
BaryonBridge: Stochastic Interpolant Model for Fast Hydrodynamical Simulations

Benjamin Horowitz^{1,2} Carolina Cuesta-Lazaro^{3,4,5} Omar Yehia^{1,2}

Abstract

Constructing a general-purpose framework for mapping between dark matter simulations and observable hydrodynamical simulation outputs is a long-standing problem in modern astrophysics. In this work, we present a new approach utilizing stochastic interpolants to map between cheap fast particle mesh simulations and baryonic quantities in three dimensions, requiring a total of 7 GPU minutes per 256^3 grid size simulation. Using the CAMELS multifield dataset, we are able to condition our mapping on both cosmological and astrophysical properties. We focus this work on hydrodynamical quantities suitable for Ly- α observables finding excellent agreement up to small spatial scales, $k \sim 10.0$ (h^{-1} Mpc) at $z = 2.0$, for Ly- α flux statistics. Our approach is fully convolutional, allowing training on comparatively small volumes and application to larger volumes, which was tested on TNG50.

1. Introduction

Understanding the large-scale structure of the universe requires modeling both the nonlinear evolution of the dark matter cosmic web and the complex baryonic physics behind observable tracers. Cosmological hydrodynamical simulations have become essential tools for capturing these effects, but their high computational cost limits their use for next-generation surveys that demand large volumes and

high fidelity. Estimating statistical quantities like power spectra and their covariances across wide parameter ranges becomes especially expensive, motivating the development of efficient surrogate models.

A parallel line of development has focused on the reconstruction of initial conditions (ICs) from late-time observations via forward modeling. These approaches, often relying on differentiable structure formation models, enable constrained simulations that match observed data and support field-level inference. While powerful, their effectiveness is currently limited by the lack of efficient differentiable models that incorporate baryonic physics. Bridging this gap is critical for enabling more realistic reconstructions that account for feedback processes and their impact on observables.

Recently, deep learning-based surrogate models have emerged to accelerate simulation workflows. These methods typically either (1) augment low-cost dark matter simulations with baryonic effects (Harrington et al., 2022; Horowitz et al., 2022), or (2) generate full hydrodynamical outputs directly (Tröster et al., 2019; Bernardini et al., 2025).

Our work builds on these by introducing; (1) A differentiable and efficient simulation backbone using `fastPM` (Feng et al., 2016), enabling integration into inverse modeling frameworks (Seljak et al., 2017; Li et al., 2022; Modi et al., 2018; Horowitz et al., 2019; 2021), (2) a fully convolutional architecture that generalizes to arbitrary volume sizes while training on smaller boxes (Harrington et al., 2022; Horowitz et al., 2022) (3) conditional modeling on cosmological and astrophysical parameters, crucial for capturing the effects of baryonic feedback and reducing uncertainties in cosmological inferences.

Unlike unconditional generative models (Zamudio-Fernandez et al., 2019; Hassan & Andrianomena, 2023; Andrianomena et al., 2024), our conditional framework leverages dark matter inputs to ensure generalization and independence across realizations—features necessary for covariance estimation and field-level reconstruction tasks.

*Equal contribution ¹Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan ²Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan ³Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA ⁴The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Massachusetts Institute of Technology, Cambridge, MA 02139, USA ⁵Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. Correspondence to: Benjamin Horowitz <horowitz.ben@gmail.com>.




Figure 1. We show the input density FastPM generated density field and generated neutral hydrogen (N_{HI}), temperature (T), and inferred Ly- α flux from our BaryonBridge model compared to the simulated true result in the TNG50 simulation.

2. Simulations

2.1. CAMELS and TNG50

CAMELS (Villaescusa-Navarro et al., 2021) is a large-scale simulation effort designed to connect cosmological structure formation with astrophysical feedback processes using both N-body and hydrodynamic simulations. In this work, we focus exclusively on the IllustrisTNG suite of CAMELS, which employs the moving-mesh AREPO code (Springel, 2010; Weinberger et al., 2020) and adopts the same sub-grid physics model as the original IllustrisTNG simulations (Weinberger et al., 2017; Pillepich et al., 2018).

Each simulation in the IllustrisTNG suite follows the coupled evolution of dark matter and baryonic matter in a $(25 h^{-1}\text{Mpc})^3$ periodic volume, using 256^3 particles per component. Within this suite, the *Latin Hypercube (LH)* set consists of 1,000 distinct simulations sampling a wide range of cosmological (Ω_m and σ_8) and astrophysical feedback ($A_{SN,1}$, $A_{SN,2}$, $A_{AGN,1}$, $A_{AGN,2}$) parameters using a Latin hypercube design. This enables broad coverage of the parameter space with relatively few samples.

We use the CAMELS Multifield Dataset (CMD) (Villaescusa-Navarro et al., 2022), which provides 3D grids derived from the CAMELS simulation outputs. Specifically, we specialize to the 3D volumes from the IllustrisTNG-LH set, leveraging its diversity in cosmological and feedback parameters to train and validate our model. For this work we focus on Ly- α Forest observables, so we are interested in neutral hydrogen density, temperature, and line of sight (LOS) velocity. The grid LOS velocity is not available in the CMD dataset, so we construct it from available particle data.

To test our model on larger simulation boxes, we also use the TNG50 simulation (Nelson et al., 2019). Like the CAMELS boxes, it uses AREPO hydrodynamics code along with a TNG subgrid physics model. However, it has ~ 3 times the volume ($35 h^{-1}\text{Mpc}$) at significantly higher particle density of 2160^3 . This change of both simulation volume and resolution will provide a robust test for the generalization of our model.

2.2. Particle Mesh Simulations

For our conditioning field, we use particle mesh simulations with matched initial conditions of the full hydrodynamical simulations. We use a JAX implementation (Lanzieri et al., 2022) of the FastPM algorithm (Feng et al., 2016). Fast particle mesh algorithms have become a standard tool for field level inference since they are straightforward to take analytical derivatives of output quantities with respect to initial conditions, this enables derivative based optimization/sampling methods for initial condition inference. (Seljak et al., 2017; Modi et al., 2018; Horowitz et al., 2019; 2021)

For our particle mesh simulations, we take the particle/velocity initial conditions given in the CAMELS simulation set and use them as ICs for our particle mesh code. We down-sample the particles uniformly to 256^3 , and use a force resolution $B = 2$ to provide more accurate diffuse large scale structure necessary for Ly- α statistics. We use a CIC painting scheme to readout the particle masses, total velocities, and line of sight velocities on a 128^3 grid to match those in the CAMELS multifield set. We use the matter density and total velocity as inputs to our interpolant model.




Figure 2. Comparison of redshift space Ly- α flux for both the original TNG50 simulation and that mapped via BaryonBridge code. We also compare to an analytical approximation given via the Fluctuating Gunn-Peterson Approximation (FGPA) on the original dark matter field.

2.3. Mapping to Ly- α Forest

For our example observable, we focus on Ly- α forest since it is a field-level quantity which has high dependance on both astrophysics (Kooistra et al., 2022) and cosmology (Lukić et al., 2014). The expression for optical depth, τ , in real space, as given in (Lukić et al., 2014), is:

$$\tau_\nu \equiv \int n_{HI} \sigma_\nu dr. \quad (1)$$

In this expression, ν denotes the frequency, σ_ν is the cross-section for the Ly- α transition, and dr represents an element of proper distance. When adopting a Doppler-broadened line profile—characterized by a Gaussian shape—the expression becomes:

$$\tau_\nu = \frac{\pi e^2}{m_e c} f_{12} \int \frac{n_{HI}}{\Delta \nu_D} \frac{\exp \left[- \left(\frac{\nu - \nu_0}{\Delta \nu_D} \right)^2 \right]}{\sqrt{\pi}} dr, \quad (2)$$

Here, e is the elementary charge, m_e is the electron mass, and c is the speed of light. The quantity ν_0 refers to the central frequency of the transition, f_{12} is the oscillator strength for the Ly- α resonance, and $\Delta \nu_D = (b/c)\nu_0$ is the Doppler width, with the Doppler parameter defined as $b = \sqrt{2k_B T/m_H}$. We use a JAX version of the differentiable GPU-accelerated code THALAS (Ding et al., 2024) to calculate these quantities from simulations.

3. Stochastic Interpolant Models

We employ Stochastic Interpolants, as formulated in (Albergo et al., 2023; Chen et al., 2024) and first applied in the cosmological context in (Cuesta-Lazaro et al., 2024), to model the conditional distribution $p(\delta_{\text{baryons}}|\delta_{\text{darkmatter}}, \mathcal{C}, \mathcal{A})$. Given a dark matter

density field, $\delta_{\text{darkmatter}}$, and a set of cosmological and astrophysical parameters $\mathcal{C}, \mathcal{A}$, (in our case, $\Omega_m, \sigma_8, A_{SN1}, A_{SN2}, A_{AGN1}, A_{AGN2}$) we seek to generate a probabilistic ensemble of possible baryon distributions δ_{baryons} . A stochastic interpolant I_s can be constructed as:

$$X_s = \alpha_s x_0 + \beta_s x_1 + \sigma_s W_s, \quad (3)$$

where W_s is a Wiener process that can be sampled as $W_s = \sqrt{s}z$ with $z \sim \mathcal{N}(0, I)$. The interpolant maps a point mass measure at x_0 to the conditional distribution $p(x_s|x_0)$ as s varies from 0 to 1. In our application, $x_0 = \delta_{\text{darkmatter}}$ and $x_1 = \delta_{\text{baryons}}$.

The interpolant satisfies boundary conditions $\alpha_0 = \beta_1 = 1$ and $\alpha_1 = \beta_0 = \sigma_0 = 0$. For our implementation, we choose $\alpha_s = 1 - s$, $\beta_s = s^2$, and $\sigma_s = 1 - s$.

The conditional distribution defined by the interpolant, $p(I_s|x_0)$, corresponds to the law of a solution to a stochastic differential equation (SDE) that can be used as a generative model (Chen et al., 2024):

$$dX_s = b_s(X_s, x_0)ds + \sigma_s dW_s, \quad X_{s=0} = x_0, \quad (4)$$

where the drift term b_s is learned by parameterizing a neural network $\hat{b}_s(x, x_0, s)$ and minimizing the objective:

$$L_b[\hat{b}_s] = \int_0^1 ds \mathbb{E} \left[|\hat{b}_s(I_s, x_0, s) - R_s|^2 \right]. \quad (5)$$

Here, R_s is determined by the interpolant as $R_s = \dot{\alpha}_s x_0 + \dot{\beta}_s x_1 + \dot{\sigma}_s W_s$, where the dots represent derivatives with respect to s .

Unlike conditional diffusion models, like the ones used for mapping dark matter fields to galaxies (Bourdin et al., 2024), stochastic interpolants leverage a base distribution x_0 that is already structurally similar to the target distribution, simplifying the mapping that the neural network needs to learn.

We implement the drift term using a fully convolutional 3D U-Net architecture. The fully convolutional nature allows us to train the model on small-volume hydrodynamical simulations (e.g., CAMELS) while performing inference on larger volumes. The interpolant time parameter s is encoded together with the cosmological and astrophysical parameters and incorporated into the network input alongside X_s .

4. Results

We have trained our network, described in Sec. 3 to map from the fast particle mesh simulations to the CAMELS simulations described in Sec. 2. We then map our output baryon fields to the Ly- α flux statistics described in Sec. 2.3. For redshift space distortions, we use the dark matter line of sight velocities from the particle mesh without machine learning augmentation.

4.1. Inference over larger volumes: TNG50

For a robust test of the utility of our trained model, we run a particle mesh simulation with matched initial conditions to that of the TNG50 simulation and calculate a binned density and total velocity field. We apply our trained BaryonBridge on this volume. From the output N_{HI} and T fields, combined with the line of sight velocity from the particle mesh dark matter simulation, we calculate the Ly- α forest as described in Sec 2.3. We show our qualitative performance in a slice in Figure 1, finding excellent visual agreement, including predictions of thermal shocks which were difficult to recover in past machine-learning based approaches (Harrington et al., 2022), although with some loss of sharpness on the edges. We show the calculated Ly- α power spectra in Figure 2, finding excellent quantitative agreement with less than 3 % error up to a $k \sim 8 h \text{ Mpc}^{-1}$.

In total, our model requires only 7 GPU-minutes for inference, one minute on four GPUs for particle mesh and three minutes on one GPU for a model sample, and 96 GPU-hours for training, compared to TNG50’s 130 million CPU core hours.

For an analytical comparison, we use the Fluctuating-Gunn Peterson Approximation for the temperature and baryon density field as described in Harrington et al. (2022). This method goes beyond the “standard” FGPA method used (e.g. that used in Horowitz et al. (2019); Kooistra et al. (2022)) by predicting a temperature and density field with an exponential mapping, and then passing through an equation of state code to calculate the neutral hydrogen fraction and line of sight integration code.

4.2. Parameter variations: CAMELS Boxes

We also test our trained BaryonBridge model on validation boxes in the original IllustrisTNG-LH set held back from training. This set spans a range of cosmological and astrophysical properties, allowing us to validate our conditional mapping. Across our conditioning set, we find high accuracy in our output quantities of N_{HI} and T in terms of quantity power spectra and cross correlation coefficient. In particular, we find $< 10\%$ error in the power spectra across the validation set up to a $k \sim 10 h \text{ Mpc}^{-1}$ as well as a correlation coefficient > 0.9 up to a $k \sim 4 h \text{ Mpc}^{-1}$ for both of these fields.

We select a few extreme values, shown in Table 4.2, to further analyze in terms of the downstream task of Ly- α analysis. We show the Ly- α power spectra in Fig 3, finding excellent agreement, with less than 5% error up to a $k = 10 h \text{ Mpc}^{-1}$ across the parameter space. The worst performing validation case, Parameters C, is quite extreme in both cosmological and astrophysical parameters, but still has acceptable error properties. We have also calculated the

	Ω_m	σ_8	A_{S1}	A_{S2}	A_{A1}	A_{A2}	$\frac{\Delta(F)}{\langle F \rangle}$
A	0.46	0.87	3.81	0.70	1.50	1.78	0.2%
B	0.22	0.79	0.82	0.27	0.55	1.11	0.2%
C	0.14	0.68	2.34	2.83	0.79	0.54	0.3%

Table 1. Validation parameter sets used to demonstrate our conditional model. We also show the relative error on the inferred mean Ly- α flux.




Figure 3. Power spectra statistics corresponding to the models in Tab. 4.2. The shaded region shows the variance of independent samples drawn from the same initial conditions; i.e. a model prediction of uncertainty in the underlying mapping.

cross correlation coefficient, finding $r > 0.9$ up to $k \sim 5.0 h \text{ Mpc}^{-1}$ across all validation parameter sets.

5. Discussion and Conclusion

In this work, we have demonstrated a fast and accurate way to map from relatively low-resolution dark matter simulations to full hydrodynamical quantities in the case of Ly- α observables. We have trained our methods on relatively small $(25 h^{-1} \text{ Mpc})^3$ boxes from the CAMELS set and tested them on the larger TNG50 box, finding excellent performance on downstream tasks like Ly- α power spectra analysis. This method can be used not only for fast generation of mock catalogs for next generation surveys, but could also play a critical role in dynamical forward modeling field level approaches.

Remaining variance between BaryonBridge and the CAMELS results could in part be caused by poor resolution in low-mass regions from the underlying JaxPM simulation and full dark matter simulations. These difference could in part be alleviated via (conditional) refinement methods or neural augmentations on the particle mesh methods (Lanzieri et al., 2022; Payot et al., 2023).

In this work, we have focused on neutral hydrogen and temperature as they are required quantities for Ly- α flux statistics. However, in preliminary tests we have found no fundamental restriction on target fields available in the

CAMELS multifield dataset. In addition, thematically similar techniques have been used to map to highly non-linear quantities like stellar mass (Ono et al., 2024). Currently, memory limitations and training-time expediency require some specialization to only a few target fields at once.

We calculate the Ly- α statistics from the predicted gas temperature and density fields rather than directly inferring the transmitted flux for several key reasons. First, predicting the underlying physical fields enables us to study correlations with other observables. Second, this approach transforms the problem into a local mapping between the dark matter distribution and gas properties, avoiding the need to model the complex long-range correlations that arise when targeting the Ly- α field directly in redshift space. Finally, the causal relationship between baryonic feedback processes and the observed Ly- α absorption is more naturally captured by modeling their separate effects on the temperature and density fields, which then combine to determine the Lyman alpha field.

It is notable that all steps in our analysis framework, JaxPM particle mesh outputs to BaryonBridge baryon quantities to THALAS Ly- α flux are differentiable allowing an end-to-end pipeline for explicit field level inference (e.g. Horowitz et al. (2019); Porqueres et al. (2019)). We will further explore this application in future works.

References

- Albergo, M. S., Boffi, N. M., and Vanden-Eijnden, E. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL <https://arxiv.org/abs/2303.08797>.
- Andrianomena, S., Hassan, S., and Villaescusa-Navarro, F. Cosmological multifield emulator. *arXiv e-prints*, art. arXiv:2402.10997, February 2024. doi: 10.48550/arXiv.2402.10997.
- Bernardini, M., Feldmann, R., Gensior, J., Anglés-Alcázar, D., Bassini, L., Bieri, R., Cenci, E., Tortora, L., and Faucher-Giguère, C.-A. EMBER-2: emulating baryons from dark matter across cosmic time with deep modulation networks. , 538(2):1201–1215, April 2025. doi: 10.1093/mnras/staf341.
- Bourdin, A., Legin, R., Ho, M., Adam, A., Hezaveh, Y., and Perreault-Levasseur, L. Inpainting galaxy counts onto n-body simulations over multiple cosmologies and astrophysics, 2024. URL <https://arxiv.org/abs/2408.00839>.
- Chen, Y., Goldstein, M., Hua, M., Albergo, M. S., Boffi, N. M., and Vanden-Eijnden, E. Probabilistic forecasting with stochastic interpolants and föllmer processes, 2024. URL <https://arxiv.org/abs/2403.13724>.
- Cuesta-Lazaro, C., Bayer, A. E., Albergo, M. S., Mishra-Sharma, S., Modi, C., and Eisenstein, D. J. Joint cosmological parameter inference and initial condition reconstruction with stochastic interpolants. *NeurIPS 2024 Workshop: Machine Learning and the Physical Sciences*, 2024.
- Ding, J., Horowitz, B., and Lukić, Z. TensorFlow Hydrodynamics Analysis for Ly- α Simulations. *Astronomy and Computing*, 48:100858, July 2024. doi: 10.1016/j.ascom.2024.100858.
- Feng, Y., Chu, M.-Y., Seljak, U., and McDonald, P. FASTPM: a new scheme for fast simulations of dark matter and haloes. , 463(3):2273–2286, December 2016. doi: 10.1093/mnras/stw2123.
- Harrington, P., Mustafa, M., Dornfest, M., Horowitz, B., and Lukić, Z. Fast, High-fidelity Ly α Forests with Convolutional Neural Networks. , 929(2):160, April 2022. doi: 10.3847/1538-4357/ac5faa.
- Hassan, S. and Andrianomena, S. HIDM: Emulating Large Scale HI Maps using Score-based Diffusion Models. *arXiv e-prints*, art. arXiv:2311.00833, November 2023. doi: 10.48550/arXiv.2311.00833.
- Horowitz, B., Lee, K.-G., White, M., Krolewski, A., and Ata, M. TARDIS. I. A Constrained Reconstruction Approach to Modeling the $z \sim 2.5$ Cosmic Web Probed by Ly α Forest Tomography. , 887(1):61, December 2019. doi: 10.3847/1538-4357/ab4d4c.
- Horowitz, B., Zhang, B., Lee, K.-G., and Kooistra, R. TARDIS. II. Synergistic Density Reconstruction from Ly α Forest and Spectroscopic Galaxy Surveys with Applications to Protoclusters and the Cosmic Web. , 906(2):110, January 2021. doi: 10.3847/1538-4357/abca35.
- Horowitz, B., Dornfest, M., Lukić, Z., and Harrington, P. HYPHY: Deep Generative Conditional Posterior Mapping of Hydrodynamical Physics. , 941(1):42, December 2022. doi: 10.3847/1538-4357/ac9ea7.
- Kooistra, R., Lee, K.-G., and Horowitz, B. Constraining the Fluctuating Gunn-Peterson Approximation using Ly α Forest Tomography at $z = 2$. , 938(2):123, October 2022. doi: 10.3847/1538-4357/ac92e8.
- Lanzieri, D., Lanusse, F., and Starck, J.-L. Hybrid Physical-Neural ODEs for Fast N-body Simulations. In *Machine Learning for Astrophysics*, pp. 60, July 2022. doi: 10.48550/arXiv.2207.05509.
- Li, Y., Lu, L., Modi, C., Jamieson, D., Zhang, Y., Feng, Y., Zhou, W., Pok Kwan, N., Lanusse, F., and Greengard, L. pmwd: A Differentiable Cosmological Particle-Mesh N-body Library. *arXiv e-prints*, art. arXiv:2211.09958, November 2022. doi: 10.48550/arXiv.2211.09958.

- Lukić, Z., Stark, C. W., Nugent, P., White, M., Meiksin, A. A., and Almgren, A. The Lyman alpha forest in optically thin hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 446(4): 3697–3724, 12 2014. ISSN 0035-8711. doi: 10.1093/mnras/stu2377. URL <https://doi.org/10.1093/mnras/stu2377>.
- Modi, C., Feng, Y., and Seljak, U. Cosmological reconstruction from galaxy light: neural network based light-matter connection. , 2018(10):028, October 2018. doi: 10.1088/1475-7516/2018/10/028.
- Nelson, D., Pillepich, A., Springel, V., Pakmor, R., Weinberger, R., Genel, S., Torrey, P., Vogelsberger, M., Marinacci, F., and Hernquist, L. First results from the TNG50 simulation: galactic outflows driven by supernovae and black hole feedback. , 490(3):3234–3261, December 2019. doi: 10.1093/mnras/stz2306.
- Ono, V., Park, C. F., Mudur, N., Ni, Y., Cuesta-Lazaro, C., and Villaescusa-Navarro, F. Debiasing with Diffusion: Probabilistic Reconstruction of Dark Matter Fields from Galaxies with CAMELS. , 970(2):174, August 2024. doi: 10.3847/1538-4357/ad5957.
- Payot, N., Lemos, P., Perreault-Levasseur, L., Cuesta-Lazaro, C., Modi, C., and Hezaveh, Y. Learning an Effective Evolution Equation for Particle-Mesh Simulations Across Cosmologies. *arXiv e-prints*, art. arXiv:2311.18017, November 2023. doi: 10.48550/arXiv.2311.18017.
- Pillepich, A., Springel, V., Nelson, D., Genel, S., Naiman, J., Pakmor, R., Hernquist, L., Torrey, P., Vogelsberger, M., Weinberger, R., and Marinacci, F. Simulating galaxy formation with the IllustrisTNG model. , 473(3):4077–4106, January 2018. doi: 10.1093/mnras/stx2656.
- Porqueres, N., Jasche, J., Lavaux, G., and Enßlin, T. Inferring high-redshift large-scale structure dynamics from the Lyman- α forest. , 630:A151, October 2019. doi: 10.1051/0004-6361/201936245.
- Seljak, U., Aslanyan, G., Feng, Y., and Modi, C. Towards optimal extraction of cosmological information from nonlinear data. , 2017(12):009, December 2017. doi: 10.1088/1475-7516/2017/12/009.
- Springel, V. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. , 401(2):791–851, January 2010. doi: 10.1111/j.1365-2966.2009.15715.x.
- Tröster, T., Ferguson, C., Harnois-Déraps, J., and McCarthy, I. G. Painting with baryons: augmenting N-body simulations with gas using deep generative models. , 487(1): L24–L29, July 2019. doi: 10.1093/mnras/slz075.
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., Spergel, D. N., Somerville, R. S., Dave, R., Pillepich, A., Hernquist, L., Nelson, D., Torrey, P., Narayanan, D., Li, Y., Philcox, O., La Torre, V., Maria Delgado, A., Ho, S., Hassan, S., Burkhart, B., Wadekar, D., Battaglia, N., Contardo, G., and Bryan, G. L. The CAMELS Project: Cosmology and Astrophysics with Machine-learning Simulations. , 915(1):71, July 2021. doi: 10.3847/1538-4357/abf7ba.
- Villaescusa-Navarro, F., Genel, S., Anglés-Alcázar, D., Thiele, L., Dave, R., Narayanan, D., Nicola, A., Li, Y., Villanueva-Domingo, P., Wandelt, B., Spergel, D. N., Somerville, R. S., Zorrilla Matilla, J. M., Mohammad, F. G., Hassan, S., Shao, H., Wadekar, D., Eickenberg, M., Wong, K. W. K., Contardo, G., Jo, Y., Moser, E., Lau, E. T., Machado Poletti Valle, L. F., Perez, L. A., Nagai, D., Battaglia, N., and Vogelsberger, M. The CAMELS Multifield Data Set: Learning the Universe’s Fundamental Parameters with Artificial Intelligence. , 259(2):61, April 2022. doi: 10.3847/1538-4365/ac5ab0.
- Weinberger, R., Springel, V., Hernquist, L., Pillepich, A., Marinacci, F., Pakmor, R., Nelson, D., Genel, S., Vogelsberger, M., Naiman, J., and Torrey, P. Simulating galaxy formation with black hole driven thermal and kinetic feedback. , 465(3):3291–3308, March 2017. doi: 10.1093/mnras/stw2944.
- Weinberger, R., Springel, V., and Pakmor, R. The AREPO Public Code Release. , 248(2):32, June 2020. doi: 10.3847/1538-4365/ab908c.
- Zamudio-Fernandez, J., Okan, A., Villaescusa-Navarro, F., Bilaloglu, S., Derin Cengiz, A., He, S., Perreault Levasseur, L., and Ho, S. HIGAN: Cosmic Neutral Hydrogen with Generative Adversarial Networks. *arXiv e-prints*, art. arXiv:1904.12846, April 2019. doi: 10.48550/arXiv.1904.12846.