Boletim da Sociedade Astronômica Brasileira, **35**, no. 1, 206-210

Graph neural networks for robust parameter inference in cosmology: the first steps before real data

Natalí S. M. de Santi[1,2], Francisco Villaescusa-Navarro[2], & L. Raul Abramo[1]

[1] Instituto de Física, Universidade de São Paulo, R. do Matão 1371, 05508-900, São Paulo, Brasil
e-mail: natalidesanti@gmail.com
[2] Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY, 10010, USA

Abstract. Parameter inference in cosmology remains a significant challenge, lacking an optimal method that encompasses all cosmological scales. In an effort to provide an alternative approach, we are translating galaxy catalogs into graphs, effectively compressing information from small to large scales while considering rotational and translational symmetries, and training graph neural networks to do field-level likelihood-free inference. Our methodology enables precise and accurate predictions of the matter content of these catalogs. Notably, our approach exclusively utilizes the phase-space information of galaxies, marking the first time that we have a model which exhibits robustness across five distinct sub-grid-physical models (Astrid, IllustrisTNG, SIMBA, Magneticum, and SWIFT-EAGLE). Furthermore, our model surpasses challenges arising from differences in halo/subhalo finders and variations in cosmological and astrophysical parameters that deviate from those in the training set. We also consider the ability of the methodoly to deal with observational effects such as masking, peculiar velocity uncertainties, and color selection. Remarkably the suite works with these condiderations. By leveraging graph neural networks and focusing on galaxy phase-space data, our framework offers a promising avenue for addressing parameter inference in cosmology, showcasing its resilience to different sub-grid physical models, diverse astrophysical effects, and real observational challenges.

Resumo. Um dos grandes desafios da Cosmologia ainda é a estimativa de parâmetros cosmológicos, ainda sem um método que seja válido em todas as escalas. No presente trabalho nós propomos um modo para preencher tal lacuna, traduzindo catálogos de galáxias em grafos, o que, efetivamente, comprime a informação de pequenas e grandes escalas, afim de treinar redes neurais gráficas para fazer inferência em níveis de campo sem a necessidade de uma probabilidade de verossimilhança. Nossa metodologia proporciona predições acuradas e precisas do conteúdo de matéria desses catálogos. Notavelmente, nós apenas fazemos uso do espaço de fase das galáxias assegurando um modelo robusto para cinco modelos de sub-gride físicos diferentes (Astrid, IllustrisTNG, SIMBA, Magneticum e SWIFT-EAGLE). Além disso, nosso modelo é capaz de superar diferentes buscadores de halos e sub-halos e variações nos parâmetros astrofísicos diferentes dos considerados no conjunto de treinamento. Nós também consideramos a habilidade do modelo de lidar com efeitos observacionais de máscara, incertezas na velocidade peculiar e diferentes seleções de galáxias. O modelo é capaz de fazer boas predições inclusive com esses efeitos. Nosso método oferece um promissor meio de inferir parâmetros cosmológicos usando redes neurais gráficas utilizando apenas as posições e velocidades das galáxias, mostrando-se capaz de lidar com diferentes modelos de física de sub-gride, diversos efeitos astrofísicos e desafios observacionais.

Keywords. cosmological parameter inference – machine learning – galaxies

1. Introduction

Parameter inference is key to probing cosmology and testing our models. Great progress has come with the development of field-level likelihood-free inference (Lemos et al. (2023)), which do not require summary statistics and the resulting loss of information. In particular, there are the predictions coming directly from galaxy or halo catalogs, whose properties can be converted into graphs and used to feed graph neural networks (GNNs), as showed in de Santi et al. (2023); Shao et al. (2023); Villanueva-Domingo at al. (2022); Shao et al. (2022). In the context of cosmology, this kind of analysis does not impose any cut on scale and can easily incorporate different physical symmetries, as well as already being permutationally invariant.

In this work we make use of the GNNs to developed a ML suite able to infer the value of $\Omega_{\rm m}$ directly from galaxy catalogs, using only positions and the z component velocity. The model aims to be robust (without needing any additional parameters) to cosmology, astrophysics, and subgrid models, as it was tested using thousands of galaxy catalogs produced from hydrodynamic simulations run with 5 different codes. To tackle real data we are taking into account several observational and systematic effects in these catalogs, analysing the ability of the method to consider them.

The manuscript is organized as follows: in Section 2, we present the data set we use to train and test the models and how we implement the systematic effects. In Section 3, we discuss the way we construct graphs from the galaxy catalogs and how we process them through GNNs. In Section 4, we present the main results and, in Section 5, we discuss and summarize the results.

2. Data

2.1. Simulations

The galaxy catalogs we are working with were constructed from the hydrodynamic simulations of the Cosmology and Astrophysics with MachinE Learning Simulations — CAMELS project (Villaescusa-Navarro et al. (2023)). These simulations are done in periodic boxes of 25 h^{-1}Mpc, and we focus our attention into the $z = 0$ results.

The CAMELS simulations can be classified into different sets and suites. In this work, we use 2 different sets and 5 suites.

The main characteristics of the different sets are:

- *Latin Hypercube (LH).* Have been run with different random seeds and have the value of their cosmological and astrophysical parameters arranged in a LH for $\Omega_{\rm m}$, σ_8, $A_{\rm SN1}$, $A_{\rm SN2}$, $A_{\rm AGN1}$, and $A_{\rm AGN2}$. These simulations have been used for training, validating, and testing the GNNs.
- *Sobol Sequence (SB).* The catalogs in this set have their cosmological and astrophysical parameters being varied following a Sobol sequence. A total of 28 parameters are varied: 5 cosmological and 23 astrophysical (Ni et al. (2023)). We use these simulations for testing the models.

The CAMELS simulations can also be classified into suites, according to the code used to run them:

- *Astrid.* The code used was MP-Gadget, applying some modifications to the subgrid model employed in the Astrid simulation (Ni et al. (2023)). This set contains 1000 LH simulations.
- *SIMBA.* The code used was Gizmo, employing the same subgrid physics as the SIMBA simulation (Davé et al. (2019)). This subset contains 1000 LH simulations.
- *IllustrisTNG.* The code used was Arepo, applying the same subgrid physics as the IllustrisTNG simulations (Weinberger et al. (2020)). This suite contains 1000 LH and 2048 SB simulations.
- *Magneticum.* The code used was the parallel cosmological Tree-PM code P-Gadget3, employing some modifications on it (de Santi et al. (2023)). This set contains 50 LH simulations.
- *SWIFT-EAGLE.* The code used was Swift (Schaller et al. (2023)), using a new subgrid physics model based on the original Gadget-EAGLE simulations (Schaye et al. 2015), with some changes. This subset contains 64 LH simulations.

Astrid, SIMBA, IllustrisTNG, and Magneticum have their halos and subhalos found using SubFind (Dolag et al. (2009)). In the case of SWIFT-EAGLE it was used VELOCIraptor (Elahi et al. (2019)). We used this difference to construct a model that remains robust across various methods for identifying halos and subhalos.

2.2. Observational effects

We now describe the different observational/systematic effects that we consider, and how we simulate them.

- *Masking.* In real surveys, some fraction of the galaxies are masked out due to bright stars, cosmic rays, bad pixels, etc. Here we simulate the effect of masks by randomly removing 5% of the galaxies in the catalogs.
- *Peculiar velocity uncertainties.* Peculiar velocity cannot be precisely measured, since observations are unable to distinguish between radial (line-of-sight) positions and radial velocities (Howlett et al. (2017)). We simulate this effect by adding a random error to the line-of-sight peculiar velocity of each galaxy according to: $xsv_z \rightarrow v_z\,[1 + 0.15\mathcal{N}(0, 1)]$, representing a relative error on the peculiar velocities of 15%.
- *Galaxy selection.* In real surveys, galaxies are selected according to some criteria: e.g., their color. Here we employ a color selection based on the "quenched" and "not-quenched" galaxies, using the specific star formation rate (sSFR = SFR/$\rm M_\star$ [$\rm yr^{-1}M_\odot$], where a galaxy's SFR is defined as the sum of the individual SFR of all gas cells in its subhalo), according to Davé et al. (2019), i.e.:
  - *Blue:* sSFR < -10.8 $\rm yr^{-1}\,M_\odot$,
  - *Red:* sSFR > -10.8 $\rm yr^{-1}\,M_\odot$.

In all these scenarios we are modifying the CAMELS catalogs (train, validation and test) in order to include these effects.

3. Methodology

This section is devoted to describing: (1) how we built the graphs from the galaxy catalogs (Section 3.1); (2) the details behind the architecture of the GNNs (Section 3.2); (3) the use of the moment neural networks (MNN) to do the likelihood-free inference (Section 3.3); (4) the training procedure and optimization choices (Section 3.4); and (5) the evaluation of the models, where we present the metrics we analyze (Section 3.5).

3.1. Galaxy graphs

Graphs are mathematical units composed by nodes ($\mathbf{n}_i$), edges ($\mathbf{e}_{ij}$, connecting a node i to a node j), and global properties ($\mathbf{g}$), each one of them characterized by a set of properties, usually called as attributes (Zhou et al. (2018)). The graphs are the input for the GNNs, not the galaxies per se. Then, we built the graphs from galaxy catalogs, using the galaxy positions (to find the edges and the edge properties); their peculiar velocities (only the z component), as node attributes; and the logarithm of the number of galaxies in the graph: $\log_{10}(N_g)$, as global attribute. This was done as presented in de Santi et al. (2023), where we link galaxies if they are close enough to be inside a radius centered on each of them, r_{link}. The value of the linking radius was a tuned hyperparameter, see Section 3.4.

The edge attributes are related to the spatial distribution of galaxies. We choose these features to make the graph invariant under rotations and translations, in the way that the edge features are:

$$\mathbf{e}_{ij} = \left[\frac{|\mathbf{d}_{ij}|}{r_{link}}, \alpha_{ij}, \beta_{ij}\right], \tag{1}$$

where: $\mathbf{d}_{ij} = \mathbf{r}_i - \mathbf{r}_j$, $\boldsymbol{\delta}_i = \mathbf{r}_i - \mathbf{c}$, $\alpha_{ij} = \frac{\boldsymbol{\delta}_i}{|\boldsymbol{\delta}_i|} \cdot \frac{\boldsymbol{\delta}_j}{|\boldsymbol{\delta}_j|}$, $\beta_{ij} = \frac{\boldsymbol{\delta}_i}{|\boldsymbol{\delta}_i|} \cdot \frac{\mathbf{d}_{ij}}{|\mathbf{d}_{ij}|}$. here we denote by $\mathbf{r}_i$ the position of a galaxy i and $\mathbf{c} = \sum_i^N \mathbf{r}_i/N$ the catalog centroid. The *distance* $\mathbf{d}_{ij}$ is the difference of 2 galaxy (i and j) positions, the *difference vector* $\boldsymbol{\delta}_i$ denotes the position of a galaxy i with respect to the centroid, α_{ij} is the (cosine of) the angle between the difference vectors of 2 galaxies, while β_{ij} represents the angle between the difference vector of a galaxy i and its distance to another galaxy j. We account for periodic boundary conditions (PBC) when computing distances, angles, and reverse edges (de Santi et al. 2023; Villanueva-Domingo at al. 2022).

3.2. GNNs architecture

The architecture we employ here is the *message passing scheme* (de Santi et al. (2023); Villanueva-Domingo at al. (2022)). Primarily, the GNNs are associated with a MNN and trained to infer the value of $\Omega_{\rm m}$ for any input graph. This task is done by transforming the graph attributes (node $\mathbf{n}_i$ and edge $\mathbf{e}_{ij}$ properties are updated using the information contained in the neighboring nodes $\mathbf{n}_i$, edges $\mathbf{e}_{ij}$, and global $\mathbf{g}$ characteristics), while the graph structure (edge indexes) is preserved. More especifically, the node and edge features at layer $\ell + 1$ are updated as:

- *Edge model:*

$$\mathbf{e}_{ij}^{(\ell+1)} = \mathcal{E}^{(\ell+1)}\left(\left[\mathbf{n}_i^{(\ell)}, \mathbf{n}_j^{(\ell)}, \mathbf{e}_{ij}^{(\ell)}\right]\right), \tag{2}$$

being $\mathcal{E}^{(\ell+1)}$ a MLP;

– *Node model:*

$$\mathbf{n}_i^{(\ell+1)} = \mathcal{N}^{(\ell+1)}\left(\left[\mathbf{n}_i^{(\ell)}, \bigoplus_{j\in\mathfrak{N}_i} \mathbf{e}_{ij}^{(\ell+1)}, \mathbf{g}\right]\right), \tag{3}$$

with $\mathfrak{N}_i$ representing all neighbors of node i, $\mathcal{N}^{(\ell+1)}$ a MLP, and $\oplus$ a multi-pooling operation responsible to concatenate several permutation invariant operations:

$$\bigoplus_{j\in\mathfrak{N}_i} \mathbf{e}_{ij}^{(\ell+1)} = \left[\max_{j\in\mathfrak{N}_i} \mathbf{e}_{ij}^{(\ell+1)}, \sum_{j\in\mathfrak{N}_i} \mathbf{e}_{ij}^{(\ell+1)}, \frac{\sum_{j\in\mathfrak{N}_i} \mathbf{e}_{ij}^{(\ell+1)}}{\sum_{j\in\mathfrak{N}_i}}\right]. \tag{4}$$

The number of layers to perform this update is a hyperparameter to be chosen in the optimization scheme.

Then, the updated version of the initial graph, after N message passing layers, is collapsed it into a 1D feature vector according to

$$\mathbf{y} = \mathcal{F}\left(\left[\bigoplus_{i\in\mathfrak{F}} \mathbf{n}_i^N, \mathbf{g}\right]\right), \tag{5}$$

where $\mathcal{F}$ is the last MLP, $\oplus_{i\in\mathfrak{F}}$ the last multi-pooling operation (operating over all nodes in the graph $\mathfrak{F}$), and $\mathbf{y}$ the target of the GNNs (i.e. $\Omega_{\rm m}$).

3.3. Likelihood-free inference and the loss function

The final product of the ML models presented is the inference of $\Omega_{\rm m}^i$, by predicting the marginal posterior mean μ_i and standard deviation σ_i. These predictions are done without making any assumption about the form of the posterior, using a specific loss function according to Jeffrey & Wandelt (2020):

$$\mathcal{L} = \log\left[\sum_{i\in\text{batch}} \left(\Omega_{\rm m}^i - \mu_i\right)^2\right] + \log\left\{\sum_{i\in\text{batch}} \left[\left(\Omega_{\rm m}^i - \mu_i\right)^2 - \sigma_i^2\right]^2\right\}, \tag{6}$$

where i represents the samples in a given batch.

3.4. Training procedure and optimization

de Santi et al. (2023) and Ni et al. (2023) found that the ML models trained on Astrid were precise and robust because the Astrid simulations encompass a diverse range in some galaxy properties and galaxy number density. For this reason, all the models presented in this paper were trained on graphs from the LH set of the Astrid simulations.

The division of the 1000 LH simulations was: 850 for training, 100 boxes for validation, and 50 boxes for testing (all of these subsets built considering the observational effects). The number of epochs were 300, we have used the Adam optimizer (Kingma & Ba (2014)), and a batch size of 25. The optimization of the hyperparameters (learning rate, weight decay, linking radius, number of message passing layers, and number of hidden channels per layer of the MLPs) was done using Optuna (Akiba et al. (2019)). We used at least 100 trials to sample the hyperparameter space, minimizing the validation loss, computed using an early-stopping scheme. Then, we save only the model with the minimum validation error and use it for robustness tests. Note that the value for the linking radius was found, most of the time, equal to $r_{link} \sim 1.25\ h^{-1}$ Mpc, as in the previous work de Santi et al. (2023).

3.5. Performance Metrics

The scores used to validate each model were:

– *Mean relative error:*

$$\epsilon = \frac{1}{N}\sum_{i=1}^{N} \frac{|\Omega_{\rm m}^i - \mu_i|}{\mu_i}. \tag{7}$$

Low values indicate the model is precise.

– *Reduced chi squared:*

$$\chi^2 = \frac{1}{N}\sum_{i=1}^{N} \left(\frac{\Omega_{\rm m}^i - \mu_i}{\sigma_i}\right)^2. \tag{8}$$

This statistic quantifies the accuracy of the estimated errors. Values of χ^2 close to 1 indicate that the errors are accurately predicted, while values larger/smaller than 1 indicate the model is under/over predicting the errors.

4. Results

In this section we present the results of training and testing our models without and with the different considered systematic effects. The scores are taken on galaxy catalogs with different cosmologies, astrophysical parameters, and subgrid physic models, in order to test the coverage of the predictions and the robustness of the different ML models.

In all the analyses, we are presenting the metrics measured in the complete data set but we are presenting a unique point considering the average over the predictions in the whole set to allow a compression of the results in the paper. In some cases, our results are highly affected by a few outliers. To avoid contamination from them, we report the scores without them (for results with $\chi^2 < 10$).

4.1. Considering all the galaxies

We start by showing the results of training the GNNs on catalogs considering all the galaxies and their 3D positions and 1D velocities. This is presented in the first plot of Figure 1. Overall, all the model trained on Astrid and tested on Astrid achieves good results, specially after removing only one point, getting $\epsilon \sim 11.8\%$ and $\chi^2 \sim 1.6$. This model was tested on all the other sub-grid physical models, presenting good scores for all of them: $\epsilon \leq 12\%$ and $\chi^2 \lesssim 1.7$. For the first time we have achieved a robust model accross 5 different hydrodynamical models (Astrid, SIMBA, IllustrisTNG, Magneticum, and SWIFT-EAGLE), different halo-subhalo finders (SubFind and VELOCIraptor), and parameter variations different from the training set (which is the case of SB28). A deeper discussion of these findings is done in de Santi et al. (2023).

4.2. Masking

The first real observational effect considered was masking. We have used a random mask applied to the galaxy catalogs that randomly removed 5% of them. This analysis considers the power of the GNNs to constrain $\Omega_{\rm m}$ from catalogs of reduced information (less galaxies).

The results for the LH set are presented in the second plot of Figure 1. In a whole the network is behaving robustly across nearly all hydrodynamic suites and the outliers are representing only $\in [3, 7]\%$ of the whole data set. The only exception is Magneticum, for which the predictions are poor ($\chi^2 \sim 2.25$).

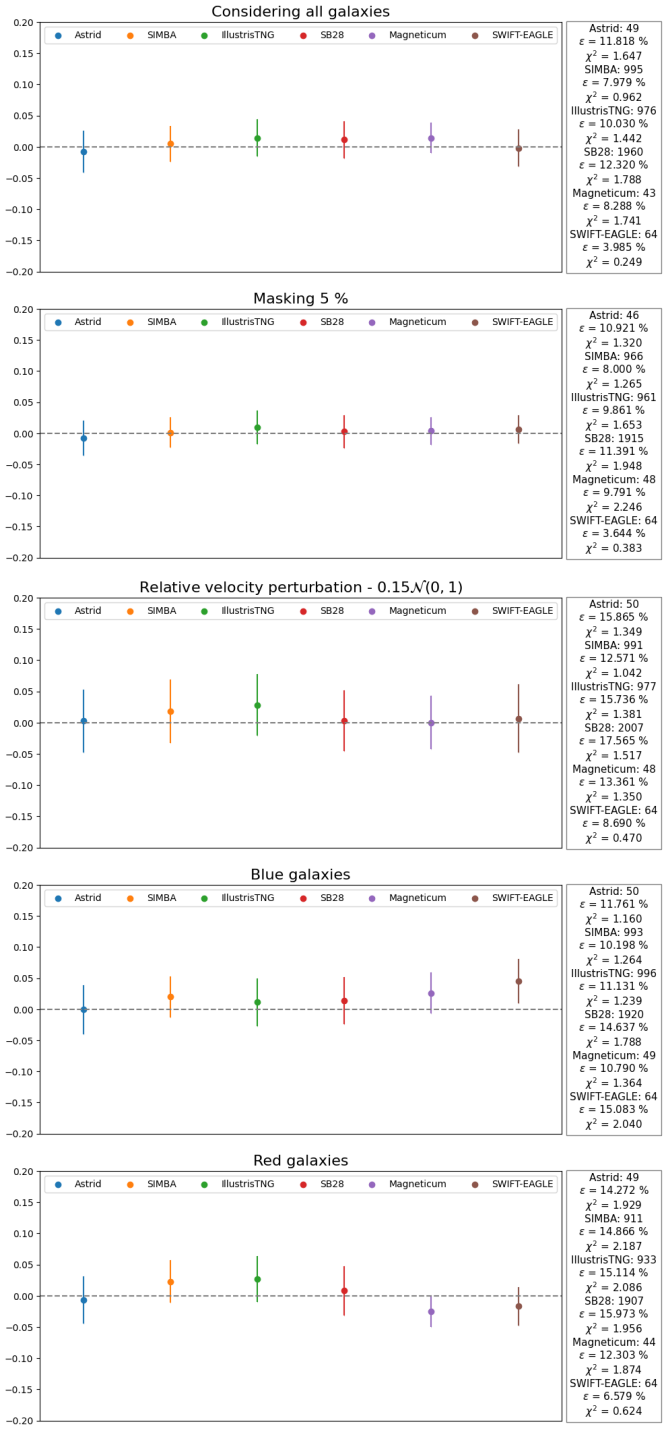

Figure 1. Predictions for: (1) all the galaxies in all the catalogs; (2) masking 5 % of the galaxies in all the catalogs; (3) considering a relative velocity perturbation of 15%; (4) considering blue galaxies; (5) considering red galaxies. We present the Truth - Inference of Ω_m for galaxy catalogs from Astrid, SIMBA, IllustrisTNG, SB28, and Magneticum alongside the scores for the predictions selected by their χ^2 values (indicating the number of catalogs into account by the name of the set in question).

4.3. Peculiar velocity uncertainties

We now discuss the results related to errors in the galaxy's peculiar velocities, that can be seen in the third plot of Figure 1. It is noticeble the fact that the error bars incread in all the tests (what can be seen by the higher values of $\epsilon \in [8.7, 17.6]\%$), which demonstrates the difficulty of the predictions by the GNN due to the changes in the galaxy velocities. Even though, the model presents very robust: the disparity in the predictions for Magneticum is comparable to the other sub-grid physical models and the fraction of outliers were always less than 4%.

4.4. Galaxy selection: color

In the fourth and fifth plots of Figure 1 we show the impact of selecting galaxies based on their color. We find that the GNNs perform well for both types, though the model is slightly better and more robust for the blue galaxies. Even after discarding the outliers, the predictions for the red galaxies are not as good as for the blue galaxies: $\epsilon \in [6.6, 15.9]\%$ (for red galaxies) $\times$, $\epsilon \in [10.2, 15.0]\%$ (for blue galaxies), and $\chi^2 \in [0.62, 2.19]$ (for red galaxies) $\times$ $\chi^2 \in [1.16, 2.04]$ (for blue galaxies). Also, the percentage of outliers is always large for the predictions in the red catalogs (6% $\rightarrow$ 12%, from blue to red galaxies). A reason for this behavior may be the low number density and higher clustering associated to the red galaxies.

5. Discussion and conclusions

Various approaches have been employed in attempting to constrain cosmological parameters through surveys of galaxy redshifts (de Santi & Abramo (2022); Chartier & Wandelt (2022)). In particular, the longstanding and successful history of utilizing galaxy positions and peculiar velocities to estimate parameters like Ω_m (Howlett et al. (2017); Peebles (1980)) has prompted significant investment in peculiar velocity surveys (Howlett et al. (2017)). However, traditional methods have been the primary means of analyzing these datasets until now (Lai et al. (2023); Howlett et al. (2017)). The emergence of ML techniques holds promise as a toolkit capable of addressing the challenge of deducing cosmological parameters from large-scale structures (Perez et al. (2023)). Notably, the field has undergone a revolution with the recent incorporation of GNNs (de Santi et al. (2023); Villanueva-Domingo at al. (2022); Shao et al. (2022)).

GNNs have already been shown to lead to fairly accurate predictions for Ω_m while using phase-space information together with other properties, for both galaxies (Villanueva-Domingo at al. (2022)) and halos (Shao et al. (2022)). An important caveat related to these works was that, when using galaxies the model was not robust across different subgrid physical models, whereas when halos were employed, the methods had to rely on unobservable halo properties. However, de Santi et al. (2023) created a GNN that was designed to take only the galaxy positions and velocities, obtaining a model able to extrapolate the predictions even when changes are made to the astrophysics modelling, subgrid physics, and subhalo/galaxy finder, which was reproduced here. The next step now is to consider real-world and systematic effects in these galaxy catalogs, in order to show that GNNs are viable tools that can handle observational data, effects already here tested.

In this work we have trained different GNNs to be tested on galaxy catalogs from CAMELS in order to infer Ω_m at the field level, using a likelihood-free approach and considering different observational effects, such as masking, velocity errors, and galaxies selected by their color. In our analyses we sought models that are capable of making robust and accurate predictions. The main conclusions of those tests are:

– All the models presented here were trained on Astrid catalogs, using only galaxy positions and velocities. When tested on Astrid catalogs, they were able to predict Ω_m with scores

(mean relative error and reduced chi equared — see equations 7 and 8) up to $\epsilon \sim 16\%$, and $\chi^2 \sim 1.9$, even when including systematic effects, as well as different cosmologies and astrophysical parameters.
- The robustness of those models were put to the test on SIMBA, IllustrisTNG, SB28, Magneticum, and SWIFT-EAGLE, showing that, indeed, we still achieved models that are able to extrapolate their predictions despite the inclusion of systematics and different subgrid physical methods. These robust models achieved scores up to $\epsilon \sim 17\%$ and $\chi^2 \sim 2.2$.
- For each of the tests we have presented the scores corresponding to selecting the predictions with $\chi^2 < 10$ (i.e., removing the outliers). The typical percentage of outliers correspond to $\sim 10\%$ of the samples.
- Our models can handle all the observational effects considered in the present work (masking effects, peculiar velocity errors, and galaxy selection effects), with different impacts on the accuracies of the predictions. The effect with the least impact on the results, compared to the model considering all the galaxies in the catalogs (see first plot of Figure 1) and de Santi et al. (2023) was masking 5% of the galaxies: $\epsilon = 10.92\%$ and $\chi^2 = 1.32$, for testing on Astrid and removing only $\in [3, 7]\%$ of outliers in the robustness tests (i.e., while tested on SIMBA, IllustrisTNG, SB28, Magneticum, and SWIFT-EAGLE). On the other hand, the selection of red gaalxies led to the worst scores.

In conclusion, we have shown that the method proposed by de Santi et al. (2023) to recover cosmological parameters from galaxies, and further developed here, is relatively robust to observational effects. For some of those real-world effects the results are still very good, while for others there is a larger impact on the accuracy of the recovered parameters. We believe that further improvements can be made — by, e.g., training on an even wider parameter space that includes not only cosmology and astrophysics but also systematic effects. Moreover, we can also design models which are more accurate within a given range of scales and with specific selection criteria. This paper represents an important first step towards applying these methods to real galaxy catalogs.

Acknowledgements. We would like to thank the CAMELS team for the enlightening discussions and valuable comments. We thank the São Paulo Research Foundation (FAPESP), the Brazilian National Council for Scientific and Technological Development (CNPq), and the Simons Foundation for financial support. NSMS acknowledges financial support from FAPESP, grants 2019/13108-0 and 2022/03589-4. The training of the GNNs has been carried out using graphic processing units from Simons Foundation, Flatiron Institute, Center of Computational Astrophysics.

References

Lemos, P., Coogan, A., Hezaveh, Y., et al. 2023, 40th International Conference on Machine Learning, 202, 19256
de Santi, N. S. M., Shao, H., Villaescusa-Navarro, F., et al. 2023, ApJ, 952, 69
Shao, H., de Santi, N. S. M., Villaescusa-Navarro, F., et al. 2023, ApJ, 956, 149
Villanueva-Domingo, P. & Villaescusa-Navarro, F. 2022, ApJ, 937, 115
Shao, H., Villaescusa-Navarro, F., Villanueva-Domingo, P., et al. 2022, arXiv:2209.06843
Villaescusa-Navarro, F., Genel, S., Anglés-Alcázar, D., et al. 2023, ApJS, 265, 54
Ni, Y., Genel, S., Anglés-Alcázar, D., et al. 2023, arXiv:2304.02096
Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, MNRAS, 486, 2827
Weinberger, R., Springel, V., & Pakmor, R. 2020, ApJS, 248, 32
Schaller, M., Borrow, J., Draper, P. W., et al. 2023, arXiv:2305.13380
Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, MNRAS, 446, 521
Dolag, K., Borgani, S., Murante, G. and et al. 2009, MNRAS, 399, 2
Elahi, P. J., Cañas, R., Poulton, R. J. J., et al. 2019, PASA, 36, e021
Howlett, C., Staveley-Smith, L., & Blake, C. 2017, MNRAS, 464, 2517
Zhou, J., Cui, G., Hu, S., et al. 2018, arXiv:1812.08434
Jeffrey, N. & Wandelt, B. D. 2020, arXiv:2011.05991
Bronstein, M. M., Bruna, J., Cohen, T., et al. 2021, arXiv:2104.13478
Kingma, D. P. & Ba, J. 2014, arXiv:1412.6980
Akiba, T., Sano, S., Yanase, T., et al. 2019, arXiv:1907.10902
de Santi, N. S. M. & Abramo, L. R. 2022, KCAP, 2022, 013
Chartier, N. & Wandelt, B. D. 2022, MNRAS, 509, 2220
Peebles, P. J. E. 1980, Large-Scale Structure of the Universe by Phillip James Edwin Peebles. Princeton University Press, 1980. ISBN: 978-0-691-08240-0
Lai, Y., Howlett, C., & Davis, T. M. 2023, MNRAS, 518, 1840
Perez, L. A., Genel, S., Villaescusa-Navarro, F., et al. 2023, ApJ, 954, 11