activation function and refers to the non-linear change in the vertical distribution of data points in Fig. 6. While common features in these types of models, the present exploratory study demonstrates the viability of our approach and is discussed further in Section 4. MLP predictions without extensions to the loss function severely underpredict the scatter across the board at higher stellar masses. The overall distribution of predictions stays within the SIMBA-indicated areas. The recovery of secondary correlations between SFR (or gas content), metallicity, and stellar mass, known as the fundamental metallicity relation, is an important success of this machine learning framework.

3.3 Correlations for separate prediction targets

In this section, we further analyse the quality of model outputs. Similar to our visualizations in Section 3.2, we wish to compare the model predictions to the underlying target data from the SIMBA suite of cosmological simulations, but separately for different baryonic properties. Fig. 7 shows kernel density estimates (KDEs) for all six target properties. For the bandwidth optimization, we make use of Scott’s rule as the default heuristic in a variety of statistical software

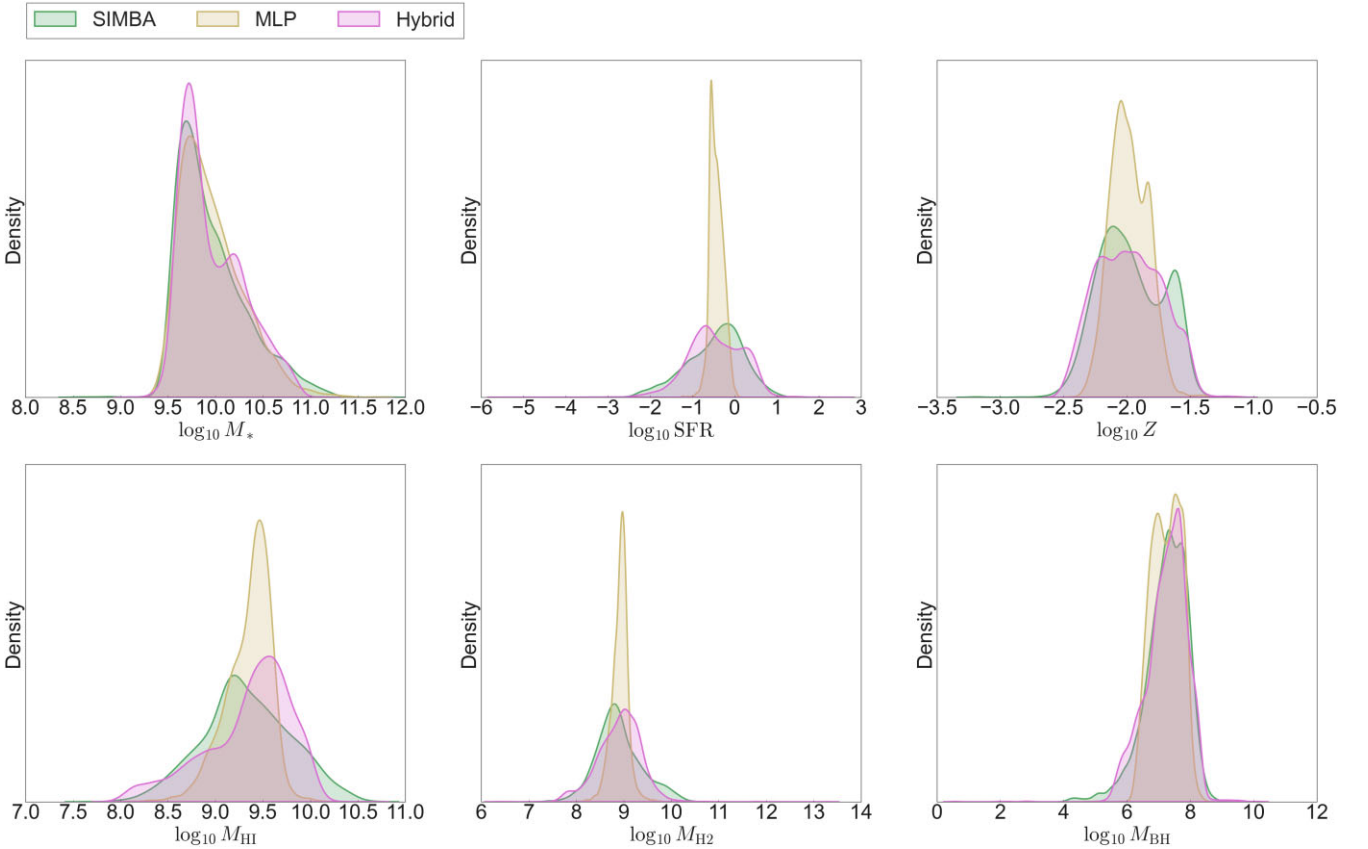


Figure 7. Density plots for predictions on all six baryonic target properties. The panels show separate KDEs for stellar mass (M_*), SFR, metallicity (Z), neutral hydrogen (M_{H_1}), molecular hydrogen (M_{H_2}), and black hole mass (M_{BH}). The corresponding values from the underlying SIMBA test data set are plotted in yellow, baseline MLP predictions in green, and results for the hybrid model are in red.

Table 3. Statistical validation for the full experimental run of hybrid model. The table lists the MSE, the coefficient of determination (R^2), Pearson’s correlation coefficient (ρ), and the KLD for each output variable.

Variable	MSE	MLP			Hybrid			
		R^2	ρ	KLD	MSE	R^2	ρ	KLD
M_*	0.023	0.827	0.901	0.049	0.020	0.847	0.920	0.045
SFR	0.638	0.050	0.267	0.245	0.423	0.080	0.284	0.226
Z	0.052	0.277	0.527	0.114	0.044	0.352	0.593	0.078
M_{H1}	0.221	0.215	0.450	0.215	0.162	0.235	0.480	0.191
M_{H2}	0.308	0.087	0.284	0.293	0.208	0.097	0.313	0.290
M_{BH}	0.276	0.466	0.682	0.090	0.214	0.502	0.709	0.091

packages for a data set length $|X|$ with a given dimensionality,

$$\hat{\beta} = |X|^{-(\dim(X)+4)^{-1}}. \quad (16)$$

In addition, we also plot the same for the baseline multilayer perception used in previous comparisons. These density plots reveal a reasonably close agreement regarding scatter between SIMBA and our model’s predictions for all baryonic properties. While the baseline model is already performing well on the stellar mass and, to a degree, on the black hole mass, it visibly struggles with the scatter of the remainder of the target properties.

Our hybrid model shows major improvements in the distributions of SFR, Z , M_{H1} , and M_{H2} . For the MLP, the predicted distributions form a sharp peak around the mean, demonstrating especially difficult-to-retrieve scatter for the SFR and the molecular hydrogen mass. While the injected physical knowledge on the SHMR does not aid with scatter predictions but instead with the accuracy of associated properties, the KLD loss component contains extra information on distributions that allows for the production of the necessary scatter.

To quantify these improvements, we compare the baseline model and our hybrid approach in Table 3, providing MSE, R^2 , ρ , and KLD values as for previous comparisons. Here, we can see the expectedly high correlations for M_* , M_{BH} , and, to a lesser degree, Z , comparable to prior research on different hybrid approaches in this area (Moews et al. 2021).

While the latter does not use neural network architectures or direct injection of information into a loss function, they build an additional analytical model into the prediction pipeline and use full largest-progenitor merger trees, which demonstrates the capability of our loss function extensions.

Despite the correct prediction of mean relations, results for SFR, M_{H1} , and M_{H2} are not quite as perfect when viewed next to the remaining baryonic properties. One reason is that we use a single neural network model to predict all variables simultaneously instead of using separate models. The downside is that the model cannot focus solely on a 1D prediction, but this is not desirable, as it also bars the model from learning the connections between different target variables. The result, however, is that properties with a smaller correlation coefficient in terms of input variables, as shown in Fig. 4, are more difficult to recover.

For the SFR and both hydrogen masses, we can see in Fig. 6 that there is more scatter at higher dark matter halo masses for these targets in particular, which usually have higher scatter compared to M_* , Z , and M_{BH} (Shankar et al. 2017). Another reason for prediction errors is likely found in the normal assumption. To calculate the KLD loss and implant it into the backpropagation, we assume a 6D Gaussian distribution, which is not completely true for our data. More complex approximations, however, are beyond the scope of this initial study due to often prohibitive computational costs and

the challenge of maintaining differentiability but are discussed in Section 4.

There are potential solutions, which include the injection of more information, for example, on scatter relations, into the loss function. Another pathway is applying more data, such as merger tree information, into the input features. The main purpose of this paper is to show that PINNs, which make use of extra physical knowledge, together with a second distributional loss component, can improve model performance in this challenging area and demonstrate the utility that current developments in deep learning approaches can provide for the simulation of baryonic properties.

4 DISCUSSION AND LIMITATIONS

The application of modern machine learning methods to the completion of N -body information has emerged as a growing area of interest in recent years. Here, the sometimes mentioned unreasonable effectiveness of tree-based ensembles, most commonly random forests and extra trees, is that these models are comparatively simple in their functionality, and yet, as covered in Section 1, are frequently found to outperform neural network architectures in this area. For standard feed-forward frameworks, the universal approximation theorem even guarantees the ability to represent arbitrary functions under very limited assumptions (Cybenko 1989; Hornik, Stinchcombe & White 1989; Maierov & Pinkus 1999). These theoretical capabilities do not, however, touch upon the ‘learnability’ of these functions, as there are a wide variety of hyperparameters to consider.

While tree-based ensembles are simple, they too are not immune to such limitations, and more recent research has discussed expected error rates as non-monotonous functions concerning the number of ensemble constituents to be set by the user (Probst & Boulesteix 2018). With the first application of machine learning to this specific area having arguably started with Kamdar et al. (2016), recent works such as Moster et al. (2021), Jespersen et al. (2022), and Stiskalek et al. (2022) have demonstrated the performance potential of suitably chosen neural network architectures.

One issue that remains is the limitations set by the data used for training and prediction in supervised machine learning models. A popular adage from the early days of computer science still widely used is ‘garbage in, garbage out’, expressing the constraints on arbitrarily powerful models by insufficient data sets. While the results of cosmological simulations are, of course, by no means garbage, they are also limited by the physics they implement and the chosen assumptions and simplifications. PINNs provide a way to go beyond these ingredients, forcing the model to consider further explicitly specified physical relationships.

Extensions of this exploratory study on the application of PINNs in this area could address complications from the black hole seeding process, including zero-mass black hole values in the data set. As

Table 4. Data set size and input variables for related works and this paper. Here, $S_{x,y,z}$ denotes the different components of the spin, V_c the maximum circular velocity in the subhalo, N_h the number of dark matter particles bound to the subhalo, λ_h the halo spin, and ρ_h the local halo density. Sets for additional information at higher redshifts are indicated using braces.

	Size	Input
Kamdar et al. (2016)	249 370	$M_h, S_x, S_y, S_z, \sigma_h, N_h, V_c$
Agarwal et al. (2018)	3400	$\{M_h\}_i, \{\rho_h\}_j, \lambda_h, \sigma_h$
Moews et al. (2021)	13 132	$\{M_h\}_i, r_h, \sigma_h$
Dai et al. (this work)	14 247	M_h, r_h, σ_h

this inclusion is likely to distort the predictive distributions due to a sharp cut-off, the question is whether models can be trained to sufficiently differentiate between these two regimes or an additional distinction step is required. A secondary discriminative model is a potential solution for this challenge, which would take the place of a filter that assigns feature vectors to two different models; one for zero-mass black hole predictions and one for the remaining majority of data points. While beyond the scope of this study, this could avoid smearing effects in the space between those two cases.

An additional extension to the presented research is the inclusion of other hydrodynamic simulations, as varying implementation details can lead to differences in the produced data sets. One example for this is the presence of quenching in SIMBA, which impacts the correlation between SFR and M_* . Depending on the goal of work in this area, for example, when focusing on the main sequence, other simulations can be considered to suit these needs.

Our additional implementation of distributional compliance as part of the loss function puts a further constraint on the model, which targets scatter fidelity. The direct incorporation of this additional knowledge into the architecture also goes beyond prior work on the inclusions of analytical models, which relies on physical computations outside of and before the application of machine learning models, as previously implemented by Moews et al. (2021).

At the same time, the predictive power of our architecture relies on the suitability of the included loss function components in terms of both the physical domain knowledge provided and the choice of a density approximation method. As covered in Section 3.1, the inclusion of baryon conversion efficiency through the relationship between M_* and M_h presents a simplification. This is especially the case due to stellar feedback resulting in lowered efficiency at lower masses and AGN feedback resulting in lowered efficiency at high masses. For future research, if this particular constraint is to be kept as part of the model, this should be further refined to more closely follow physical theory and represent a limitation.

The natural extension of our work, although beyond the scope of this paper, is the identification and analysis of additional domain knowledge to be injected into the loss function, thus allowing the model to root its learning in a more complete set of physics. Another pathway is replacing our normal assumption with more complex approximations that allow for non-Gaussian features and multimodal distributions.

In Table 4, we list the data set sizes and different dark matter properties used as inputs in similar works providing the same performance metrics (Kamdar et al. 2016; Agarwal et al. 2018; Moews et al. 2021). In contrast to this paper, all three comparable studies employ tree-based models, with the largest data set used by Kamdar et al. (2016). The latter also provide their model, in addition to M_h and σ_h , with the three spin components, number of dark matter particles bound to the subhalo, and maximum circular velocity in the

Table 5. Performance metrics for related works. The table lists the coefficient of determination (R^2) and Pearson’s correlation coefficient (ρ) for each output variable. Here, in addition to the variables described throughout this paper, M_{gas} denotes the total gas mass of a subhalo in the indicated paper.

Variable	Kamdar et al. (2016)		Agarwal et al. (2018)		Moews et al. (2021)	
	R^2	ρ	R^2	ρ	R^2	ρ
M_*	0.91	0.95	0.90	0.95	0.82	0.94
SFR	0.63	0.79	0.55	0.74	0.73	0.87
Z	0.93	0.96	0.73	0.86	0.21	0.66
M_{gas}	0.67	0.85	–	–	–	–
M_{HI}	–	–	0.35	0.59	0.36	0.65
M_{H_2}	–	–	0.51	0.71	0.54	0.75
M_{BH}	0.72	0.85	–	–	0.71	0.88

subhalo, but omit r_h as used in our work. This supplies additional information on the dark matter halo of their model.

Similarly, Agarwal et al. (2018) do not include r_h , but add the halo local density, with $\{M_h\}_i$ denoting not only the current halo mass at $z = 0$ but also the five preceding snapshots at higher redshifts, which provides information on the recent merger history. Conversely, $\{\rho\}_j$ indicates the set of nearby halo mass densities within radii $r \in \{200, 500, 1000\}$ (in kpc) focused on the halo’s mass centre, thus providing additional information in a similar way. That being said, their data set is the smallest in this line-up.

The work most closely aligned with ours in terms of inputs is Moews et al. (2021), although largest-progenitor merger trees over the entirety of a halo’s evolution are provided instead of only the current halo mass. More importantly, their hybrid approach uses an analytical formalism to pre-compute a subset of baryonic properties with a physical model, feeding the merger trees into both the latter and the subsequent machine learning model.

In Table 5, we list the performance metrics, R^2 and ρ , for these related works to enable a discussion of differences in data and model approach. With regard to Kamdar et al. (2016), our model yields a close performance on M_* and M_{BH} , but with a less robust prediction of SFR and Z . This could be due to the size of the ILLUSTRIS data set used by the authors, which is about eight times larger than our available SIMBA data, as more suitable data usually lead to a better performance in machine learning models (Zhou et al. 2014).

Agarwal et al. (2018) use MUFASA, the predecessor of SIMBA, and benefit from their inclusion of prior halo masses and halo mass densities at different distance radii. Our model’s results are comparable for M_* and M_{HI} , with a slight underperformance for Z , but there is a notable degradation in M_{H_2} and SFR predictions. Based on the correlation matrix in Fig. 4, these two properties are highly correlated, meaning that an improvement in one in future research should impact the other. One potential reason for this degradation is the larger scatter in SIMBA higher halo masses, while Agarwal et al. (2018) pre-select their galaxies to be star forming, which strongly reduces the scatter at high masses.

Lastly, Moews et al. (2021) develop a hybrid approach that combines an extra trees ensemble with the equilibrium model and incorporates merger trees into the latter. The same data and model inputs are used in both works, save for the largest-progenitor merger trees that are fed into both the physical and machine learning models. While our model achieves close results on M_* , Z , M_{HI} , and, to a lesser degree, M_{BH} , metrics of SFR and M_{H_2} are lower. Here, the difficulty in predicting the SFR due to the large scatter could be data-driven, while Moews et al. (2021) utilize a physical model for this property. At the same time, our model outperforms these results on Z , which reflects the equilibrium model’s difficulty with this quantity.

Overall, unlike our method focusing on the retrieval of accurate scatter, the above models are optimized for accuracy and make use of more information on the dark matter halo. The results of both Agarwal et al. (2018) and Moews et al. (2021) show a considerably less accurate scatter retrieval than can be seen in our experiments, with noticeable underprediction of the tails. The aim of this work is not to replace existing research but to demonstrate modifications that can increase the predictive ability of deep learning models in this area. Since the data set, model, and purpose are different, this comparison provides a direction for future research targeting the combination of these strengths. With regard to the model architecture, the closer approximation of scatter along each variable’s range instead of a general distribution parameter per variable could further improve the retrieval of scatter. While beyond the scope of this work, this should be combined with a closer look at alternative activation functions to potentially alleviate the wave-like nature of some predictions as described in Section 3.2.

For follow-up research not using PINNs, or machine learning models that do not require the loss function to be differentiable for backpropagation as described in Section 2.4, we recommend extensions to the comparison of target and prediction distributions. This work uses the KLD to calculate the difference between those distributions under the normal assumption. While a reasonably close proposition for the data used in our experiments, this limits our model’s effectiveness when directly transferred to markedly non-Gaussian data sets. In such cases, the spread of predictions through the variance would still be enforced, but in the case of, for example, a starkly multimodal distribution, the recreation of these distributional features would not be a major component of the optimization. This is also the reason for concerns covered in Section 2.4 with regard to potential negative effects on the mean accuracy, for example in case of strongly skewed distributions. Such negative effects are not present in our results, but could be argued to be a more relevant concern for starkly non-normal distributions in other research.

One point that should be discussed beforehand is the improvement in MSE, R^2 , and ρ values in Table 2. While no decrease in these key performance indices, as covered above, shows that the distributional constraint allows for retaining the level of accuracy achieved only through the inclusion of the SHMR, an improvement is not immediately intuitive, as the KLD component targets the scatter reproduction, not the predictive accuracy. The underprediction of tails is a common problem in the application of machine learning models to predict baryonic properties from dark matter halo information; this is the case not only in this research, but also in other publications, as mentioned above.

One possibility is that a more accurate distribution that places the data and predictions within the same range reduces, on average, the point-by-point error through a more accurate distribution width. Another reason could be found in the reduced underprediction of tails, which can be viewed as a subcase of the first possibility, as it results in a decreased potential for errors stemming from a lack of accurate predictions in these regions.

As the comparison needs to be reasonably fast, approaches such as Bayesian mixture models as well as associated methods that are more complex (see, for example, Moews & Zuntz 2020) are likely to slow down the training process too much. As a compromise, a KDE, also known as the Parzen–Rosenblatt window after Rosenblatt (1956) and Parzen (1962), can be used, although this is limited to lower dimensionalities. We propose two ways to circumvent the latter limitation. The first is applying 1D KDEs on a variable-by-variable basis and then averaging the KLDs between these estimates. While the advantage is the good fit in \mathbb{R}^1 , this can lead to a subset of variables

not being forced to follow the target distribution as long as the KLD average remains small.

The alternative is to make use of dimensionality reduction methods such as principal component analysis, which collapses the n -dimensional coordinate space, for n variables, into orthogonal vectors ranked by their ability to explain the variance (for a recent overview, see Jolliffe & Cadima 2016). Reducing the coordinate space to a subspace in \mathbb{R}^2 would, for example, still allow for reasonably good KDE approximations while retaining each variable’s contribution to a combined KLD. The same line of thought does, of course, also apply to other density approximation and dimensionality reduction methods.

5 CONCLUSION

In this paper, we transfer the paradigm of PINNs to predicting baryonic properties for associated dark matter halo variables. We adapt this approach in two different ways. The first includes the SHMR, a double power law previously used to parametrize the instantaneous baryon conversion efficiency. While this is a more established way to include physical theory into PINNs, our second extension is the enforcement of baryonic scatter in simulations under a normal assumption using the KLD between the underlying cosmological simulation and predictions. In doing so, and separately from improvements in the mean accuracy of predictions, we solve the common problem of scatter reproduction in this area, which is merged directly into the machine learning model.

We first test the improvement for the more traditional approach to PINNs, meaning the injection of physical domain knowledge into the loss function, and demonstrate a positive effect on the model’s performance. The hybrid approach combines these strengths by including the measurement of distributional differences and outperforms the standard PINN model. Subsequent tests of the scatter retrieval show more faithful reproductions for baryonic properties. These improvements in scatter are especially notable for molecular hydrogen masses and SFRs but can also be seen for neutral hydrogen masses and metallicities. In particular, our model successfully recovers the fundamental metallicity relation.

Our experiments demonstrate that PINNs, a rapidly expanding area of research across various subfields of physics, offer a way to directly bake theoretical constraints and distributional adherence into neural network architectures when painting baryonic properties into galactic dark matter haloes. As such, they can be used to complete cosmological N -body simulations based on full hydrodynamic simulation suites, although this comes with the same caveats as other research in this area.

The inference of physics from simulations operates under the assumption that such simulations are a sufficiently close approximation of the real world. Any machine learning models learning from those simulations are subject to the same assumption. That said, including physical models in the learning process enables these algorithms to include domain information beyond the underlying cosmological simulations.

Potential follow-ups include additional physical models specific to galaxy formation and evolution into the loss function, and further constraints based on observational data or other simulations to diversify the data sources. Our presented framework is widely applicable to large-scale cosmological simulations and the study of the utility and effect of physical domain knowledge on galaxy evolution emulators. It provides a further piece in the puzzle of fully using modern machine learning in astrophysics.

ACKNOWLEDGEMENTS

This work received support from the Hewlett Packard Enterprise Data Science Institute at the University of Houston. BM acknowledges support from a McWilliams Fellowship at Carnegie Mellon University. SIMBA was run on the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility. The equipment was funded by BEIS (Department for Business, Energy & Industrial Strategy) capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The SIMBA simulation data and galaxy catalogues underlying this article are publicly available at <https://simba.roe.ac.uk>.

REFERENCES

- Agarwal S., Davé R., Bassett B. A., 2018, *MNRAS*, 478, 3410
- Akiba T., Sano S., Yanase T., Ohta T., Koyama M., 2019, in A.Teredesai, ed., Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Optuna: A Next-Generation Hyperparameter Optimization Framework. Association for Computing Machinery, New York, USA, p. 2623
- Baugh C. M. et al., 2018, *MNRAS*, 483, 4922
- Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
- Ben-David A., Liu H., Jackson A. D., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 051
- Berlind A. A., Weinberg D. H., 2002, *ApJ*, 575, 587
- Blumenthal G., Faber S. M., Primack J. R., Rees M. J., 1984, *Nature*, 311, 517
- Bouché N. et al., 2010, *ApJ*, 718, 1001
- Boylan-Kolchin M., Springel V., White S. D. M., Jenkins A., Lemson G., 2009, *MNRAS*, 398, 1150
- Branca L., Pallottini A., 2022, *MNRAS*, 518, 5718
- Campbell D., van den Bosch F. C., Padmanabhan N., Mao Y.-Y., Zentner A. R., Lange J. U., Jiang F., Villarreal A. S., 2018, *MNRAS*, 477, 359
- Cao J., Tinker J. L., Mao Y., Wechsler R. H., 2020, *MNRAS*, 498, 5080
- Carleo G., Cirac I., Cranmer K., Daudet L., Schuld M., Tishby N., Vogt-Maranto L., Zdeborová L., 2019, *Rev. Mod. Phys.*, 91, 045002
- Cattaneo A. et al., 2017, *MNRAS*, 471, 1401
- Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, 319, 168
- Cornell A. S., Neube A., Harmsen G., 2022, *Phys. Rev. D*, 106, 124047
- Croton D. J. et al., 2016, *ApJS*, 222, 22
- Cui W. et al., 2018, *MNRAS*, 480, 2898
- Cuomo S., Di Cola V., Giampaolo F., Rozza G., Raissi M., Piccialli F., 2022, *J. Sci. Comput.*, 92, 88
- Cybenko G., 1989, *Math. Control Signals Syst.*, 2, 303
- Davé R., Finlator K., Oppenheimer B. D., 2012, *MNRAS*, 421, 98
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieeferantsa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- de Andres D., Yepes G., Sembolini F., Martínez-Muñoz G., Cui W., Robledo F., Chuang C., Rasia E., 2022, *MNRAS*, 518, 111
- de Santi N. S. M., Rodrigues N. V. N., Montero-Dorta A. D., Abramo L. R., Tucci B., Artale M. C., 2022, *MNRAS*, 514, 2463
- Desmond H., Mao Y., Wechsler R. H., Crain R. A., Schaye J., 2017, *MNRAS*, 471, L11
- Dissanayake M. W. M. G., Phan-Thien N., 1994, *Commun. Numer. Methods Eng.*, 10, 195
- Dolag K., Borgani S., Schindler S., Diaferio A., Bykov A. M., 2008, *Space Sci. Rev.*, 134, 229
- Dubois Y., Peirani S., Pichon C., Devriendt J., Gavazzi R., Welker C., Volonteri M., 2016, *MNRAS*, 463, 3948
- Efstathiou G., Davis M., White S. D. M., Frenk C. S., 1985, *ApJS*, 57, 241
- Ferdosi B. J., Buddelmeijer H., Trager S. C., Wilkinson M. H. F., Roerdink J. B. T. M., 2011, *A&A*, 531, A114
- Frenk C. S., White S. D. M., 2012, *Ann. Phys.*, 524, 507
- Fussell L., Moews B., 2019, *MNRAS*, 485, 3203
- Genel S. et al., 2014, *MNRAS*, 445, 175
- Hastie T., Tibshirani R., Friedman J., 2009, The Elements of Statistical Learning: Data Mining, Inference and Prediction, 2nd edn. Springer New York, NY
- Hatton S., Devriendt J. E. G., Ninin S., Bouchet F. R., Guiderdoni B., Vibert D., 2003, *MNRAS*, 343, 75
- Hee S., Vázquez J. A., Handley W. J., Hobson M. P., Lasenby A. N., 2016, *MNRAS*, 466, 369
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hornik K., Stinchcombe M., White H., 1989, *Neural Netw.*, 2, 359
- Jespersen C. K., Cranmer M., Melchior P., Ho S., Somerville R. S., Gabrielpillai A., 2022, *ApJ*, 941, 7
- Jin P., Zhang Z., Zhu A., Tang Y., Karniadakis G. E., 2020, *Neural Netw.*, 132, 166
- Jo Y., Kim J., 2019, *MNRAS*, 489, 3565
- Jolliffe I. T., Cadima J., 2016, *Phil. Trans. R. Soc. A*, 374, 20150202
- Kamdar H. M., Turk M. J., Brunner R. J., 2016, *MNRAS*, 457, 1162
- Karniadakis G. E., Kevrekidis I. G., Lu L., Perdikaris P., Wang S., Yang L., 2021, *Nat. Rev. Phys.*, 3, 422
- Kashefi A., Rempe D., Guibas L. J., 2021, *Phys. Fluids*, 33, 027104
- Klypin A. A., Trujillo-Gomez S., Primack J., 2011, *ApJ*, 740, 102
- Kravtsov A. V., Vikhlinin A. A., Meshcheryakov A. V., 2018, *Astron. Lett.*, 44, 8
- Krumholz M. R., Dekel A., 2012, *ApJ*, 753, 16
- Kullback S., Leibler R. A., 1951, *Ann. Math. Stat.*, 22, 79
- Lagaris I., Likas A., Fotiadis D., 1998, *IEEE Trans. Neural Netw.*, 9, 987
- LeCun Y., Bengio Y., 1995, Arbib M. A., ed., Handbook of Brain Theory and Neural Networks. MIT Press, Cambridge, MA, USA, p. 3361
- Lehmann B. V., Mao Y., Becker M. R., Skillman S. W., Wechsler R. H., 2016, *ApJ*, 834, 37
- Lovell C. C., Wilkins S. M., Thomas P. A., Schaller M., Baugh C. M., Fabbian G., Bahé Y., 2021, *MNRAS*, 509, 5046
- LSST Dark Energy Science Collaboration, 2020, *MNRAS*, 495, 5040
- Lu Y. et al., 2014, *ApJ*, 795, 123
- Maurov V., Pinkus A., 1999, *Neurocomputing*, 25, 81
- Martin J., Schaub H., 2022, *Celest. Mech. Dyn. Astron.*, 134, 46
- McCarn Deiana A. et al., 2022, *Frontiers Big Data*, 5, 787421
- McGibbon R. J., Khochfar S., 2022, *MNRAS*, 513, 5423
- Mishra S., Molinaro R., 2021, *J. Quant. Spectrosc. Radiat. Transfer*, 270, 107705
- Mitra S., Davé R., Finlator K., 2015, *MNRAS*, 452, 1184
- Mitra S., Davé R., Simha V., Finlator K., 2017, *MNRAS*, 464, 2766
- Moews B., Zuntz J., 2020, *ApJ*, 896, 98
- Moews B., de Souza R. S., Ishida E. E. O., Malz A. I., Heneka C., Vilalta R., Zuntz J., 2019, *Phys. Rev. D*, 99, 123529
- Moews B., Davé R., Mitra S., Hassan S., Cui W., 2021, *MNRAS*, 504, 4024
- Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macció A. V., Naab T., Oser L., 2010, *ApJ*, 710, 903
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Moster B. P., Naab T., White S. D. M., 2018, *MNRAS*, 477, 1822
- Moster B. P., Naab T., Lindström M., O’Leary J. A., 2021, *MNRAS*, 507, 2115
- Nicola A., Amara A., Refregier A., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 011
- Panther B., Jimenez R., Heavens A. F., Charlot S., 2007, *MNRAS*, 378, 1550
- Parzen E., 1962, *Ann. Math. Stat.*, 33, 1065
- Pillepich A. et al., 2018, *MNRAS*, 475, 648
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Potter D., Stadel J., Teyssier R., 2017, *Comput. Astrophys. Cosmol.*, 4, 2
- Probst P., Boulesteix A.-L., 2018, *J. Mach. Learn. Res.*, 18, 1
- Raissi M., Perdikaris P., Karniadakis G., 2019, *J. Comput. Phys.*, 378, 686

- Reddick R. M., Wechsler R. H., Tinker J. L., Behroozi P. S., 2013, *ApJ*, 771, 30
- Rees M. J., Ostriker J. P., 1977, *MNRAS*, 179, 541
- Riebe K. et al., 2013, *Astron. Nachr.*, 334, 691
- Rosenblatt M., 1956, *Ann. Math. Stat.*, 27, 832
- Saintonge A. et al., 2013, *ApJ*, 778, 2
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Shankar F. et al., 2017, *ApJ*, 840, 34
- Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
- Somerville R. S., Primack J. R., 1999, *MNRAS*, 310, 1087
- Springel V. et al., 2005, *Nature*, 435, 629
- Stiskalek R., Bartlett D. J., Desmond H., Anbajagane D., 2022, *MNRAS*, 514, 4026
- Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, *Nat. Rev. Phys.*, 2, 42
- von Martens R. et al., 2022, *MNRAS*, 516, 3924
- Wetzel A. R., White M., 2010, *MNRAS*, 403, 1072
- Wu Z., Pan S., Chen F., Long G., Zhang C., Yu P. S., 2021, *IEEE Trans. Neural Netw. Learn. Syst.*, 32, 4
- Yang Y., Perdikaris P., 2019, *Comput. Mech.*, 64, 417
- Yarotsky D., 2017, *Neural Netw.*, 94, 103
- Yasin T., Desmond H., Devriendt J., Slyz A., 2023, *MNRAS*, 525, 5066
- Zhou Z., Chawla N. V., Jin Y., Williams G. J., 2014, *IEEE Comput. Intell. Mag.*, 9, 62
- Zhou J. et al., 2020, *AI Open*, 1, 57

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