
A Halo Merger Tree Generation and Evaluation Framework

Sandra Robles^{1,2} Jonathan S. Gómez³ Adín Ramírez Rivera⁴ Jenny A. González³ Nelson D. Padilla³
Diego Dujovne⁵

Abstract

Semi-analytic models are best suited to compare galaxy formation and evolution theories with observations. These models rely heavily on halo merger trees, and their realistic features (i.e., no drastic changes on halo mass or jumps on physical locations). Our aim is to provide a new framework for halo merger tree generation that takes advantage of the results of large volume simulations, with a modest computational cost. We treat halo merger tree construction as a matrix generation problem, and propose a Generative Adversarial Network that learns to generate realistic halo merger trees. We evaluate our proposal on merger trees from the EAGLE simulation suite, and show the quality of the generated trees.

1. Introduction

In cosmology and astrophysics, galaxy formation and evolution is a complex non-linear problem that entails very different physical phenomena spanning a wide range of scales, from large scale structure formation, dark matter (DM) gravitational collapse in halos down to star formation, and metal enrichment among others. Two strategies are used to tackle this problem, semi-analytic models (SAMs) (White & Frenk, 1991) and hydrodynamical simulations (Carlberg et al., 1990; Katz et al., 1992).

Large-scale hydrodynamical simulations solve the equations of gravity and fluid dynamics simultaneously. Nevertheless, high-resolution simulations are extremely computational

resource intensive. SAMs, on the other hand, describe physical processes behind galaxy formation and evolution in a simpler way and require less CPU time. Furthermore, they can simulate much larger cosmological volumes than hydrodynamical simulations and consequently they can be tuned to observations. These models can be populated with halo merger trees and then phenomenological prescriptions of the baryonic physics governing galaxy formation are introduced.

A halo merger tree describes the hierarchical formation history of a DM halo by tracing all its progenitors back in time. In the standard model of cosmology, large DM halos are formed by the coalescence of smaller progenitors. Therefore, halo merger trees encapsulate the growth and merger history of DM halos. If the progenitors contain central galaxies, halo mergers¹ eventually give rise to galaxy mergers. In this way, galaxy evolution is directly linked to the halo merger history.

Two different methodologies are used to produce halo merger trees. The first method takes advantage of the extended Press-Schechter formalism (Bond et al., 1991) and Monte Carlo simulations. It allows rapid merger tree construction in large volumes with high mass resolutions (Somerville & Primack, 1999; Cole et al., 2000; Somerville et al., 2008; Benson & Bower, 2010; Ricciardelli & Franceschini, 2010), but a single tree per halo can be built at a time. In addition, this method often yields mass assembly histories that are in disagreement with simulations (Jiang & van den Bosch, 2014).

Despite being computationally intensive, the most popular method for merger tree construction is based on high-resolution DM only (N -body) simulations, since it results in a more realistic evolutionary history of the halos and yields thousands of merger trees at once (Kauffmann et al., 1999; Hatton et al., 2003; De Lucia et al., 2004; Croton et al., 2006; Bower et al., 2006; Guo et al., 2011; Lee et al., 2014). An important caveat of merger tree construction using simulations is the limitation in mass resolution. Since massive halos are better resolved than those with low mass, the former can be traced down to progenitor masses that

¹Departamento de Física Teórica, Universidad Autónoma de Madrid, Madrid, Spain. ²ARC Centre of Excellence for Particle Physics at the Terascale, School of Physics, The University of Melbourne, Victoria, Australia. ³Instituto de Astrofísica, Pontificia Universidad Católica de Chile, Santiago, Chile. ⁴Institute of Computing, University of Campinas, São Paulo, Brazil. ⁵Escuela de Informática y Telecomunicaciones, Universidad Diego Portales, Santiago, Chile. Correspondence to: Sandra Robles <sandra.robles@unimelb.edu.au>, Adín Ramírez Rivera <adin@ic.unicamp.br>.

¹We refer to them as ‘mergers’ throughout this paper. In that sense, we are describing the event of halos merging together.

are a smaller fraction of the final halo mass than the latter. Merger trees are built in two steps. First, each output time-step of the simulation (snapshot) is scanned to find DM collapsed structures, halos and produce halo catalogs. Then, or simultaneously to the DM halo search, halos are linked across different snapshots giving us the merger history of each halo (the structure known as merger tree), where the last leaf (last descendent) of the main branch (the longest branch) is the halo whose assembly history the tree encloses. The remaining leaves of the tree are the progenitors of the aforementioned halo.

Halo merger trees play an important role in modern galaxy formation theory, and are the key ingredient of SAMs. More specifically, mock halo catalogs from one step in time (snapshot) of a DM only cosmological simulation are the usual inputs of SAMs. Next, these DM halos are populated with galaxies using a given baryonic physics prescription. A self-consistent evolution of these galaxies requires the knowledge of the growth and mass assembly history of the DM halo that hosts them, which is precisely what halo merger trees provide. In this sense, SAMs rely heavily on both a precise halo identification and, realistic and well-constructed merger trees, with no drastic changes in the halo mass, or jumps in physical location from one snapshot to the next.

Nowadays, SAMs are best suited to compare theoretical predictions with galaxy surveys, thanks to their flexibility to explore physical phenomena in different ways with relatively computational inexpensive runs (see e.g. Cole, 1991; Lacey & Silk, 1991; White & Frenk, 1991; Kauffmann & White, 1993; Kauffmann et al., 1993; 1999; Cole et al., 1994; 2000; Bower et al., 2006; Somerville et al., 2008; Lagos et al., 2018). As mentioned before, the standard method to produce realistic merger trees is based on high resolution N -body simulations that require long runs. Our main contribution is to provide astrophysicists with a new framework for halo merger tree construction, taking advantage of the best features of large volume simulations, but with a modest computational expense. To this end, we designed a Generative Adversarial Network (GAN) (Goodfellow et al., 2014) that generates merger trees (represented as matrices) with characteristics of a subset of trees from the EAGLE simulation suite (Schaye et al., 2015; Crain et al., 2015). Additionally, we introduce a set of criteria to assess the quality of the generated merger trees when compared to those obtained from simulations.

2. Halo Merger Tree Generation

Halo merger trees are the backbone of SAMs. With the upcoming of new major observational facilities, robust predictions from SAMs will be needed to understand the nature of the processes that could imprint observational features on galaxies. Hence, our goal is to learn a lean generative

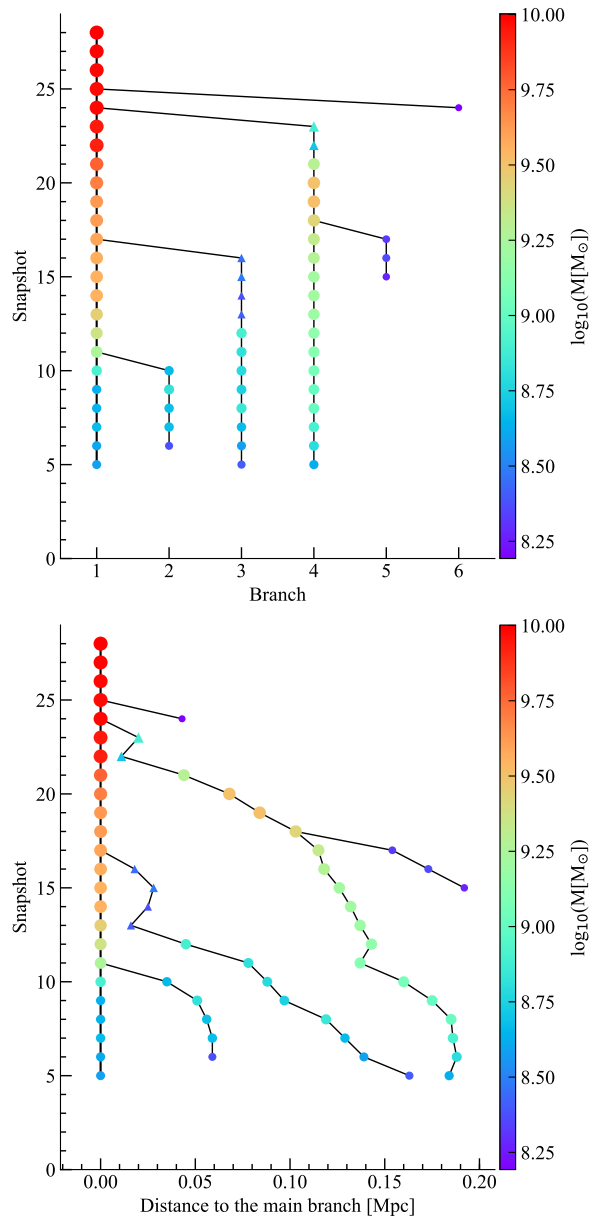


Figure 1. Top: Generated merger tree, the color map denotes progenitor masses and the progenitor type is represented by circles (main halos) and triangles (subhalos). Bottom: Same merger tree in the plane snapshot vs. distance.

model, based on GANs, that generates merger trees closer to the training data. To that end, first, we represent merger trees as matrices. Next, we propose a learning framework to train our model with different sets of variables.

Merger Tree Representation. We selected halo merger trees for central subhalos at $z = 0$ (current time of Universe) with masses in the range $10^{10}M_{\odot} \leq M \leq 10^{11}M_{\odot}$ from the EAGLE reference simulation with comoving cubic box length of 100 Mpc and 1504^3 particles. This simula-

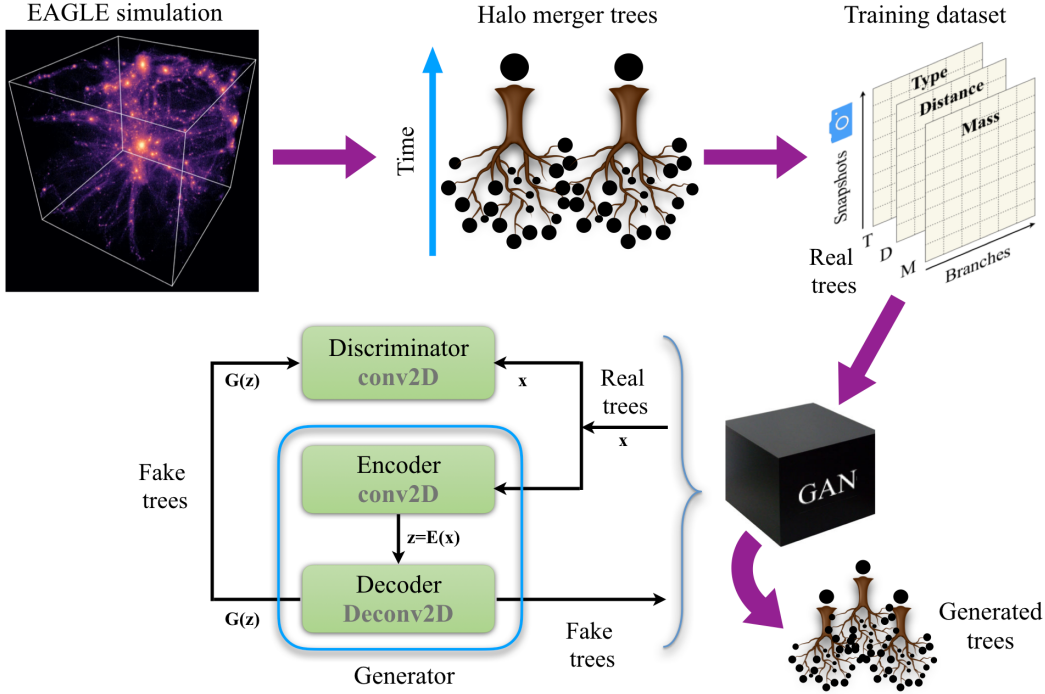


Figure 2. Halo merger tree generation process.

tion supplies a large enough amount of trees for training purposes. We obtained the trees by traversing the subhalo table of the EAGLE public catalog (McAlpine et al., 2016). The EAGLE simulation suite is a set of cosmological hydrodynamical simulations with cosmological parameters taken from Planck Collaboration et al. (2014), that track the evolution of both, dark matter and baryonic particles, from redshift $z = 127$ to $z = 0$.

Next, we stored each relevant variable for consistent merger tree description, namely mass of the halo progenitors, distance of every progenitor in a branch to the corresponding progenitor in the main branch and a discrete variable that indicates if the progenitor is a main halo or a subhalo,² in a matrix format that reflects the tree structure (see Fig. 1), one matrix per each variable. Columns in this format represent the branches of a tree, and rows correspond to 29 snapshots between redshifts $z = 20$ (snapshot 0) and $z = 0$ (snapshot 28, present day). The first column is the so-called main branch that is the longest branch of a tree. At snapshot 28, this branch contains the last descendant, i.e., the halo whose mass assembly history the merger tree tells about. Note that the main branch not necessarily starts at snapshot 0, indeed the location in time of the first progenitor depends on the mass of the final halo and the resolution of the simulation.

²Note, that halos can be nested, i.e., they can contain substructures. The halo at the center of the gravitational potential is referred to as main or central halo. The others are called subhalos, often regarded as orbiting main halos.

Recall that massive halos can be traced down to low mass progenitors at earlier steps in time and hence their merger trees tend to have longer main branches than those of low massive halos. Progenitors in other branches eventually merge with the corresponding progenitor in the main branch (same snapshot) in the next time step. Matrix elements are filled following the structure of a tree. Zeros in each matrix denote no progenitor in that particular branch and snapshot.

Neural Network Model. To generate merger trees, we selected a deep convolutional GAN architecture, based on DCGAN (Dosovitskiy et al., 2015; Radford et al., 2015). More specifically, the generator features an encoder-decoder architecture (Bengio et al., 2013) that learns to produce the matrix representation of the halo merger trees. We detail the architecture in Appendix A. Our design for this model uses row- and column-wise convolutional filters. These filters intend to reproduce the operations within a branch (column-wise filters), and between the branches to produce mergers (row-wise filters).

We trained this model with the dataset in matrix form, where the merger tree variables (mass, distance to the main branch, and progenitor type) correspond to input channels. The reconstruction loss function that drives the learning process comprises cross-entropy losses applied individually to each of the variables. We have restricted our initial dataset for training purposes to trees with only six branches, since they are the most abundant in the aforementioned mass

range, and the reconstruction of trees with a fixed number of branches is a good test of the viability of GANs to learn merger tree representations. In Fig. 2, we show a diagram that summarizes the entire merger tree generation process.

3. Evaluation of Merger Trees

As stated before, SAMs require well-constructed merger trees in order to make consistent theoretical predictions. Therefore, we developed a set of criteria to evaluate both the quality of the generated trees, and the effect of adding variables or inputs to the GAN on the generated trees.

First of all, GAN outputs should resemble a real merger tree, i.e., the learned representation should contain no reappearing halos after a merger, the absence of a progenitor in a branch must be represented by a zero value, the last row (snapshot 28) should contain only one halo, the last descendant. This criterion holds for all the variables considered in this work.

Merger trees tell the mass assembly history of halos. Hence, the first variable to reproduce is the progenitor mass. The second criterion takes into account a fair reproduction of the masses, i.e., masses in a branch should be within the mass range of the dataset used for training. It is worth mentioning that progenitor masses in a branch are usually expected to grow monotonically or at least preserve its mass, which is an assumption for SAMs (Lacey & Silk, 1991; White & Frenk, 1991; Cole et al., 1994; 2000; Gonzalez-Perez et al., 2014; Lacey et al., 2016; Lagos et al., 2018). However, there are other variables which can alter that expected behavior, such as the distance to the main branch and the progenitor type.

The second variable considered is the distance between the centers of a progenitor in a given branch with respect to the progenitor in the main branch at the same step in time. In general, this distance should decrease as the merging event approaches, but the exact step in time when this should occur varies with the merger tree. Nevertheless, we can extract statistic distributions about the behavior of this input variable for several snapshots before the merger occurs, for both samples the real (training dataset) and generated merger trees (see Appendix C). A way to evaluate whether or not the distances are well reproduced by the learned representation, i.e., there are no dramatic jumps in location, is by comparing the real and fake distributions with the Kolmogorov-Smirnov (KS) test. Unlike the mass, there is no criterion for a correct range in which the prediction of the distance between two merging progenitors should lie. Note that, together, mass and distance are a measure of the gravitational pull between merging progenitors.

The third variable accounts for the progenitor type, i.e., a progenitor in a branch can be either a main (central) halo or a subhalo. Progenitors in the main branch are expected to

Table 1. Kolmogorov-Smirnov test for mass, distance to the main branch and number of snapshots that a progenitor spend as a subhalo before the merging event for merger trees generated with 1, 2 and 3 variables. Dashes denote unavailable tests.

KS Test	1 Variable	2 Variables	3 Variables
Mass	0.43	0.57	0.21
MH mass	–	–	0.14
SH mass	–	–	0.05
Distance	–	0.05	0.06
N. snaps. as subhalo	–	–	0.04

be central halos, with some exceptions at earlier snapshots. On the other hand, progenitors in other branches can be always main halos or become subhalos a few snapshots before the merger as a consequence of gravitational infall. Therefore, the fusion is strongly linked to the condition of the progenitors involved being central or subhalos (Diemand et al., 2006; Muldrew et al., 2011; Han et al., 2012; Onions et al., 2012; Elahi et al., 2011; Onions et al., 2013). Once again, the exact step in time this should happen depends on the specific tree. The condition of a progenitor being a main or a sub-halo affects the behavior of the other two variables. Subhalos are expected to be closer to the progenitor in the main branch they are going to merge, than central halos. Hence, subhalo masses are allowed to diminish in time as long they approach the merging event. We show an example of the interplay among the three variables in Fig. 1, where progenitors become subhalos before the merger in the second and third branches and their masses decrease. In fact, in both cases the progenitor mass diminishes one snapshot before they become subhalos, as they approach the progenitor in the main branch (see bottom panel of Fig. 1). Therefore, the mass gain and loss in a generated tree cannot be penalized arbitrarily in a particular tree, we should compare sample distributions instead. More examples of merger trees generated with 3 variables can be found in Appendix B.

In light of the foregoing, a final test to probe that the progenitor type is well generated by our GAN framework is to perform a KS test on the distribution of the number of snapshots a progenitor spend as a subhalo before merging with another progenitor.

4. Experiments and Results

We consider three different scenarios. The first configuration accounts for the mass reconstruction of halo progenitors, denoted henceforth as 1 variable. Then, we have trained our GAN with only the mass dataset. Next, we have introduced a second input, the distance to the main branch and trained the GAN with the mass and distance input dataset. This is

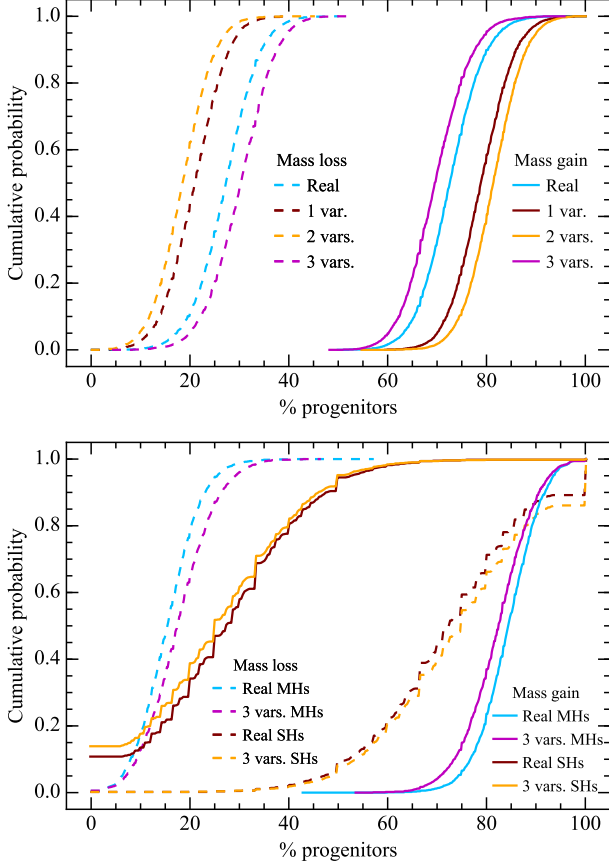


Figure 3. Top: Cumulative probability of the mass gain and loss for all the progenitors in all the merger trees in the training dataset (real), and in the samples of generated merger trees with 1, 2, and 3 variables. Bottom: Same for the progenitors that are main halos (MH) and subhalos (SH) in all the merger trees in the training dataset (real), and in the generated merger trees with three variables.

the 2 variable configuration. Afterwards, the 3 variable configuration introduces the progenitor type as a new variable, and, once again, we trained our GAN with the dataset for the three aforementioned variables.

Then, we evaluated the progenitor masses of the merger trees generated with 1, 2 and 3 variables according to Section 3. We found that this variable is well reconstructed except for small numerical artifacts when there is no progenitor in a branch and the corresponding mass is not exactly zero, but a very small number.

Next, we compared the distributions of the mass gain and loss for all the progenitors in the full sample of reconstructed merger trees for each of the above mentioned configurations with that of the training dataset (real trees). Note that the sample size of the generated merger trees is as large as that of the real trees. We represent the cumulative probability of these distributions in the top panel of Fig. 3. The KS test (see

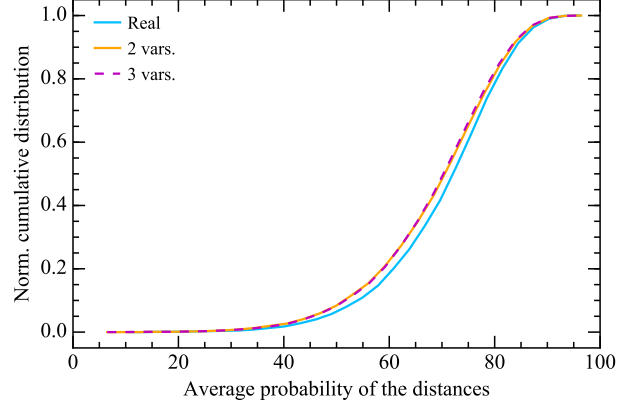


Figure 4. Normalized cumulative distribution of the averaged probabilities of the distances for the real merger trees and those generated with 2 and 3 variables.

first row of Table 1) shows that the 3 variable configurations yield the best mass reconstruction. To understand why the introduction of the progenitor type produce the best results, we have analyzed the mass gain and loss of the progenitors that are main halos (MH) and subhalos (SH) separately with the KS test. We show the cumulative probabilities of the MH and SH mass distributions in the bottom panel of Fig. 3. According to the results of the KS test (second, MH, and third, SH, rows of Table 1), the 3 variables configuration gives the best fit to the trained data due to the GAN’s ability to reproduce the mass gain and loss of SHs than those of the progenitors that are main halos.

To evaluate the reconstruction of the distance for the 2 and 3 variables configurations, as previously stated, we computed the probability distribution of the distance between merging progenitors for several snapshots before the merger takes place, for the training dataset and an equally large sample of reconstructed merger trees (see Appendix C). More specifically, we calculated the probability distribution (at all snapshots) of the distance to the main branch for the real sample, which will be our reference distribution in order to compute the probability of every distance in a merger tree of any sample. Then, for a given value of the distance, we calculated the average probability for each merger tree of the real, 2 and 3 variables samples. We show the normalized cumulative distribution of these average probabilities in Fig. 4. We compared the real and fake distributions with the KS test, and show it on the fourth row of Table 1. Both generated distributions yield equally good results with the 2 variables configurations giving a slightly better goodness of the fit.

Finally, we can perform a last test on the ensemble of generated merger trees with 3 variables. This test focus only on the fair reproduction of the third variable, the progenitor type. We show in Fig. 5 the cumulative distribution of the

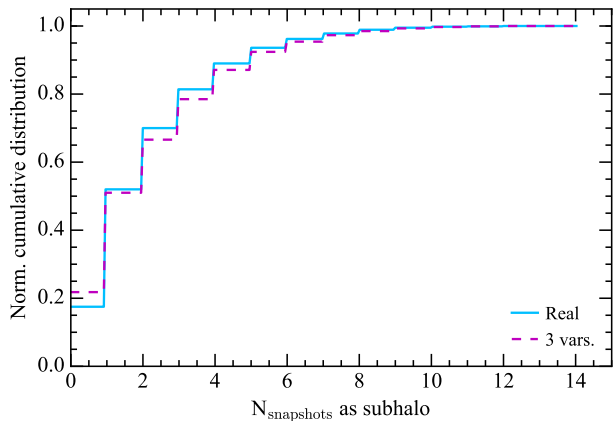


Figure 5. Normalized cumulative distribution of the number of snapshots that a progenitor spend as a subhalo before the merger takes place.

number a snapshot a progenitor spend as subhalo before merging with another progenitor, for generated and real merger trees. Once again, we compared both distribution with the KS test (see last row of Table 1) and found a very good agreement.

As final remark, note that there are unavailable tests in Table 1 due to the requirement of additional variables which are out of the scope on the corresponding configuration.

5. Conclusions

We proposed a new framework for halo merger tree generation based on Generative Adversarial Networks that intends to learn the properties of more expensive simulations with a lower computational cost. Our proposal successfully generates the most important merger trees’ properties, namely progenitor masses, distance of the progenitors to those in the main branch and the progenitor type, and inherits the best features of the trees of the EAGLE simulation suite. Despite restricting trees to have six branches in our experiments, our framework can straightforwardly be expanded to trees with a variable number of branches, provided that there is a large number of tree samples in the training dataset.

A final test for our GAN generated merger trees would be to apply them to populate a semi-analytic model of galaxy formation and then evaluate the quality of the resulting galaxies by comparing them with observations and/or hydrodynamically simulated galaxies. Nevertheless, given that the generated halo merger trees reproduce fairly the properties of the training dataset, specially merger trees generated with 3 variables, we expect as good outcomes as those obtained with trees from N -body simulations.

In addition, we have developed for the first time a set of criteria that takes into account up to three merger trees’

properties to compare different samples of trees. We applied these criteria to assess the quality of the learned representation of generated merger trees when compared to those in the training dataset and to measure the effect of training our GAN framework with additional merger trees’ properties. It is worth remarking that this evaluation framework can also be used to compare merger trees of different large volume cosmological simulations and to assess the importance of the set of variables to take into account for a fair merger construction.

Acknowledgements

SR acknowledges support from MINECO/FEDER (Spain) under grant AYA2015-63810-P and from the Australian Research Council. SR also acknowledges funding from the European Union’s Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie grant agreement No. 734374 (LACEGAL-RISE) for a secondment at the Pontificia Universidad Católica de Chile. JG acknowledges support from CONICYT project Basal AFB-170002. This project has also received funding from the São Paulo Research Foundation (FAPESP) under grant No. 2016/19947-6, and from the Brazilian National Council for Scientific and Technological Development (CNPq) under grant No. 307425/2017-7.

References

- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8):1798–1828, Aug 2013. ISSN 0162-8828. doi: 10.1109/TPAMI.2013.50.
- Benson, A. J. and Bower, R. Galaxy formation spanning cosmic history. *Mon. Not. Roy. Astron. Soc.*, 405(3):1573–1623, Jul 2010. doi: 10.1111/j.1365-2966.2010.16592.x.
- Bond, J. R., Cole, S., Efstathiou, G., and Kaiser, N. Excursion set mass functions for hierarchical gaussian fluctuations. *Astrophys. J.*, 379:440, 1991. doi: 10.1086/170520.
- Bower, R. G., Benson, A. J., Malbon, R., Helly, J. C., Frenk, C. S., Baugh, C. M., Cole, S., and Lacey, C. G. Breaking the hierarchy of galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 370:645–655, August 2006. doi: 10.1111/j.1365-2966.2006.10519.x.
- Carlberg, R. G., Couchman, H. M. P., and Thomas, P. A. Cosmological velocity bias. *Astrophys. J.*, 352:L29–L32, April 1990. doi: 10.1086/185686.
- Cole, S. Modeling Galaxy Formation in Evolving Dark Matter Halos. *Astrophys. J.*, 367:45, Jan 1991. doi: 10.1086/169600.

- Cole, S., Aragon-Salamanca, A., Frenk, C. S., Navarro, J. F., and Zepf, S. E. A recipe for galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 271:781–806, Dec 1994. doi: 10.1093/mnras/271.4.781.
- Cole, S., Lacey, C. G., Baugh, C. M., and Frenk, C. S. Hierarchical galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 319:168–204, November 2000. doi: 10.1046/j.1365-8711.2000.03879.x.
- Crain, R. A., Schaye, J., Bower, R. G., Furlong, M., Schaller, M., Theuns, T., Dalla Vecchia, C., Frenk, C. S., McCarthy, I. G., Helly, J. C., Jenkins, A., Rosas-Guevara, Y. M., White, S. D. M., and Trayford, J. W. The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations. *Mon. Not. Roy. Astron. Soc.*, 450(2):1937–1961, Jun 2015. doi: 10.1093/mnras/stv725.
- Croton, D. J., Springel, V., White, S. D. M., De Lucia, G., Frenk, C. S., Gao, L., Jenkins, A., Kauffmann, G., Navarro, J. F., and Yoshida, N. The many lives of active galactic nuclei: cooling flows, black holes and the luminosities and colours of galaxies. *Mon. Not. Roy. Astron. Soc.*, 365:11–28, January 2006. doi: 10.1111/j.1365-2966.2005.09675.x.
- De Lucia, G., Kauffmann, G., and White, S. D. M. Chemical enrichment of the intracluster and intergalactic medium in a hierarchical galaxy formation model. *Mon. Not. Roy. Astron. Soc.*, 349(3):1101–1116, Apr 2004. doi: 10.1111/j.1365-2966.2004.07584.x.
- Diemand, J., Kuhlen, M., and Madau, P. Early Supersymmetric Cold Dark Matter Substructure. *Astrophys. J.*, 649(1):1–13, Sep 2006. doi: 10.1086/506377.
- Dosovitskiy, A., Springenberg, J. T., and Brox, T. Learning to generate chairs with convolutional neural networks. In *IEEE Inter. Conf. Comput. Vis., Pattern Recog. (CVPR)*, pp. 1538–1546, June 2015. doi: 10.1109/CVPR.2015.7298761.
- Elahi, P. J., Thacker, R. J., and Widrow, L. M. Peaks above the Maxwellian Sea: a new approach to finding substructures in N-body haloes. *Mon. Not. Roy. Astron. Soc.*, 418(1):320–335, Nov 2011. doi: 10.1111/j.1365-2966.2011.19485.x.
- Gonzalez-Perez, V., Lacey, C. G., Baugh, C. M., Lagos, C. D. P., Helly, J., Campbell, D. J. R., and Mitchell, P. D. How sensitive are predicted galaxy luminosities to the choice of stellar population synthesis model? *Mon. Not. Roy. Astron. Soc.*, 439(1):264–283, Mar 2014. doi: 10.1093/mnras/stt2410.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. In *Adv. Neural Inf. Process. Sys. (NeurIPS)*, June 2014.
- Guo, Q., White, S., Boylan-Kolchin, M., De Lucia, G., Kauffmann, G., Lemson, G., Li, C., Springel, V., and Weinmann, S. From dwarf spheroidals to cD galaxies: Simulating the galaxy population in a Λ CDM cosmology. *Mon. Not. Roy. Astron. Soc.*, 413:101–131, May 2011. doi: 10.1111/j.1365-2966.2010.18114.x.
- Han, J., Jing, Y. P., Wang, H., and Wang, W. Resolving subhaloes’ lives with the Hierarchical Bound-Tracing algorithm. *Mon. Not. Roy. Astron. Soc.*, 427(3):2437–2449, Dec 2012. doi: 10.1111/j.1365-2966.2012.22111.x.
- Hatton, S., Devriendt, J. E. G., Ninin, S., Bouchet, F. R., Guiderdoni, B., and Vibert, D. GALICS- I: A hybrid n-body/semi-analytic model of hierarchical galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 343:75–106, July 2003. doi: 10.1046/j.1365-8711.2003.05589.x.
- Jiang, F. and van den Bosch, F. C. Generating merger trees for dark matter haloes: a comparison of methods. *Mon. Not. Roy. Astron. Soc.*, 440:193–207, May 2014. doi: 10.1093/mnras/stu280.
- Katz, N., Hernquist, L., and Weinberg, D. H. Galaxies and gas in a cold dark matter universe. *Astrophys. J.*, 399:L109–L112, November 1992. doi: 10.1086/186619.
- Kauffmann, G. and White, S. D. M. The merging history of dark matter haloes in a hierarchical universe. *Mon. Not. Roy. Astron. Soc.*, 261:921–928, Apr 1993. doi: 10.1093/mnras/261.4.921.
- Kauffmann, G., White, S. D. M., and Guiderdoni, B. The formation and evolution of galaxies within merging dark matter haloes. *Mon. Not. Roy. Astron. Soc.*, 264:201–218, Sep 1993. doi: 10.1093/mnras/264.1.201.
- Kauffmann, G., Colberg, J. M., Diaferio, A., and White, S. D. M. Clustering of galaxies in a hierarchical universe—I. methods and results at $z = 0$. *Mon. Not. Roy. Astron. Soc.*, 303:188–206, February 1999. doi: 10.1046/j.1365-8711.1999.02202.x.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Adv. Neural Inf. Process. Sys. (NeurIPS)*, pp. 1097–1105, 2012.
- Lacey, C. and Silk, J. Tidally triggered galaxy formation. I—evolution of the galaxy luminosity function. *Astrophys. J.*, 381:14, Nov 1991. doi: 10.1086/170625.
- Lacey, C. G., Baugh, C. M., Frenk, C. S., Benson, A. J., Bower, R. G., Cole, S., Gonzalez-Perez, V., Helly, J. C.,

- Lagos, C. D. P., and Mitchell, P. D. A unified multiwavelength model of galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 462(4):3854–3911, Nov 2016. doi: 10.1093/mnras/stw1888.
- Lagos, C. d. P., Tobar, R. J., Robotham, A. S. G., Obreschkow, D., Mitchell, P. D., Power, C., and Elahi, P. J. Shark: introducing an open source, free, and flexible semi-analytic model of galaxy formation. *Mon. Not. Roy. Astron. Soc.*, 481(3):3573–3603, Dec 2018. doi: 10.1093/mnras/sty2440.
- Lee, J., Yi, S. K., Elahi, P. J., Thomas, P. A., Pearce, F. R., Behroozi, P., Han, J., Helly, J., Jung, I., Knebe, A., Mao, Y.-Y., Onions, J., Rodriguez-Gomez, V., Schneider, A., Srisawat, C., and Tweed, D. Sussing merger trees : The impact of halo merger trees on galaxy properties in a semi-analytic model. *Mon. Not. Roy. Astron. Soc.*, 445(4):4197–4210, 2014. doi: 10.1093/mnras/stu2039.
- McAlpine, S., Helly, J. C., Schaller, M., Trayford, J. W., Qu, Y., Furlong, M., Bower, R. G., Crain, R. A., Schaye, J., Theuns, T., Dalla Vecchia, C., Frenk, C. S., McCarthy, I. G., Jenkins, A., Rosas-Guevara, Y., White, S. D. M., Baes, M., Camps, P., and Lemson, G. The EAGLE simulations of galaxy formation: Public release of halo and galaxy catalogues. *Astron. Comput.*, 15:72–89, Apr 2016. doi: 10.1016/j.ascom.2016.02.004.
- Muldrew, S. I., Pearce, F. R., and Power, C. The accuracy of subhalo detection. *Mon. Not. Roy. Astron. Soc.*, 410(4): 2617–2624, Feb 2011. doi: 10.1111/j.1365-2966.2010.17636.x.
- Onions, J., Knebe, A., Pearce, F. R., Muldrew, S. I., Lux, H., Knollmann, S. R., Ascasibar, Y., Behroozi, P., Elahi, P., Han, J., Maciejewski, M., Merchán, M. E., Neyrinck, M., Ruiz, A. N., Sgró, M. A., Springel, V., and Tweed, D. Subhaloes going Notts: the subhalo-finder comparison project. *Mon. Not. Roy. Astron. Soc.*, 423(2):1200–1214, Jun 2012. doi: 10.1111/j.1365-2966.2012.20947.x.
- Onions, J., Ascasibar, Y., Behroozi, P., Casado, J., Elahi, P., Han, J., Knebe, A., Lux, H., Merchán, M. E., Muldrew, S. I., Neyrinck, M., Old, L., Pearce, F. R., Potter, D., Ruiz, A. N., Sgró, M. A., Tweed, D., and Yue, T. Subhaloes gone Notts: spin across subhaloes and finders. *Mon. Not. Roy. Astron. Soc.*, 429(3):2739–2747, Mar 2013. doi: 10.1093/mnras/sts549.
- Planck Collaboration, Ade, P. A. R., Aghanim, N., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Baccigalupi, C., Banday, A. J., and et al. Planck 2013 results. XVI. cosmological parameters. *Astron. Astrophys.*, 571:A16, November 2014. doi: 10.1051/0004-6361/201321591.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv*, art. arXiv:1511.06434, Nov 2015.
- Ricciardelli, E. and Franceschini, A. GECCO: Galaxy Evolution COde—a new semi-analytical model of galaxy formation. *Astron. Astrophys.*, 518:A14, July 2010. doi: 10.1051/0004-6361/200913374.
- Schaye, J., Crain, R. A., Bower, R. G., Furlong, M., Schaller, M., Theuns, T., Dalla Vecchia, C., Frenk, C. S., McCarthy, I. G., Helly, J. C., Jenkins, A., Rosas-Guevara, Y. M., White, S. D. M., Baes, M., Booth, C. M., Camps, P., Navarro, J. F., Qu, Y., Rahmati, A., Sawala, T., Thomas, P. A., and Trayford, J. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Mon. Not. Roy. Astron. Soc.*, 446:521–554, January 2015. doi: 10.1093/mnras/stu2058.
- Somerville, R. S. and Primack, J. R. Semi-analytic modelling of galaxy formation: the local universe. *Mon. Not. Roy. Astron. Soc.*, 310:1087–1110, December 1999. doi: 10.1046/j.1365-8711.1999.03032.x.
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., and Hernquist, L. A semi-analytic model for the co-evolution of galaxies, black holes and active galactic nuclei. *Mon. Not. Roy. Astron. Soc.*, 391:481–506, December 2008. doi: 10.1111/j.1365-2966.2008.13805.x.
- White, S. D. M. and Frenk, C. S. Galaxy formation through hierarchical clustering. *Astrophys. J.*, 379:52–79, September 1991. doi: 10.1086/170483.

A. Details of the GAN Architecture

We selected a convolutional-based architecture, after having evaluated first a fully connected (FC) architecture with different number of layers. Even if the FC-based GAN successfully learned the merger tree structure, it failed to reproduce the correct mass range of the halos in branches different from the main branch.

The discriminator was implemented with convolutional neural networks (CNNs) (Krizhevsky et al., 2012) and the generator with deconvolutional layers. The layout of the encoder was also based on CNNs. We used a combination of column- and row-like filters in each neural network. These filters intend to reproduce the operations within a branch (column-wise filters), and between the branches to produce mergers (row-wise filters). We present the parameters for the discriminator, encoder, and generator in Table A.1. We tested different activation functions, where the Exponential Linear Unit (ELU) yielded the best results. The performance of the GAN was optimized when losses were calculated with cross

Table A.1. Layout and parameters of the (top) discriminator, (middle) encoder, and (bottom) generator networks, where n_{var} denotes the number of variables, k is the kernel structure, and s the number of strides.

Layer	Parameters	Output shape
Input		$(29, 6, n_{\text{var}})$
Conv2D	$k:(1, 3) s:1$	$(29, 6, 16)$
ELU		$(29, 6, 16)$
Conv2D	$k:(1, 3) s:1$	$(29, 6, 32)$
ELU		$(29, 6, 32)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 64)$
ELU		$(29, 6, 64)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 128)$
ELU		$(29, 6, 128)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 256)$
ELU		$(29, 6, 256)$
Flatten		(44544)
FC		(1)

Layer	Parameters	Output Shape
Input		$(29, 6, n_{\text{var}})$
Conv2D	$k:(1, 3) s:1$	$(29, 6, 16)$
ELU		$(29, 6, 16)$
Conv2D	$k:(1, 3) s:1$	$(29, 6, 32)$
ELU		$(29, 6, 32)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 64)$
ELU		$(29, 6, 64)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 128)$
ELU		$(29, 6, 128)$
Conv2D	$k:(3, 1) s:1$	$(29, 6, 256)$
ELU		$(29, 6, 256)$
Flatten		(44544)
FC		(100)

Layer	Parameters	Output Shape
Input		(100)
FC		(44544)
ELU		(44544)
Deconv2D	$k:(3, 1) s:1$	$(29, 6, 128)$
ELU		$(29, 6, 16)$
Deconv2D	$k:(3, 1) s:1$	$(29, 6, 64)$
ELU		$(29, 6, 32)$
Deconv2D	$k:(1, 3) s:1$	$(29, 6, 32)$
ELU		$(29, 6, 64)$
Deconv2D	$k:(1, 3) s:1$	$(29, 6, 16)$
ELU		$(29, 6, 128)$
Deconv2D	$k:(1, 3) s:1$	$(29, 6, n_{\text{var}})$
Sigmoid		$(29, 6, n_{\text{var}})$

B. Examples of Generated Merger Trees

In Fig. B.1, we show two additional examples of merger trees generated with three variables. The tree on the left panel has no sub-branches. While the one on the right shows a sub-branch in the third branch and a late merger (snapshot 28) of two large branches (main branch and branch 6). Note that the main branch hosts more massive progenitors.

C. Distributions of the Distance to the Main Branch

In Fig. C.1, we show the distribution of the progenitor distance to the main branch for several snapshots before the merger takes place for real and merger trees generated with 2 and 3 variables. The distribution is normalized by the maximum number of progenitors in the complete merger tree sample at the snapshot 0. Similar distributions but normalized by the number of progenitors in the merger tree sample at the timestep when the merger event occurs, were used to perform the KS test in Section 4. Note that the three distributions in Fig. C.1 are very similar, the slight differences among them are only noticeable with the KS test.

entropy with logits and a batch of 100 samples. It is worth remarking that the quality of the generated merger trees improved considerably with the introduction of an encoder.

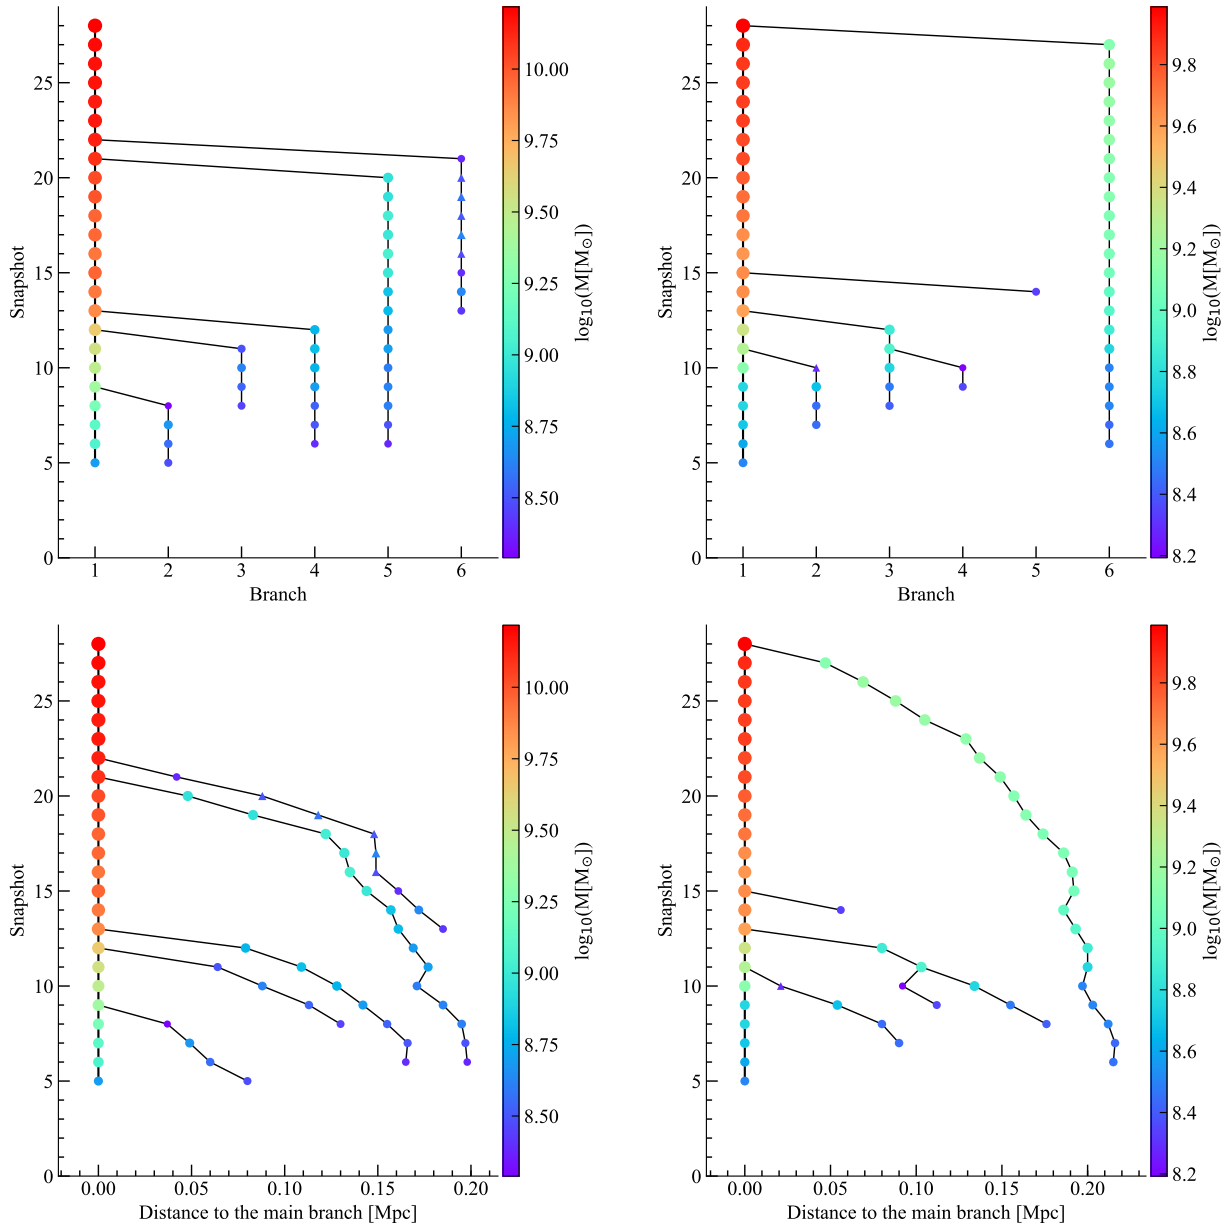


Figure B.1. Left: Merger tree generated with 3 variables, featuring no sub-branches, Right: Another example of a merger tree generated with 3 variables, a sub-branch and a late merger characterize this tree. The color map denotes progenitor masses and the progenitor type is represented by circles (main halos) and triangles (subhalos). Top panels show merger trees in the plane snapshot vs. branch, and bottom panels correspond to the same trees in the plane snapshot vs. distance.

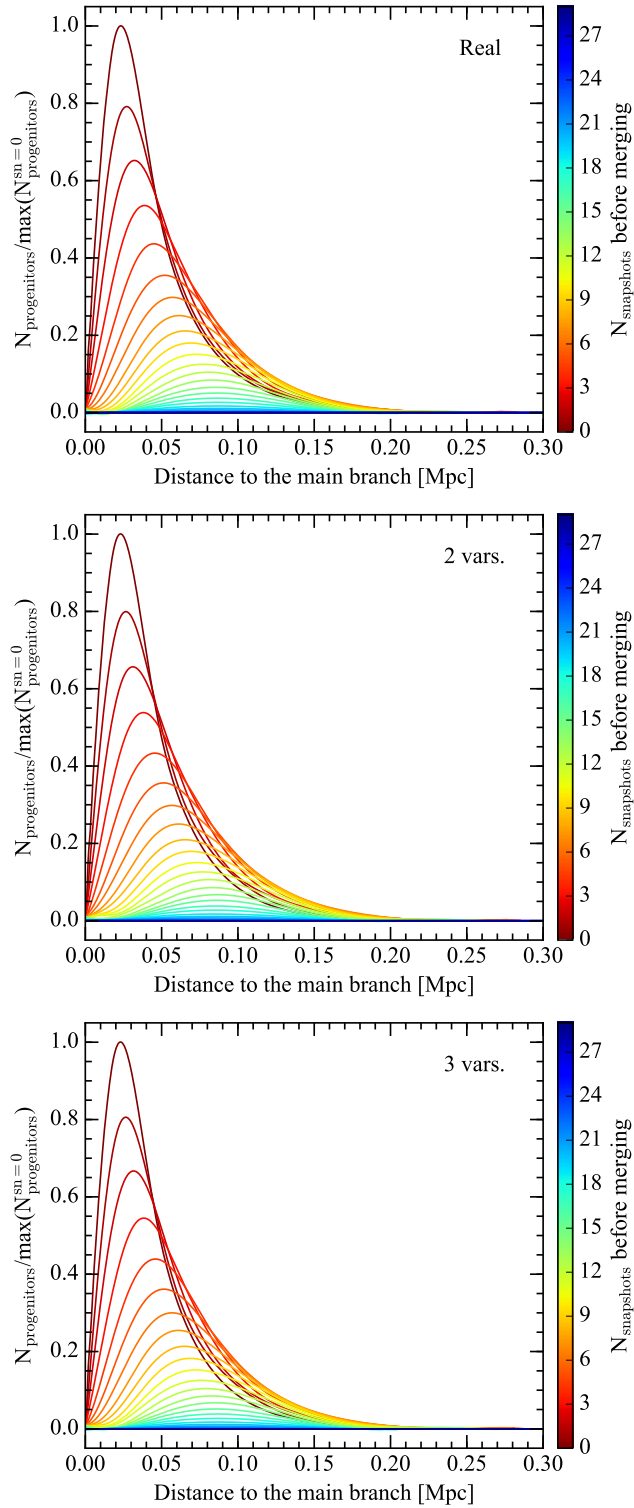


Figure C.1. Distribution of the distance to the main branch for the real trees and those generated with 2 and 3 variables and several snapshots before the merger. The number of progenitors is normalized by the maximum number of these halos at the snapshot 0. The colormap denotes the number of snapshots before the merging event.