
Estimating Dark Matter Halo Masses in Simulated Galaxy Clusters with Graph Neural Networks

Nikhil Garuda
Steward Observatory
University of Arizona
933 N Cherry Ave, Tucson, AZ, 85721
nikhilgaruda@arizona.edu

John F. Wu
Space Telescope Science Institute
3700 San Martin Dr, Baltimore, MD 21218
jowu@stsci.edu

Dylan Nelson
Universität Heidelberg
Zentrum für Astronomie, ITA,
Albert-Ueberle-Str. 2
69120 Heidelberg, Germany
dnelson@uni-heidelberg.de

Annalisa Pillepich
Max-Planck-Institut für Astronomie
Königstuhl 17, 69117 Heidelberg, Germany
pillepich@mpia.de

Abstract

Galaxies grow and evolve in dark matter halos. Because dark matter is not visible, galaxies’ halo masses (M_{halo}) must be inferred indirectly. We present a graph neural network (GNN) model for predicting M_{halo} from stellar mass (M_*) in simulated galaxy clusters using data from the IllustrisTNG simulation suite. Unlike traditional machine learning models like random forests, our GNN captures the information-rich substructure of galaxy clusters by using spatial and kinematic relationships between galaxy neighbour. A GNN model trained on the TNG-Cluster dataset and independently tested on the TNG300 simulation achieves superior predictive performance compared to other baseline models we tested. Future work will extend this approach to different simulations and real observational datasets to further validate the GNN models ability to generalise.

1 Introduction

In the Lambda Cold Dark Matter cosmological model [28, 4], galaxies form and evolve in dark matter halos. Cosmological simulations demonstrate that galaxies grow in tandem with their dark matter halos according to well-measured and tight scaling relations [39]. This interdependence between stellar mass (M_*) and subhalo mass (M_{halo}) is known as the stellarhalo mass relation (SHMR).

While M_* is observable, M_{halo} must often be inferred indirectly via the SHMR due to observational constraints. For example, galaxy clusters—the most massive gravitationally bound objects in the Universe—are dark matter dominated, but their total mass must be measured via gravitational lensing [8, 37], the Sunyaev-Zel’dovich effect [2, 22, 3], and/or visible wavelength proxies (e.g., galaxy richness, intracluster light, etc; [30, 31]). However, these methods are unable to fully leverage galaxy substructure within clusters to estimate their dark matter halo masses.

Therefore, we present a graph neural network (GNN) algorithm [32] for predicting M_{halo} for galaxies in simulated cluster environments¹. Compared to primitive machine learning (ML) methods like random forests [1], a GNN can learn the substructure in neighbouring galaxies and thereby improve halo mass predictions. Our results using the GNN demonstrate significant performance gains on the training, validation, and an independent test set.

¹https://github.com/Nikhil0504/halo_masses




Figure 1: Flow diagram of GNN architecture used for halo mass prediction. The GNN processes node features (x_i, x_j) and edge features (ϵ_{ij}) through multiple unshared layers, where each layer applies learnable functions, ϕ , which are implemented as MLPs. These unshared layers operate in parallel across the graph structure. A pooling layer then aggregates (\oplus) the information from these interactions back into each node. Subsequent repetitions of these GNN layers can give it more representational power. Finally, the output MLP, ψ , combines node features and aggregated edge features to predict each node’s halo mass.

2 IllustrisTNG Simulation Data

The simulation data we use are large-volume, cosmological, gravo-magnetohydrodynamical simulations from the IllustrisTNG simulation suite [25]. We specifically use the TNG-Cluster [24] simulation, a collection of zoom-in simulations centered 352 of the most massive halos (i.e., galaxy clusters), for training and validation. Our dataset is based on the SUBFIND [35] subhalo catalogs that were obtained from snapshot 99 ($z = 0$), focusing on the high-resolution components of the zoom-in simulation. We adopt cosmological parameters from [29], using $H_0 = 67.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$ for consistency with the IllustrisTNG simulation suite. Additional details about the TNG-Cluster training data are provided in Appendix A. The distribution of subhalos in TNG-Cluster is shown in Figure 4, and the selection criteria and number of samples are described in Table 2.

We test our ML models on an independent data set, the Illustris TNG300-1 hydrodynamic simulation (hereafter TNG300; [25]). The TNG-Cluster and TNG300 simulations use the same physics and have comparable resolutions (in the former’s zoom-in regions), but the two simulations are otherwise independent. When reporting TNG-Cluster cross-validation results TNG300 test set results, we *only* consider galaxies within 10 Mpc of all clusters with $M_{\text{halo}} > 10^{14} M_{\odot}$.

3 Methods/Experiments

The primary objective of our study is to estimate M_{halo} from M_* . Building on the work of [18], we train ML models on galaxies and dark matter halos from TNG-Cluster to probe cluster environments.


Loss Functions and Evaluation Metrics. Model performance is assessed using several metrics (presented in Table 1). Simple models are trained to minimise the Mean Squared Error (MSE), while the GNN is optimised using Gaussian negative log-likelihood (combining MSE and log-variance terms, per [15])². Validation and test performance are evaluated with Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), coefficient of determination R^2 , Normalised Median Absolute Deviation (NMAD), average offset (Bias), and Outlier Fraction (f_{outlier}).

Random Forest Baseline Models. To establish a benchmark for subsequent comparisons with our GNN model, we use Random Forest (RF) regression [14] as a baseline model due to its capability to handle complex non-linear relationships between features. To further augment the simple RF model, we compute an overdensity parameter (Δ_G), defined as the sum of stellar masses within a specified radius R_{max} . The RF models are configured with 100 estimators using `scikit-learn` [27], one of which utilises M_* , and one which uses both M_* and Δ_G as features.

Graph Neural Networks. In our GNN model, each node represents a galaxy, with the M_* as the sole node feature. We construct edges between galaxy pairs separated by less than 3 Mpc [41], connecting neighbouring nodes. These connections enable the neural network to learn interactions between the substructure and galaxy properties within the cluster. We provide two edge features to incorporate both the spatial and kinematic separations of galaxies: the squared Euclidean distance between pairs of galaxy positions, and pairs of relative line-of-sight velocities.

GNN Architecture. Our GNN follows the architecture described in [41] with 8 unshared layers and 3 sequential layers, as shown in Figure 1. Each layer is composed of a two-layer MLP with

²The negative log-likelihood objective accounts for the intrinsically varying scatter in M_{halo} .



(a) Predicted versus true M_{halo} for the TNG-Cluster validation set, coloured by distance from cluster center. (b) Predicted versus true M_{halo} for the TNG300 test set, coloured by distance from cluster center.

Figure 2: For each subfigure, we show results for the RF with only M_* as a feature (*left*), the RF with M_* and overdensity parameter Δ_G (*center*), and the GNN with M_* and graph connectivity (*right*).

Table 1: Validation and test set performance for all models. The best metrics are underlined.

Model	RMSE	MAE	R^2	Bias	f_{outlier}	NMAD
TNG-Cluster cross-validation						
(Always predict mean)	0.542	0.396	0	0	0.019	0.479
RF: M_*	0.489 ± 0.002	0.382 ± 0.003	0.186 ± 0.011	-0.067 ± 0.006	0.008 ± 0.001	0.463 ± 0.005
RF: $M_* + \Delta_G$	0.385 ± 0.002	0.301 ± 0.002	0.490 ± 0.007	-0.124 ± 0.004	0.008 ± 0.000	0.367 ± 0.002
GNN	<u>0.273 ± 0.010</u>	<u>0.209 ± 0.009</u>	<u>0.745 ± 0.019</u>	-0.085 ± 0.027	0.013 ± 0.002	<u>0.246 ± 0.013</u>
TNG-300 test set						
(Always predict mean)	0.466	0.351	0	0	0.021	0.422
RF	0.468 ± 0.009	0.365 ± 0.014	0.199 ± 0.030	-0.200 ± 0.017	0.009 ± 0.003	0.456 ± 0.033
RF: $M_* + \Delta_G$	0.344 ± 0.003	0.256 ± 0.003	0.567 ± 0.007	-0.048 ± 0.001	0.022 ± 0.001	0.293 ± 0.006
GNN	<u>0.242 ± 0.013</u>	<u>0.184 ± 0.010</u>	<u>0.785 ± 0.023</u>	-0.039 ± 0.034	0.014 ± 0.002	<u>0.217 ± 0.014</u>

Note: The intrinsic scatter in M_{halo} ranges from 0.42 (at $\log(M_{\text{halo}}) = 11 M_{\odot}$) to 0.33 (at $\log(M_{\text{halo}}) = 13 M_{\odot}$) dex in TNG-Cluster and 0.48 dex to 0.19 dex in TNG 300 respectively.

16 hidden channels, SiLU activations [12], and 16 outputs. These operate over edges connected to each node, using max pooling to aggregate edge information to each node³. The node output is concatenated with its initial feature (M_*) and passed through a 3-layer MLP. The GNN predicts two quantities [15]: M_{halo} and the expected log variance of M_{halo} at a given M_* .

GNN Optimisation. We employ the METIS algorithm to partition the training set into 48 parts (see ClusterLoader class in PyTorch Geometric [6, 11]), allowing us to handle large graph datasets efficiently. The model is trained with the AdamW optimiser [19] at an initial learning rate of 10^{-2} and weight decay of 10^{-4} . A scheduler reduces the learning rate by 0.2 if validation loss stagnates by 10^{-3} for 15 epochs. Early stopping occurs after 35 epochs of no improvement, with a maximum of 300 epochs. On an Nvidia A6000 GPU, training takes 20 minutes and inference takes under 1 second.

4 Results

Table 1 compares model performance for predicting M_{halo} from galaxies residing in clusters for the validation and test datasets. We additionally show the scatter of M_{halo} in the first row, which represents the most naive “prediction” of the sample mean. Below, we present the results for the baseline models and GNN model. We display scatter plots of the true versus predicted masses for the TNG-Cluster cross-validation data set in Figure 2a and TNG300 test set in Figure 2b.

The simplest RF model exhibits high error and very low predictive power.⁴ When we augment the RF model with Δ_G , the performance improves, demonstrating that galaxy environments contain vital information for the SHMR. Nonetheless, the RF models systematically underpredict M_{halo} for the highest-mass galaxies and yield high error.

GNNs greatly outperform RF models, as indicated by the right-most panels of Figures 2a and 2b. Running the same experiments using XGBoost (which is more prone to overfitting), we find a significant improvement over RF but not enough to surpass GNNs. We find that the GNN performance

³This helps the GNN to effectively capture the neighbouring features.

⁴In fact, for the TNG300 test set, the simplest RF model produces even higher error than the scatter inherent to the data. We ascribe this to the RF model’s significant negative bias (i.e., systematic underprediction).

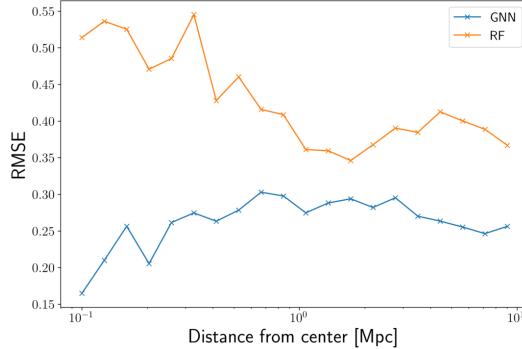


Figure 3: Validation set RMSE as a function of distance from cluster center. Results shown for the GNN (blue) and RF with M_* and Δ_G (orange).

on the training and validation sets translates to accurate predictions on the independent test set. For nearly all metrics in Table 1, the GNN outperforms the RF models for cross-validation and test sets.

5 Discussion

5.1 Model performance as a function of local environment

In Figure 3, we show the cross-validation RMSE as a function of distance from the cluster center for the GNN and RF (M_* and Δ_G) models; the GNN significantly outperforms the RF across all distance bins. Notably, the RF model performance suffers for galaxies closer to the center of the cluster. One potential explanation for this discrepancy is that the RF does not account for the dense cluster environment, where interactions such as tidal stripping can lead to significant loss of M_{halo} .⁵ In contrast, the GNN model outperforms the RF due by leveraging information from galaxy pairwise distances and line-of-sight velocities.

5.2 Comparison against previous work

Previous studies have used ML to estimate galaxy properties from dark matter halos [16, 1], i.e. the inverse of the problem we tackle. Some works employ feature importance from decision tree-based methods [21, 41], while others use reinforcement learning to connect halo properties to galaxies [23]. Convolutional neural networks (CNNs) and GNNs have also been used to predict galaxy stellar masses from simulated halos [5, 40, 41].

Several works have used ML methods to predict cluster halo masses from observable parameters such as X-ray brightness and Sunyaev-Zel'dovich decrements [26, 13]. [42] compare how different cluster observables fare when pixelised as inputs to a CNN. [18] use GNNs to predict M_{halo} directly from galaxy point clouds, but their training dataset (the much smaller TNG50 simulation) does not contain many rare galaxy clusters. Our work is the first to train and test GNNs for predicting halo masses in the extremely overdense regime of galaxy clusters.

6 Conclusions, Limitations and Future Work

In this work, we predict M_{halo} for simulated galaxies using their stellar masses, $2D$ projected positions, and line-of-sight velocities (i.e., x, y, v_z) with the TNG-Cluster simulation for training and TNG300 for testing. We evaluated both Random Forest (RF) models and Graph Neural Networks (GNNs). The key findings are:

1. The GNN model significantly outperforms RF model, even when the latter is provided Δ_G as a parameter. This suggests that GNNs capture the underlying spatial relationships and substructures within clusters, as shown in Table 1 and Figures 2 and 3.
2. The GNN maintains its predictive power when tested on the independent TNG300 dataset, demonstrating that the model generalises across the IllustrisTNG simulation suite.

Despite our promising results, models trained on one simulation may face challenges when applied to other simulations or real observational data. Machine learning models are often susceptible to do-

⁵Due to line-of-sight effects, not all galaxies at small *projected* distances experience significant tidal stripping.

main shift, where their performance degrades when applied to datasets that differ from their training data [36, 17]. In our case, the comparable performance between the TNG-Cluster cross-validation and TNG300 test datasets suggests that the GNN model may be robust to domain shift within the IllustrisTNG suite. This robustness could be attributed to the GNN’s ability to learn generalizable symbolic relationships [10]. Further tests using other simulation physics or with observed datasets (e.g., galaxies at other redshifts) are needed before we can conclude that this method is fully generalizable.

In future work, we will account for observational effects like contaminating galaxies in projection, missing data, and photometric redshift uncertainties, as well as broader concerns about domain shift in ML (see e.g. [7]). Aside from additional validation on other cosmological simulations [33], we will test on observational data using published M_{halo} estimates for well-known galaxy clusters (e.g., [20, 38]). With upcoming telescopes like the Roman Space Telescope [34] and Rubin Observatory [9], we will be able to study GNN applications to large galaxy cluster samples in the wide-field domain.

References

- [1] Shankar Agarwal, Romeel Davé, and Bruce A Bassett. Painting galaxies into dark matter haloes using machine learning. *Monthly Notices of the Royal Astronomical Society*, 478(3):3410–3422, August 2018.
- [2] Mark Birkinshaw. The Sunyaev–Zel’dovich effect. *Physics Reports*, 310(2):97–195, March 1999.
- [3] L. E. Bleem, B. Stalder, T. de Haan, K. A. Aird, S. W. Allen, D. E. Applegate, M. L. N. Ashby, M. Bautz, M. Bayliss, B. A. Benson, S. Bocquet, M. Brodwin, J. E. Carlstrom, C. L. Chang, I. Chiu, H. M. Cho, A. Clocchiatti, T. M. Crawford, A. T. Crites, S. Desai, J. P. Dietrich, M. A. Dobbs, R. J. Foley, W. R. Forman, E. M. George, M. D. Gladders, A. H. Gonzalez, N. W. Halverson, C. Hennig, H. Hoekstra, G. P. Holder, W. L. Holzappel, J. D. Hrubes, C. Jones, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, J. Liu, M. Lueker, D. Luong-Van, A. Mantz, D. P. Marrone, M. McDonald, J. J. McMahan, S. S. Meyer, L. Mocanu, J. J. Mohr, S. S. Murray, S. Padin, C. Pryke, C. L. Reichardt, A. Rest, J. Ruel, J. E. Ruhl, B. R. Saliwanchik, A. Saro, J. T. Sayre, K. K. Schaffer, T. Schrabback, E. Shirokoff, J. Song, H. G. Spieler, S. A. Stanford, Z. Staniszewski, A. A. Stark, K. T. Story, C. W. Stubbs, K. Vanderlinde, J. D. Vieira, A. Vikhlinin, R. Williamson, O. Zahn, and A. Zenteno. Galaxy Clusters Discovered via the Sunyaev-Zel’dovich Effect in the 2500-Square-Degree SPT-SZ Survey. *The Astrophysical Journal Supplement Series*, 216(2):27, February 2015.
- [4] James S. Bullock and Michael Boylan-Kolchin. Small-Scale Challenges to the Λ CDM Paradigm. *Annual Review of Astronomy and Astrophysics*, 55(1):343–387, August 2017.
- [5] Urmila Chadayammuri, Michelle Ntampaka, John ZuHone, Ákos Bogdán, and Ralph P. Kraft. Painting baryons on to N-body simulations of galaxy clusters with image-to-image deep learning. *Monthly Notices of the Royal Astronomical Society*, 526(2):2812–2829, December 2023.
- [6] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. ClusterGCN: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 257–266, July 2019.
- [7] A. Čiprijanović, D. Kafkes, K. Downey, S. Jenkins, G. N. Perdue, S. Madireddy, T. Johnston, G. F. Snyder, and B. Nord. DeepMerge - II. Building robust deep learning algorithms for merging galaxy identification across domains. *Monthly Notices of the Royal Astronomical Society*, 506(1):677–691, September 2021.
- [8] Douglas Clowe, Maruša Bradač, Anthony H. Gonzalez, Maxim Markevitch, Scott W. Randall, Christine Jones, and Dennis Zaritsky. A Direct Empirical Proof of the Existence of Dark Matter. *The Astrophysical Journal*, 648(2):L109–L113, September 2006.
- [9] LSST Dark Energy Science Collaboration. Large Synoptic Survey Telescope: Dark Energy Science Collaboration, November 2012.
- [10] Miles Cranmer, Alvaro Sanchez-Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. Discovering Symbolic Models from Deep Learning with Inductive Biases, 2020.

- [11] Matthias Fey and Jan Eric Lenssen. Fast Graph Representation Learning with PyTorch Geometric, 2019.
- [12] Dan Hendrycks and Kevin Gimpel. Gaussian Error Linear Units (GELUs), 2016.
- [13] Matthew Ho, John Soltis, Arya Farahi, Daisuke Nagai, August Evrard, and Michelle Ntampaka. Benchmarks and explanations for deep learning estimates of X-ray galaxy cluster masses. *Monthly Notices of the Royal Astronomical Society*, 524(3):3289–3302, September 2023.
- [14] Tin Kam Ho. Random decision forests. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 1, pages 278–282 vol.1, August 1995.
- [15] Niall Jeffrey and Benjamin D. Wandelt. Solving high-dimensional parameter inference: Marginal posterior densities & Moment Networks, 2020.
- [16] Harshil M. Kamdar, Matthew J. Turk, and Robert J. Brunner. Machine learning and cosmological simulations – I. Semi-analytical models. *Monthly Notices of the Royal Astronomical Society*, 455(1):642–658, January 2016.
- [17] Wouter M. Kouw and Marco Loog. An introduction to domain adaptation and transfer learning, January 2019.
- [18] Austin J. Larson, John F. Wu, and Craig Jones. Predicting dark matter halo masses from simulated galaxy images and environments, July 2024.
- [19] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, January 2019.
- [20] J. M. Lotz, A. Koekemoer, D. Coe, N. Grogin, P. Capak, J. Mack, J. Anderson, R. Avila, E. A. Barker, D. Borncamp, G. Brammer, M. Durbin, H. Gunning, B. Hilbert, H. Jenkner, H. Khandrika, Z. Levay, R. A. Lucas, J. MacKenty, S. Ogaz, B. Porterfield, N. Reid, M. Robberto, P. Royle, L. J. Smith, L. J. Storrie-Lombardi, B. Sunnquist, J. Surace, D. C. Taylor, R. Williams, J. Bullock, M. Dickinson, S. Finkelstein, P. Natarajan, J. Richard, B. Robertson, J. Tumlinson, A. Zitrin, K. Flanagan, K. Sembach, B. T. Soifer, and M. Mountain. The Frontier Fields: Survey Design and Initial Results. *The Astrophysical Journal*, 837(1):97, March 2017.
- [21] Christopher C. Lovell, Stephen M. Wilkins, Peter A. Thomas, Matthieu Schaller, Carlton M. Baugh, Giulio Fabbian, and Yannick Bahé. A machine learning approach to mapping baryons on to dark matter haloes using the EAGLE and C-EAGLE simulations. *Monthly Notices of the Royal Astronomical Society*, 509(4):5046–5061, February 2022.
- [22] Tobias A. Marriage, Viviana Acquaviva, Peter A. R. Ade, Paula Aguirre, Mandana Amiri, John William Appel, L. Felipe Barrientos, Elia S. Battistelli, J. Richard Bond, Ben Brown, Bryce Burger, Jay Chervenak, Sudeep Das, Mark J. Devlin, Simon R. Dicker, W. Bertrand Doriese, Joanna Dunkley, Rolando Dünner, Thomas Essinger-Hileman, Ryan P. Fisher, Joseph W. Fowler, Amir Hajian, Mark Halpern, Matthew Hasselfield, Carlos Hernández-Monteagudo, Gene C. Hilton, Matt Hilton, Adam D. Hincks, Renée Hlozek, Kevin M. Huffenberger, David Handel Hughes, John P. Hughes, Leopoldo Infante, Kent D. Irwin, Jean Baptiste Juin, Madhuri Kaul, Jeff Klein, Arthur Kosowsky, Judy M. Lau, Michele Limon, Yen-Ting Lin, Robert H. Lupton, Danica Marsden, Krista Martocci, Phil Mauskopf, Felipe Menanteau, Kavilan Moodley, Harvey Moseley, Calvin B. Netterfield, Michael D. Niemack, Michael R. Nolta, Lyman A. Page, Lucas Parker, Bruce Partridge, Hernan Quintana, Erik D. Reese, Beth Reid, Neelima Sehgal, Blake D. Sherwin, Jon Sievers, David N. Spergel, Suzanne T. Staggs, Daniel S. Swetz, Eric R. Switzer, Robert Thornton, Hy Trac, Carole Tucker, Ryan Warne, Grant Wilson, Ed Wollack, and Yue Zhao. The Atacama Cosmology Telescope: Sunyaev-Zel’dovich-Selected Galaxy Clusters at 148 GHz in the 2008 Survey. *The Astrophysical Journal*, 737(2):61, August 2011.
- [23] Benjamin P. Moster, Thorsten Naab, Magnus Lindström, and Joseph A. O’Leary. GalaxyNet: Connecting galaxies and dark matter haloes with deep neural networks and reinforcement learning in large volumes. *Monthly Notices of the Royal Astronomical Society*, 507(2):2115–2136, October 2021.
- [24] Dylan Nelson, Annalisa Pillepich, Mohammadreza Ayromlou, Wonki Lee, Katrin Lehle, Eric Rohr, and Nhut Truong. Introducing the TNG-Cluster simulation: Overview and the physical properties of the gaseous intracluster medium. *Astronomy & Astrophysics*, 686:A157, June 2024.

- [25] Dylan Nelson, Volker Springel, Annalisa Pillepich, Vicente Rodriguez-Gomez, Paul Torrey, Shy Genel, Mark Vogelsberger, Ruediger Pakmor, Federico Marinacci, Rainer Weinberger, Luke Kelley, Mark Lovell, Benedikt Diemer, and Lars Hernquist. The IllustrisTNG simulations: Public data release. *Computational Astrophysics and Cosmology*, 6(1):2, May 2019.
- [26] M. Ntampaka, J. ZuHone, D. Eisenstein, D. Nagai, A. Vikhlinin, L. Hernquist, F. Marinacci, D. Nelson, R. Pakmor, A. Pillepich, P. Torrey, and M. Vogelsberger. A Deep Learning Approach to Galaxy Cluster X-Ray Masses. *The Astrophysical Journal*, 876(1):82, May 2019.
- [27] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.
- [28] P. J. E. Peebles. Tests of cosmological models constrained by inflation. *The Astrophysical Journal*, 284:439, September 1984.
- [29] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, R. Battye, K. Benabed, A. Benoît, A. Benoit-Lévy, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J.-F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, R.-R. Chary, H. C. Chiang, J. Chluba, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. De Bernardis, A. De Rosa, G. De Zotti, J. Delabrouille, F.-X. Désert, E. Di Valentino, C. Dickinson, J. M. Diego, K. Dolag, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, M. Farhang, J. Fergusson, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, M. Gerbino, T. Ghosh, M. Giard, Y. Giraud-Héraud, E. Giusarma, E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Hamann, F. K. Hansen, D. Hanson, D. L. Harrison, G. Helou, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, Z. Huang, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J.-M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, J. P. Leahy, R. Leonardi, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marchini, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Masi, S. Matarrese, P. McGehee, P. R. Meinhold, A. Melchiorri, J.-B. Melin, L. Mendes, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M.-A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Novello, D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, G. Prézeau, S. Prunet, J.-L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, C. Rosset, M. Rossetti, G. Roudier, B. Rouillé d’Orfeuil, M. Rowan-Robinson, J. A. Rubiño-Martín, B. Rusholme, N. Said, V. Salvatelli, L. Salvati, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, P. Serra, E. P. S. Shellard, L. D. Spencer, M. Spinelli, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, D. Sutton, A.-S. Suur-Uski, J.-F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, M. Tucci, J. Tuovinen, M. Türlér, G. Umana, L. Valenziano, J. Valiviita, F. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Wilkinson, D. Yvon, A. Zacchei, and A. Zonca. *Planck* 2015 results: XIII. Cosmological parameters. *Astronomy & Astrophysics*, 594:A13, October 2016.
- [30] E. S. Rykoff, E. Rozo, M. T. Busha, C. E. Cunha, A. Finoguenov, A. Evrard, J. Hao, B. P. Koester, A. Leauthaud, B. Nord, M. Pierre, R. Reddick, T. Sadibekova, E. S. Sheldon, and R. H. Wechsler. redMaPPer. I. Algorithm and SDSS DR8 Catalog. *The Astrophysical Journal*, 785(2):104, April 2014.

- [31] H. Sampaio-Santos, Y. Zhang, R. L. C. Ogando, T. Shin, Jesse B. Golden-Marx, B. Yanny, K. Herner, M. Hilton, A. Choi, M. Gatti, D. Gruen, B. Hoyle, M. M. Rau, J. De Vicente, J. Zuntz, T. M. C. Abbott, M. Aguena, S. Allam, J. Annis, S. Avila, E. Bertin, D. Brooks, D. L. Burke, M. Carrasco Kind, J. Carretero, C. Chang, M. Costanzi, L. N. da Costa, H. T. Diehl, P. Doel, S. Everett, A. E. Evrard, B. Flaugher, P. Fosalba, J. Frieman, J. García-Bellido, E. Gaztanaga, D. W. Gerdes, R. A. Gruendl, J. Gschwend, G. Gutierrez, S. R. Hinton, D. L. Hollowood, K. Honscheid, D. J. James, M. Jarvis, T. Jeltema, K. Kuehn, N. Kuropatkin, O. Lahav, M. A. G. Maia, M. March, J. L. Marshall, R. Miquel, A. Palmese, F. Paz-Chinchón, A. A. Plazas, E. Sanchez, B. Santiago, V. Scarpine, M. Schubnell, M. Smith, E. Suchyta, G. Tarle, D. L. Tucker, T. N. Varga, and R. H. Wechsler. Is diffuse intracluster light a good tracer of the galaxy cluster matter distribution? *Monthly Notices of the Royal Astronomical Society*, 501(1):1300–1315, February 2021.
- [32] F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 20(1):61–80, January 2009.
- [33] Joop Schaye, Roi Kugel, Matthieu Schaller, John C. Helly, Joey Braspenning, Willem Elbers, Ian G. McCarthy, Marcel P. van Daalen, Bert Vandenbroucke, Carlos S. Frenk, Juliana Kwan, Jaime Salcido, Yannick M. Bahé, Josh Borrow, Evgenii Chaikin, Oliver Hahn, Filip Huško, Adrian Jenkins, Cedric G. Lacey, and Folkert S. J. Nobels. The FLAMINGO project: Cosmological hydrodynamical simulations for large-scale structure and galaxy cluster surveys. *Monthly Notices of the Royal Astronomical Society*, 526(4):4978–5020, October 2023.
- [34] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, B. Macintosh, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, D. Benford, M. Hudson, W. S. Jeong, Y. Mellier, W. Traub, T. Yamada, P. Capak, J. Colbert, D. Masters, M. Penny, D. Savransky, D. Stern, N. Zimmerman, R. Barry, L. Bartusek, K. Carpenter, E. Cheng, D. Content, F. Dekens, R. Demers, K. Grady, C. Jackson, G. Kuan, J. Kruk, M. Melton, B. Nemati, B. Parvin, I. Poberezhskiy, C. Peddie, J. Ruffa, J. K. Wallace, A. Whipple, E. Wollack, and F. Zhao. Wide-field infrared survey telescope-astronomy focused telescope assets wfirst-afta 2015 report, 2015.
- [35] Volker Springel, Simon D. M. White, Giuseppe Tormen, and Guinevere Kauffmann. Populating a cluster of galaxies - I. Results at $z=0$. *Monthly Notices of the Royal Astronomical Society*, 328(3):726–750, December 2001.
- [36] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pages 2058–2065, Phoenix, Arizona, February 2016. AAAI Press.
- [37] S. Vegetti, S. Birrer, G. Despali, C. D. Fassnacht, D. Gilman, Y. Hezaveh, L. Perreault Levasseur, J. P. McKean, D. M. Powell, C. M. O’Riordan, and G. Vernardos. Strong Gravitational Lensing as a Probe of Dark Matter. *Space Science Reviews*, 220(5):58, July 2024.
- [38] John R. Weaver, Sam E. Cutler, Richard Pan, Katherine E. Whitaker, Ivo Labbé, Sedona H. Price, Rachel Bezanson, Gabriel Brammer, Danilo Marchesini, Joel Leja, Bingjie Wang, Lukas J. Furtak, Adi Zitrin, Hakim Atek, Iryna Chemerynska, Dan Coe, Pratika Dayal, Pieter van Dokkum, Robert Feldmann, Natascha M. Förster Schreiber, Marijn Franx, Seiji Fujimoto, Yoshinobu Fudamoto, Karl Glazebrook, Anna de Graaff, Jenny E. Greene, Stéphanie Juneau, Susan Kassin, Mariska Kriek, Gourav Khullar, Michael V. Maseda, Lamiya A. Mowla, Adam Muzzin, Themiyi Nanayakkara, Erica J. Nelson, Pascal A. Oesch, Camilla Pacifici, Casey Papovich, David J. Setton, Alice E. Shapley, Heath V. Shipley, Renske Smit, Mauro Stefanon, Edward N. Taylor, Andrea Weibel, and Christina C. Williams. The UNCOVER Survey: A First-look HST + JWST Catalog of 60,000 Galaxies near A2744 and beyond. *The Astrophysical Journal Supplement Series*, 270:7, January 2024.
- [39] Risa H. Wechsler and Jeremy L. Tinker. The Connection Between Galaxies and Their Dark Matter Halos. *Annual Review of Astronomy and Astrophysics*, 56(1):435–487, September 2018.
- [40] John F. Wu and Christian Kragh Jespersen. Learning the galaxy-environment connection with graph neural networks, June 2023.
- [41] John F. Wu, Christian Kragh Jespersen, and Risa H. Wechsler. How the Galaxy-Halo Connection Depends on Large-Scale Environment, February 2024.

- [42] Z Yan, A J Mead, L Van Waerbeke, G Hinshaw, and I G McCarthy. Galaxy cluster mass estimation with deep learning and hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 499(3):3445–3458, November 2020.

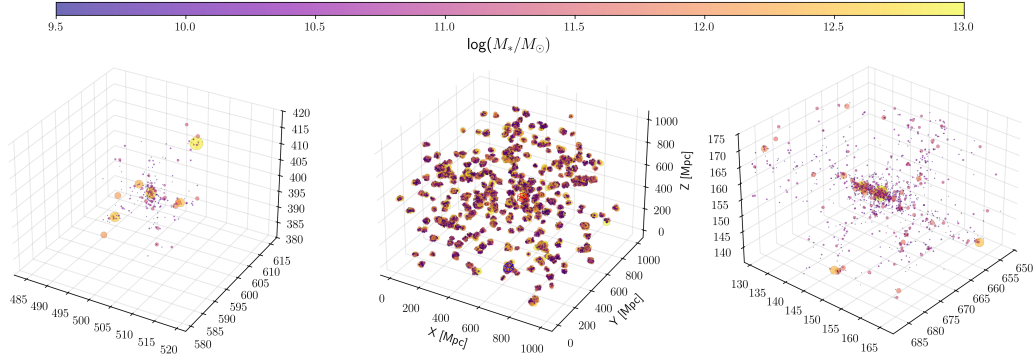


Figure 4: Spatial distribution of halos within the TNG-Cluster simulation. The middle panel shows the full simulation, and the left and right panels highlight two example galaxy clusters. The boundaries of these clusters are marked as blue and red boxes in the middle panel.

Table 2: Summary of cuts applied to the TNG-Cluster data. Here, N_* refers to number of stellar particles, M_\odot refers to solar mass, R_{200} refers to the virial radius of the halo.

Sample	Number of Subhalos
Full TNG-Cluster catalog	10,378,451
— within mass cuts - $N_* > 50$; $\log(M_*/M_\odot) > 9.5$; $\log(M_{\text{halo}}/M_\odot) > 10.5$	154,120
— within $< 10 \times R_{200}$ of the cluster halo	127,165
Selection Criteria - $\log(M_{\text{halo}}/M_\odot) > 11$; within 10 Mpc of the cluster halo	
TNG-Cluster cross-validation	60,756
TNG300 Test Set	34,689

A TNG-Cluster Additional Details

Galaxies in the TNG-Cluster training data are shown in Figure 4. To mimic astronomical observations of galaxies, we project the galaxy clusters along the z axis, which is chosen to be the line of sight. This procedure bridges the gap between simulation data and spectroscopic observations, which typically capture two spatial dimensions (x, y) and line-of-sight velocities (v_z). We also apply quality cuts to the simulation in Table 2 to ensure a complete sample of massive, well-resolved galaxies.

We split the TNG-Cluster data into training and validation sets by implementing a k -fold cross-validation strategy based on cluster IDs rather than traditional random splits. This method isolates subhalos according to their cluster IDs while ensuring that all subhalos from a single cluster remain within the same fold. One potential caveat of this method is that we do not include the contaminating structure along the line-of-sight from other clusters which might be in a different k -fold.