Continuous Representations of Baryonic Feedback for Robust Inference from Multiple Simulation Suites

Ming-Shau Liu
Cavendish Laboratory
University of Cambridge
msl63@cam.ac.uk

Carolina Cuesta-Lazaro
Flatiron Institute
Institute for Advanced Studies
ccuesta-lazaro@flatironinstitute.org

Abstract

Accurate modeling of baryonic physics remains a major challenge for precision cosmology due to our incomplete understanding of complex subgrid processes, like star formation and feedback from supernovae and active galactic nuclei below 10 Mpc scales. This uncertainty leads to different hydrodynamical simulation suites to implement fundamentally different prescriptions for these unresolved physics. Current simulation-based inference approaches rely therefore on discrete sets of simulators, each encoding specific physical assumptions, making it difficult to robustly quantify theoretical uncertainties and learn about the underlying physics from observations. We introduce a machine learning framework that learns continuous representations of baryonic feedback across multiple simulation suites, to enable interpolation between different physical implementations while providing robust uncertainty quantification. Our approach addresses the key challenge of marginalizing over theoretical uncertainties represented by various simulators while simultaneously constraining the underlying baryonic physics from observations. We frame this as learning a shared continuous latent representation of the physics implemented across different simulators, allowing us to both marginalize over and constrain a continuous baryonic parameter space. Using the CAMELS simulation suite, we demonstrate our method on several baryonic fields including stellar mass, gas density, temperature, and pressure fields. This framework provides a path toward more robust cosmological inference by properly accounting for theoretical uncertainties in baryonic modeling while extracting maximum information about the underlying physical processes from current and future surveys.

1 Introduction

Understanding how matter is distributed over small scales in the Universe is crucial for extracting cosmological information from current and future surveys, yet it remains one of the most challenging aspects of precision cosmology [4, 23]. While the large-scale distribution of dark matter is well-described by gravity alone, complex baryonic physics—including star formation, and feedback from supernovae and active galactic nuclei—significantly alters the distribution of matter below ~ 10 Mpc [20]. Currently, we rely on hydrodynamical simulations for accurate predictions of baryonic physics, [5–7, 14, 16] but these must make unresolved assumptions about subgrid processes, with different simulation suites implementing fundamentally different prescriptions for feedback mechanisms. It remains unclear which approaches best capture the true underlying physics, and likely none of the existing suites fully captures the reality, making empirical approaches to modeling baryons interesting for robust inference.

In this work, we introduce a machine learning framework to learn a continuous representation of baryonic physics from multiple simulation suites, enabling both interpolation between different

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

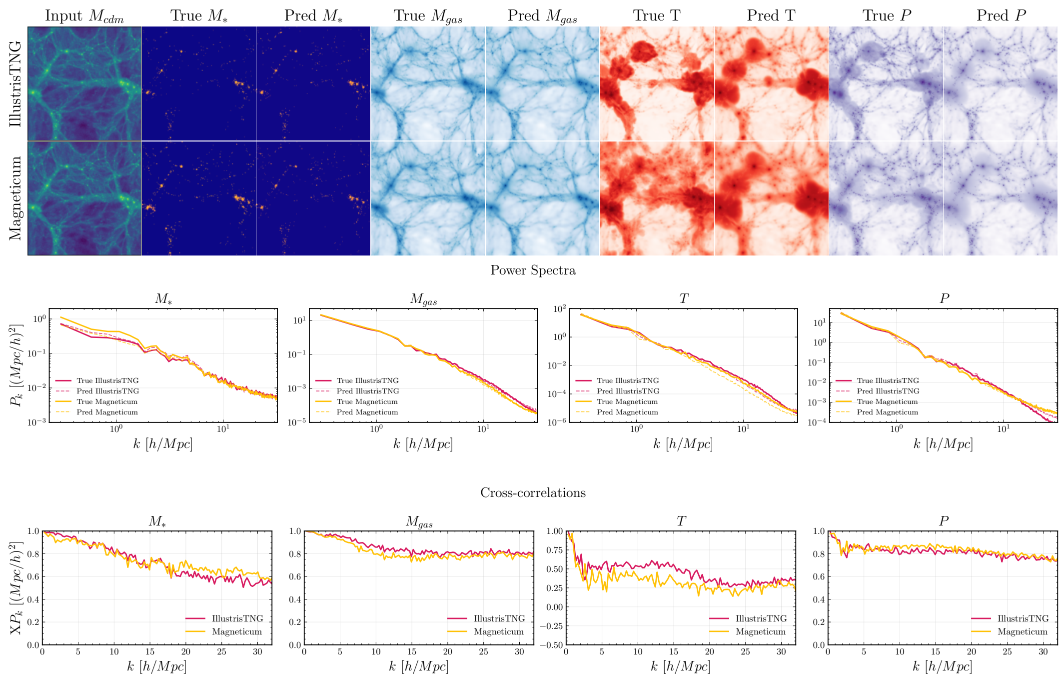

Figure 1: Reconstruction of baryonic maps from dark matter fields on two different simulation suites starting from the same initial conditions, IllustrisTNG (used for training) and Magneticum (used for testing). On the top panel, we show the dark matter fields used for conditioning the model and both true and generated baryonic maps for $\{M_*, \delta_{\text{gas}}, T, P\}$. On the middle panel, we compare true and predicted power spectra for each field and simulation suite. On the bottom panel, we show the cross-correlation coefficient as a function of scale computed between the true and reconstructed maps across simulations. In general, we find the reconstruction accuracy on the test Magneticum suite to be comparable to that of the in distribution IllustrisTNG.

baryonic implementations and robust uncertainty quantification. This addresses the more general challenge [9, 11] in simulation-based inference of working with a discrete set of simulators, where each parametrises different physical assumptions. Here, we assume one wishes to also marginalise over a more general parameter space that would allow us to interpolate in between these assumptions.

We frame this problem as learning a shared continuous representation [2] of the underlying physics implemented across different simulators. Our goal is two-fold: we wish to robustly marginalize over the theoretical uncertainties represented by the various simulators (i.e., marginalize over z_{baryons}) while simultaneously learning about the underlying physical processes from observations of the real Universe (i.e., constrain z_{baryons} via observations). We demonstrate this approach using the CAMELS simulation suite, focusing on key baryonic observables including stellar mass, gas density, temperature, and pressure fields.

2 Learning a representation for baryonic physics

We aim to learn a low dimensional representation for baryonic physics, z_{baryons}, that allows us to constrain the mapping between the dark matter distribution, δ_{dm}, whose evolution is dominated by gravity, and the baryons δ_{baryons}, whose evolution is impacted by baryonic feedback,

$$p(\delta_{\text{baryons}} | \delta_{\text{dm}}, \mathcal{C}, z_{\text{baryons}}) \tag{1}$$

where δ_{baryons} is a diverse set of baryonic fields, and $\mathcal{C}$ are the cosmological ΛCDM parameters that determine the composition of the Universe, Ω_m and σ_8. In this paper, the δ_{baryons} we use includes the distribution of stars, $M_\star$, the gas density distribution, δ_{gas}, the gas temperature distribution, T, and the gas pressure distribution, P.

Architecture. To model the probability distribution in Equation (1), we use flow matching [1, 10, 12, 13], a continuous normalizing flow that learns velocity fields $v_\theta(x,t)$ that transform the source distribution (dark matter density $\delta_{\rm dm}$) to target distributions (baryonic fields , $\{M_*, \delta_{\rm gas}, T, P\}$) along time-dependent paths $x_t = (1-t)x_0 + tx_1$. The loss we minimize is $\mathcal{L} = \mathbb{E}_{t,x_0,x_1}[\|v_\theta(x_t, t, \delta_{\rm dm}, \mathcal{C}, z_{\rm baryon}) - (x_1 - x_0)\|^2]$, enabling both target fields reconstruction and new data generation.

We jointly learn the parameters of a ResNet encoder [8] that compresses multi-channel baryonic inputs $\{M_*, \delta_{\rm gas}, T, P\}$ into an 8-dimensional bottleneck representation $z_{\rm baryon}$ that captures shared and simulation-specific variations in baryonic feedback prescriptions.

We parametrize the learned velocity field, $v_\theta(x,t)$, with a UNet [18], featuring an encoder-decoder structure and skip connections for multi-scale feature learning, modified with Feature-wise Linear Modulation (FiLM) [17] layers in each block that condition the network through affine transformations $\gamma(\mathbf{c}) \odot \mathbf{h} + \beta(\mathbf{c})$, where $\mathbf{c}$ represents a unified conditioning embedding. This conditioning vector concatenates four complementary sources: (i) a time embedding from the flow-matching timestep t, (ii) a parameter embedding from cosmological parameters (Ω_m, σ_8) processed through a small MLP, (iii) a spatial embedding from a CNN applied to the input dark matter density field, and (iv) a baryonic physics embedding $z_{\rm baryon}$ extracted from the ResNet bottleneck applied to stacked baryonic reference maps.

CAMELS Simulation. The CAMELS Multifield Dataset [26] comprises hydrodynamical simulations with varied cosmological parameters $\mathcal{C} = \{\Omega_m, \sigma_8\}$ and distinct baryonic feedback implementations across simulation suites. We partition the dataset using four simulations for training–Astrid [3, 15, 16], IllustrisTNG [14], SIMBA [6], and EAGLE [5, 19]–while reserving Magneticum [7] as a held-out test set to evaluate cross-simulation generalization. Each simulation generates projected 2D density fields over $(25\ h^{-1}{\rm Mpc})^2$ areas at redshift $z = 0$, with spatial resolution of 256^2 pixels. The only varying cosmological parameters are: Ω_m (the matter density parameter), ranging from 0.1 to 0.5 and σ_8 (the amplitude of matter fluctuation), ranging from 0.6 to 1.0. All other cosmological parameters are held constant across the simulations. The training is completed using the Latin hypercube (LH) and Sobol sequence (SBn) sets where available. Each simulation suites implements distinct baryonic feedbacks and subgrid physics that manifests in baryonic fields captured by our learned $z_{\rm baryon}$ embedding. See Ref. [27] for more details on the simulation suites.

3 Results

3.1 Reconstructing the baryonic fields

In Fig. 1, we demonstrate that the model successfully reconstructs four key baryonic observables—stellar mass (M_*), gas density ($\delta_{\rm gas}$), temperature (T), and pressure (P)—from an 8-dimensional baryonic latent representation. Fig. 1 presents the pixel-wise reconstruction of the different fields together with reconstruction quality metrics, the power spectrum and cross-correlation between true and predicted fields, calculated using `Pylians` [25]. Both statistics were evaluated on in-distribution IllustrisTNG samples and a held-out Magneticum test set, simulated from the same initial conditions. The model exhibits strong performance at large scales ($k < 1\,h\,{\rm Mpc}^{-1}$), where cross-correlations exceed > 0.9 for all baryonic channels, but shows degraded fidelity at small scales ($k > 5\,h\,{\rm Mpc}^{-1}$) where nonlinear baryonic processes create complex, localized features that deviate significantly from the dark matter distribution. To compute the power spectrum and cross correlations, we generate five independent samples from the learned flow and ensemble average to reduce stochastic variance.

In particular, we find similar cross-correlation coefficients for the IllustrisTNG and Magneticum suites. This generalization to the unseen Magneticum physics demonstrates the robustness of the learned $z_{\rm baryon}$ embedding, though with modestly reduced performance relative to training simulations. Stellar mass and gas density reconstructions maintain high fidelity due to their strong correlations with underlying dark matter structure, while temperature field exhibit larger reconstruction errors reflecting its weaker coupling to gravitational potentials and stronger dependence on simulation-specific feedback prescriptions. This performance hierarchy aligns with the physical expectation that

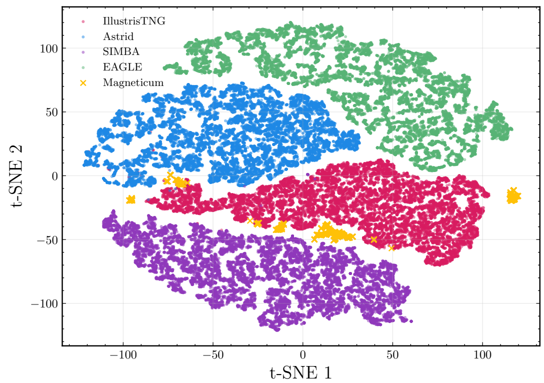

Figure 2: Two-dimensional t-SNE projection of the eight-dimensional Latent space z_{baryon}, colored by hydrodynamic simulation suite. Note that IllustrisTNG, EAGLE, SIMBA and Astrid were used for training, whereas Magneticum is a held-out test suite.

observables more sensitive to subgrid modeling choices than purely gravitational tracers are more difficult to predict from dark matter alone.

3.2 The structure of the latent space

We show a two dimensional t-SNE projection [24] of the 8-dimensional z_{baryon} embedding in Fig. 2. The learned manifold clusters samples from the same simulation suites while positioning the held-out Magneticum mostly within samples of EAGLE and Astrid. The latent geometry encodes simulation relationships: IllustrisTNG and Astrid exhibit nearby embeddings consistent with similar feedback implementations, while SIMBA occupies a distant manifold region reflecting its distinct prescriptions of AGN feedback. This clustering demonstrates successful disentanglement of simulation-specific systematics from shared features.

The interpolation property is important for robust downstream inference by enabling marginalization over simulation-specific features. The learned embedding enables the possibility of marginalization over simulation space, isolating cosmological signals from simulation systematics. This marginalization framework addresses a fundamental challenge in simulation-based inference where systematic uncertainties from subgrid physics can dominate statistical errors, providing a pathway toward cosmological parameter estimation robust to hydrodynamical modeling assumptions.

4 Conclusions and Future Work

We present a flow-matching framework that learns the probability distribution $p(\delta_{\mathrm{baryons}}|\delta_{\mathrm{dm}}, \mathcal{C}, z_{\mathrm{baryons}})$ to reconstruct four key baryonic observables—stellar mass, gas density, temperature, and pressure—from dark matter density fields. We jointly learn a ResNet encoder to extract an 8-dimensional z_{baryons} embedding that captures variations across hydrodynamical simulation codes in a continuous space. This space, in principle, enables marginalization $p(\delta_{\mathrm{baryons}}|\delta_{\mathrm{dm}}, \mathcal{C}) = \int p(\delta_{\mathrm{baryons}}|\delta_{\mathrm{dm}}, \mathcal{C}, z_{\mathrm{baryons}}) p(z_{\mathrm{baryons}})\, \mathrm{d}z_{\mathrm{baryons}}$ over simulation systematics, given a prior over $p(z_{\mathrm{baryons}})$, potentially yielding robust inference invariant to hydrodynamical simulator choice and addressing challenging biases introduced by subgrid physics disparities across simulations.

Future extensions will incorporate an application to galaxy clustering and Sunyaev-Zel'dovich observables [21, 22] to enhance constraining power through multi-probe analysis. Additionally, we plan to develop a gravitational physics embedding disentangled from z_{baryons} to detect gravitational anomalies while maintaining robustness to baryonic uncertainties—a capability particularly valuable for constraining alternative dark matter models by isolating gravitational signatures from confounding baryonic feedback effects. Our framework demonstrates that flow-matching can effectively separate cosmological signals from simulation-specific systematics, providing a pathway toward unbiased inference in the era of precision cosmology.

5 Acknowledgment

This work is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, http://iaifi.org/). The research reported in this paper was done on the Tufts University High Performance Computing Cluster (`https://it.tufts.edu/high-performance-computing`). We are grateful to the NSF AI Institute for Artificial Intelligence and Fundamental Interactions and the AstroAI Institute at the Center for Astrophysics | Harvard & Smithsonian for providing a working space and hosting this project. M.-S. L. was supported by Institute of Physics Early Career Researchers fund and the Cambridge Philosophical Society.

References

[1] M. S. Albergo and E. Vanden-Eijnden. Building normalizing flows with stochastic interpolants, 2023.

[2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives, 2014.

[3] S. Bird, Y. Ni, T. Di Matteo, R. Croft, Y. Feng, and N. Chen. The ASTRID simulation: galaxy formation and reionization. , 512(3):3703–3716, May 2022.

[4] N. E. Chisari, A. J. Mead, S. Joudaki, P. G. Ferreira, A. Schneider, J. Mohr, T. Tröster, D. Alonso, I. G. McCarthy, S. Martin-Alvarez, J. Devriendt, A. Slyz, and M. P. van Daalen. Modelling baryonic feedback for survey cosmology. *The Open Journal of Astrophysics*, 2(1), June 2019.

[5] R. A. Crain, J. Schaye, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, and J. W. Trayford. The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations. , 450(2):1937–1961, June 2015.

[6] R. Davé, D. Anglés-Alcázar, D. Narayanan, Q. Li, M. H. Rafieferantsoa, and S. Appleby. simba: Cosmological simulations with black hole growth and feedback. *Monthly Notices of the Royal Astronomical Society*, 486(2):2827–2849, Apr. 2019.

[7] K. Dolag, R.-S. Remus, L. M. Valenzuela, L. C. Kimmig, B. Seidel, S. Fortune, J. Stoiber, A. Ivleva, T. Hoffmann, V. Biffi, I. Marini, P. Popesso, and S. Vladutescu-Zopp. Encyclopedia Magneticum: Scaling Relations from Cosmic Dawn to Present Day. *arXiv e-prints*, page arXiv:2504.01061, Apr. 2025.

[8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

[9] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, V. Begy, and G. Louppe. A trust crisis in simulation-based inference? your posterior approximations can be unfaithful, 2022.

[10] S. Kannan, T. Qiu, C. Cuesta-Lazaro, and H. Jeong. Cosmoflow: Scale-aware representation learning for cosmology with flow matching, 2025.

[11] P. Lemos, A. Coogan, Y. Hezaveh, and L. Perreault-Levasseur. Sampling-based accuracy testing of posterior estimators for general inference, 2023.

[12] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling, 2023.

[13] X. Liu, C. Gong, and Q. Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022.

[14] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel, M. Vogelsberger, R. Pakmor, F. Marinacci, R. Weinberger, L. Kelley, M. Lovell, B. Diemer, and L. Hernquist. The illustristng simulations: Public data release, 2021.

[15] Y. Ni, N. Chen, Y. Zhou, M. Park, Y. Yang, T. DiMatteo, S. Bird, and R. Croft. The astrid simulation: Evolution of black holes and galaxies to z=0.5 and different evolution pathways for galaxy quenching, 2024.

[16] Y. Ni, T. Di Matteo, S. Bird, R. Croft, Y. Feng, N. Chen, M. Tremmel, C. DeGraf, and Y. Li. The ASTRID simulation: the evolution of supermassive black holes. , 513(1):670–692, June 2022.

[17] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. C. Courville. Film: Visual reasoning with a general conditioning layer. In *AAAI*, 2018.

[18] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing.

[19] J. Schaye, R. A. Crain, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, M. Baes, C. M. Booth, P. Camps, J. F. Navarro, Y. Qu, A. Rahmati, T. Sawala, P. A. Thomas, and J. Trayford. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. , 446(1):521–554, Jan. 2015.

[20] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci, and J. Naiman. First results from the illustristng simulations: matter and galaxy clustering. *Monthly Notices of the Royal Astronomical Society*, 475(1):676–698, Dec. 2017.

[21] R. A. Sunyaev and Y. B. Zeldovich. Small-Scale Fluctuations of Relic Radiation. , 7(1):3–19, Apr. 1970.

[22] R. A. Sunyaev and Y. B. Zeldovich. The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173, Nov. 1972.

[23] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia. The effects of galaxy formation on the matter power spectrum: a challenge for precision cosmology: Galaxy formation and the matter power spectrum. *Monthly Notices of the Royal Astronomical Society*, 415(4):3649–3665, July 2011.

[24] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.

[25] F. Villaescusa-Navarro. Pylians: Python libraries for the analysis of numerical simulations. Astrophysics Source Code Library, record ascl:1811.008, Nov. 2018.

[26] F. Villaescusa-Navarro, S. Genel, D. Anglés-Alcázar, L. A. Perez, P. Villanueva-Domingo, D. Wadekar, H. Shao, F. G. Mohammad, S. Hassan, E. Moser, E. T. Lau, L. F. Machado Poletti Valle, A. Nicola, L. Thiele, Y. Jo, O. H. E. Philcox, B. D. Oppenheimer, M. Tillman, C. Hahn, N. Kaushal, A. Pisani, M. Gebhardt, A. M. Delgado, J. Caliendo, C. Kreisch, K. W. K. Wong, W. R. Coulton, M. Eickenberg, G. Parimbelli, Y. Ni, U. P. Steinwandel, V. La Torre, R. Dave, N. Battaglia, D. Nagai, D. N. Spergel, L. Hernquist, B. Burkhart, D. Narayanan, B. Wandelt, R. S. Somerville, G. L. Bryan, M. Viel, Y. Li, V. Irsic, K. Kraljic, F. Marinacci, and M. Vogelsberger. The camels project: Public data release. *The Astrophysical Journal Supplement Series*, 265(2):54, apr 2023.

[27] F. Villaescusa-Navarro, S. Genel, D. Anglés-Alcázar, L. Thiele, R. Dave, D. Narayanan, A. Nicola, Y. Li, P. Villanueva-Domingo, B. Wandelt, D. N. Spergel, R. S. Somerville, J. M. Zorrilla Matilla, F. G. Mohammad, S. Hassan, H. Shao, D. Wadekar, M. Eickenberg, K. W. K. Wong, G. Contardo, Y. Jo, E. Moser, E. T. Lau, L. F. Machado Poletti Valle, L. A. Perez, D. Nagai, N. Battaglia, and M. Vogelsberger. The camels multifield data set: Learning the universe's fundamental parameters with artificial intelligence. *The Astrophysical Journal Supplement Series*, 259(2):61, apr 2022.