
Mean-Field Simulation-Based Inference for Cosmological Initial Conditions

Oleg Savchenko

GRAPPA Institute
University of Amsterdam, The Netherlands
o.savchenko@uva.nl

Florian List

Department of Astrophysics
University of Vienna, Austria
florian.list@univie.ac.at

Guillermo Franco Abellán

GRAPPA Institute
University of Amsterdam, The Netherlands
g.francoabellan@uva.nl

Noemi Anau Montel

GRAPPA Institute
University of Amsterdam, The Netherlands
n.anaumontel@uva.nl

Christoph Weniger

GRAPPA Institute
University of Amsterdam, The Netherlands
c.weniger@uva.nl

Abstract

Reconstructing cosmological initial conditions (ICs) from late-time observations is a difficult task, which relies on the use of computationally expensive simulators alongside sophisticated statistical methods to navigate multi-million dimensional parameter spaces. We present a simple method for Bayesian field reconstruction based on modeling the posterior distribution of the initial matter density field to be diagonal Gaussian in Fourier space, with its covariance and the mean estimator being the trainable parts of the algorithm. Training and sampling are extremely fast (*training*: ~ 1 h on a GPU, *sampling*: $\lesssim 3$ s for 1000 samples at resolution 128^3), and our method supports industry-standard (non-differentiable) N -body simulators. We verify the fidelity of the obtained IC samples in terms of summary statistics.

1 Introduction

The cosmic initial conditions – minuscule matter density perturbations seeded 14 billion years ago – gradually evolved under gravity into the intricately woven filamentary structures of galaxies observed in the Universe around us (e.g. [1, 2]). Modeling this process of cosmic structure formation by means of N -body simulations is, in principle, relatively straightforward (albeit computationally demanding) – at least when neglecting all non-gravitational effects (see [3, 4] for recent reviews).

An important task in cosmology addresses the inverse problem: *given (simulated or observed) cosmological present-day observations, can the cosmic ICs be reconstructed?* Given a late-time matter density field without velocity information, this problem is ill-posed due to the non-linear nature of gravity, and the IC reconstruction is not unique [5–7]. Probabilistic (Bayesian) reconstruction techniques have successfully been employed [8, 9]; however, they rely on Hamiltonian Monte Carlo, which is slow due to the large number of parameters (= voxels) $10^6 \lesssim d \lesssim 10^7$ and the complexity of the forward model (which needs to be differentiable such that its gradient can be computed).

Recently, simulation-based inference (SBI, e.g. [10]) has gained significant traction in the astrophysics and cosmology communities, driven by the recent advancements in machine learning (e.g. [11–27]).

SBI allows tackling inverse problems in situations where a forward model is available which describes the mapping from input parameters (here: cosmic ICs) to observations (here: late-time matter densities), whereas explicitly formulating a probabilistic model may be hard or even impossible. Equipped with N -body simulations as an accurate forward model, SBI ideally lends itself for the task of probabilistic cosmic IC reconstruction.

Our contribution. We present a fast (both in terms of training and sampling), simple, yet effective SBI-based technique for drawing cosmic IC posterior samples constrained on late-time density fields. The key components are 1) a trainable maximum a-posteriori (MAP) estimator $\hat{\mu}_\theta$ and 2) a trainable likelihood covariance matrix $(\mathbf{Q}_\theta^L)^{-1}$. Specifically, we model the posterior as Gaussian with a diagonal covariance matrix in Fourier space, which enables near-instantaneous sampling.

Related work. Methods for inferring cosmological ICs from late-time observations can generally be subdivided into deterministic and probabilistic frameworks. The former class reconstructs a single ‘best-fit’ IC field and relies either on traditional (analytic and numerical) techniques [5, 6, 28–33] or, more recently, deep learning approaches (mostly involving U-Nets) [34–38], recurrent inference machines [39], and variational self-boosted sampling [40]. While probabilistic reconstruction techniques have traditionally employed Bayesian modeling paired with Hamiltonian Monte Carlo sampling [8, 9], score-based generative models have recently emerged as an accurate and much faster contender [41–43], and see [44] for a first exploratory study based on SBI alongside autoregressive modeling.

2 Methodology

Simulation data. The simplicity of our method presented below does not demand the simulator to be differentiable, which, in combination with the fact that a relatively small training set is sufficient, does not require us to rely on approximation schemes in our forward model and enables the use of full industry-standard cosmological N -body codes. As our training data, we employ 2000 pairs of initial (at redshift $z = 127$) and present-day ($z = 0$) matter overdensity fields from the Quijote N -body simulations suite [45] with varying random phases of ICs and fixed fiducial Planck cosmology [46]. These simulations evolve 512^3 collisionless particles with the Gadget-III TreePM code [47] in a periodic cubic volume of $(1 \text{ Gpc}/h)^3$. The overdensity fields are computed by interpolating the particles onto a grid with resolution of $d = 128^3$, in total amounting to a $\sim 10^6$ -dimensional parameter space for the inference.

Statistical model. The mentioned ill-posedness of the cosmological ICs reconstruction makes it natural to formulate the problem in the probabilistic Bayesian setting:

$$p(\mathbf{z}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{z})}{p(\mathbf{x})}p(\mathbf{z}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the observed final overdensity field, and $\mathbf{z} \in \mathbb{R}^d$ are its ICs, whose distribution for a given observation \mathbf{x}_{obs} we aim to infer. From observations of the Cosmic Microwave Background [46], we know that the early-universe matter overdensity fluctuations had small amplitudes and to a very high precision can be described as a Gaussian random field; therefore, we can write down the prior as $p(\mathbf{z}) \propto \exp\{-\frac{1}{2}\mathbf{z}^T \mathbf{Q}^P \mathbf{z}\}$, where the precision matrix \mathbf{Q}^P is diagonal in Fourier space: $\mathbf{Q}^P = \mathcal{F}^\dagger \mathbf{D}^P \mathcal{F}$, with $\mathbf{D}^P = 1/P(k)$ determined by the linear matter power spectrum $P(k)$ at $z = 127$, $k := |\mathbf{k}|$ being the wave vector \mathbf{k} modulus, and \mathcal{F} denoting the Fourier transform¹.

A crucial component of any Bayesian cosmological reconstruction framework, which determines its descriptive power, is the complexity level of the likelihood model. In SBI, the likelihood is implicitly encoded in the forward model, for which reason it is typically not explicitly prescribed, but rather accessed from the forward model by a neural network. In this work, however, we adopt a relatively simple approach, in which we explicitly expand the log likelihood up to quadratic terms w.r.t. the ICs \mathbf{z} , with a diagonal likelihood precision matrix in Fourier space. We found this to be sufficient to yield tight and statistically consistent posterior distributions within the selected scale range, while

¹In practice, we make use of the closely related discrete Hartley transform rather than the discrete Fourier transform because it has the advantage that it maps real values to real values, while still being unitary and diagonalizing the precision matrix [48].

vastly reducing the complexity and computational cost of our method. On sufficiently large scales, this level of modeling extracts the full information content in the data, as all Fourier modes evolve independently in the linear regime. On smaller scales, our approach neglects information from the cross-correlations between Fourier modes (however, this information is leveraged when estimating the MAP, see below), which should typically lead to conservative, rather than overconfident, posterior estimates. Thus, the likelihood we use is defined as follows:

$$p(\mathbf{x}|\mathbf{z}) \propto \exp\left\{-\frac{1}{2}(\mathbf{z} - \hat{\mathbf{z}}_{\theta}(\mathbf{x}))^T \mathbf{Q}_{\theta}^L (\mathbf{z} - \hat{\mathbf{z}}_{\theta}(\mathbf{x}))\right\}, \quad (2)$$

where $\mathbf{Q}_{\theta}^L = \mathcal{F}^{\dagger} D_{\theta}^L \mathcal{F}$ with D_{θ}^L being diagonal, $\hat{\mathbf{z}}_{\theta}(\mathbf{x})$ is the \mathbf{x} -dependent maximum likelihood estimator (MLE), and the subscript θ denotes quantities which depend on the set of trainable parameters θ .

Making use of Bayes' theorem (1) with the likelihood defined in (2), after rearranging the terms one obtains the following form of the normalized posterior distribution:

$$p(\mathbf{z}|\mathbf{x}) = \frac{\exp\left\{-\frac{1}{2}(\mathbf{z} - \hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}))^T (\mathbf{Q}^P + \mathbf{Q}_{\theta}^L) (\mathbf{z} - \hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}))\right\}}{\sqrt{(2\pi)^d \det(\mathbf{Q}^P + \mathbf{Q}_{\theta}^L)^{-1}}}, \quad (3)$$

where $\hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x})$ is the MAP estimator related to the MLE by $\hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}) = (\mathbf{Q}^P + \mathbf{Q}_{\theta}^L)^{-1} \mathbf{Q}_{\theta}^L \hat{\mathbf{z}}_{\theta}(\mathbf{x})$. Motivated by the linearity of cosmic structure growth on large scales, we model the MAP estimator $\hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x})$ as a combination of two terms, where the large-scale regime is captured by a simple multiplicative scaling factor $\alpha_{\theta}(\mathbf{k})$, while the complex small-scale evolution is described by a 3D U-Net type [49] neural network in the following way:

$$\hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}) = \mathcal{F}^{\dagger} \left\{ \alpha_{\theta}(\mathbf{k}) \odot [\mathcal{F}\{\mathbf{x}\} + \sigma_{>k_{\Lambda}}(\mathcal{F}\{\text{U-Net}_{\theta}(\mathbf{x})\})] \right\}. \quad (4)$$

Here, $\sigma_{>k_{\Lambda}}$ denotes a sigmoidal high-pass filter centered at the cut-off scale k_{Λ} , which we conservatively chose as $k_{\Lambda} = 0.03 h/\text{Mpc}$, $\alpha_{\theta}(\mathbf{k})$ is a trainable multiplier for each wave vector \mathbf{k} , and \odot is the Hadamard (elementwise) product. We found that including $\alpha_{\theta}(\mathbf{k})$ slightly improves the results while maintaining essentially the same speed of training. We implemented the U-Net from the `map2map` package [50], slightly simplifying it by reducing the number of hidden channels from 64 to 16. To facilitate the training, before passing the $z = 127$ fields to the network, we found it beneficial to divide them by the linear growth factor $D(z = 127)$, effectively making the network learn the residuals from the linear theory prediction.

Our training objective is to minimize the negative log posterior probability under the ansatz (3):

$$\mathcal{L} = \frac{1}{2} \sum_{i=1}^N \left\{ (\mathbf{z}_i - \hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}_i))^T \mathbf{Q}_{\theta} (\mathbf{z}_i - \hat{\boldsymbol{\mu}}_{\theta}(\mathbf{x}_i)) \right\} - \frac{N}{2} \text{tr} \log \mathbf{Q}_{\theta}, \quad (5)$$

where N is the number of samples in the training set, $\mathbf{Q}_{\theta} = \mathbf{Q}^P + \mathbf{Q}_{\theta}^L$, and we used the relation $\log \det \mathbf{Q}_{\theta} = \text{tr} \log \mathbf{Q}_{\theta}$, also leaving out an unimportant constant additive term. Altogether, the trainable set θ of our model consists of the MAP estimator $\hat{\boldsymbol{\mu}}_{\theta}$ parameters (see Eq. 4) and the diagonal part of the likelihood precision matrix D_{θ}^L (in order to ensure the positivity of the latter, we store the respective square root values as the parameters, and square them at inference). Due to its parallels with approximate models encountered in the contexts of statistical physics (e.g. [51, Chap. 4]) and variational inference [52], we refer to our method as '*mean-field simulation-based inference*'. While our method, which focuses on the first two moments (mean and variance), is conceptually aligned with the SBI framework outlined in [25], which employs a hierarchy of networks to estimate successive moments of the posterior, our approach differs in both the training objective and the learning strategy. Moreover, our approach is closely connected to uncertainty quantification methods that optimize the Gaussian negative log-likelihood loss, as discussed in [53].

Training procedure. We train the model using a batch size of 8 on a single 40 GB Nvidia A100 GPU, starting with an initial learning rate of 10^{-2} and decreasing it by a factor of 10 after each epoch where the validation loss does not decrease. Training continues until no further improvement is observed in the validation loss, which, under these conditions, takes approximately 30 epochs and 1.5 hours to converge. The dataset was randomly split with 80% assigned to training and the remaining 20% reserved for validation. As a wrapper for our training implementation, we use the PyTorch Lightning-based [54] code `swyft` [55, 56].

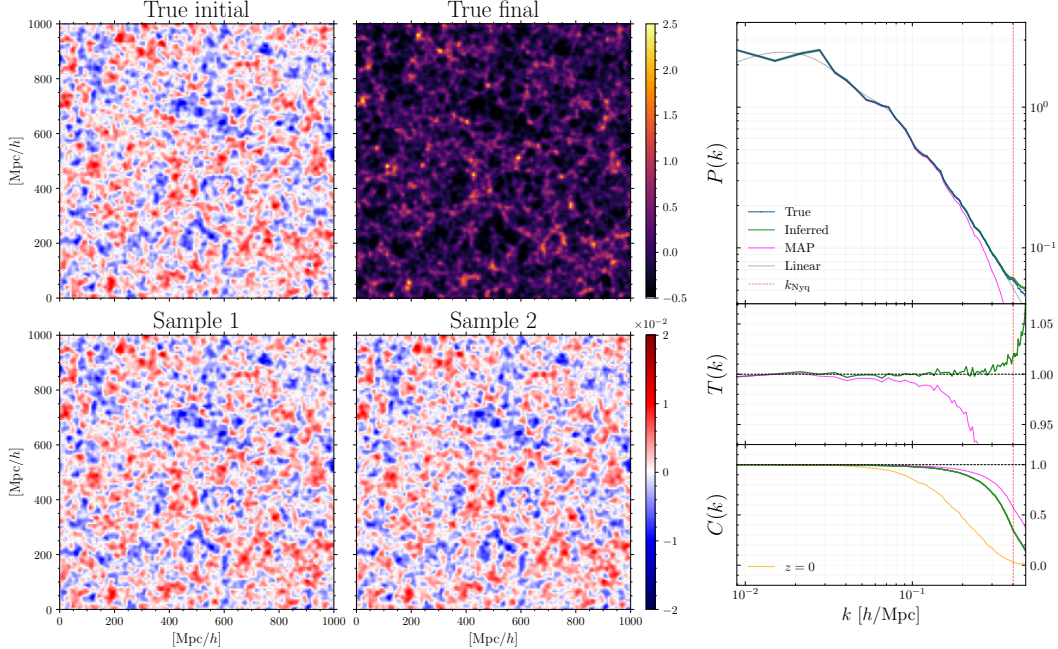


Figure 1: Our method produces posterior IC samples constrained by a given late-time density field (*training*: $\simeq 1.5$ h on a GPU, *sampling*: < 3 s for 1000 samples at resolution 128^3). *Top row*: slices of the initial and final overdensity fields of the target simulation. *Bottom row*: two examples of the generated IC samples. All the shown slices are averaged over the depth of 100 Mpc/h in the third axis direction. *Rightmost column*: power spectrum, transfer function, and cross-correlation of the generated samples and the ground truth. Shaded regions correspond to 1σ and 2σ errors, and the yellow $C(k)$ line corresponds to the cross-correlation between the final and the initial density fields. Note that the MAP estimate, which is the only quantity recovered by most other approaches, loses significant power at small scales and is not fully able to describe the field’s statistical properties.

3 Results

The fact that the prior and the likelihood (and hence also the posterior) are Gaussian and diagonal in the same Fourier basis allows us to obtain samples from the posterior in a very fast and simple way after the training is done. We demonstrate our reconstruction method for one of the fiducial Quijote simulations $\{z_{\text{truth}}, \mathbf{x}_{\text{obs}}\}$ not contained in the training set (see Fig. 1). Given such an observation of the final density field \mathbf{x}_{obs} , one evaluation of the MAP estimator (4) is required to obtain the mean of the field $\hat{\mu}_{\theta}(\mathbf{x}_{\text{obs}})$ ($\lesssim 1$ s computation time), and having the $\mathbf{Q}_{\theta} = \mathbf{Q}^P + \mathbf{Q}_{\theta}^L$ matrix then allows us to obtain over a thousand posterior samples in a batched fashion within three seconds on a single GPU, which is orders of magnitude faster than existing methods (see App. A for details on MAP estimation and the uncertainty of the modes reconstruction).

To confirm that the generated samples have the correct statistical properties, we compute the power spectrum $P(k)$, transfer function $T(k)$ and cross-correlation $C(k)$ with respect to the ground truth field z_{truth} using the Pylians [57] package (see App. B for definitions). We observe an agreement in the power spectrum with the ground truth at $\lesssim 1 - 2\%$ accuracy for all the scales up to the Nyquist scale $k_{\text{Nyq}} \simeq 0.4$ h/Mpc, and a high ($\gtrsim 50\%$) cross-correlation up to $k \simeq 0.35$ h/Mpc. Furthermore, in order to verify that the distribution of samples correctly covers the true posterior, we perform a Bayesian coverage test for different ranges of k -values (see App. C).

4 Discussion and conclusions

We have presented and implemented a simple yet precise and effective method for Bayesian field-level reconstruction of cosmological initial conditions. The simplicity of the method lies in the fact that the posterior is modeled to be Gaussian with the covariance being diagonal in Fourier space and the mean

determined by a non-linear neural network-based estimator $\hat{\mu}_\theta(\mathbf{x}_{\text{obs}})$. This turns out to be sufficient to produce accurate results in the chosen range of scales and, after training, the estimator $\hat{\mu}_\theta$ and the covariance $(Q_\theta^L)^{-1}$ allow us to produce posterior samples in a remarkably fast way. The method does not require the simulator to be differentiable and is not reliant on approximation schemes.

Our method leaves a lot of promising future directions for investigation. First, we observe that the trained Q_θ^L matrix contains a relatively small amount of scatter and can be approximately described as a function of the wavenumber k , as expected from the physical rotational symmetry (see App. A). This should make it possible to impose the constraint $Q_\theta^L(\mathbf{k}) = Q_\theta^L(k)$ from the start, significantly reducing the number of trainable parameters and strengthening the interpretability of the approach. Furthermore, since our proposed method is SBI-based, it has the ability to marginalize over cosmological parameters or, alternatively, to infer them alongside the random phases of the ICs field. Especially interesting enhancement in this context can be offered by sequential SBI techniques.

It is important to realize the present limitations of our approach. While we found the Gaussian diagonal approximation for the posterior to be sufficient in the chosen range of scales $k \lesssim 0.4 h/\text{Mpc}$, extending it to smaller scales might require a more complex modeling of the Q_θ^L matrix or to abandon the Gaussian approximation altogether. Moreover, we demonstrate our method in the idealized scenario where the observation consists of the full 3D matter overdensity field. In reality, what we observe are galaxies—biased tracers of this field—and additionally, the observations are subject to systematic effects such as redshift space distortions, lightcone effects, survey masks, and noise contamination. While the presented SBI framework is well suited to incorporate all these effects, any realistic application to observational data would require further robustness checks.

References

- [1] Tully R. B., Courtois H., Hoffman Y., Pomarède D., 2014, *Nature*, 513, 71
- [2] Tempel E., Stoica R. S., Martínez V. J., Liivamägi L. J., Castellan G., Saar E., 2014, *Monthly Notices of the Royal Astronomical Society*, 438, 3465
- [3] Angulo R. E., Hahn O., 2022, *Living Reviews in Computational Astrophysics*, 8, 1
- [4] Vogelsberger M., Marinacci F., Torrey P., Puchwein E., 2020, *Nature Reviews Physics*, 2, 42
- [5] Nusser A., Dekel A., 1992, *The Astrophysical Journal*, 391, 443
- [6] Frisch U., Matarrese S., Mohayaee R., Sobolevski A., 2002, *Nature*, 417, 260
- [7] Crocce M., Scoccimarro R., 2006, *Physical Review D*, 73, 063520
- [8] Jasche J., Wandelt B. D., 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 894
- [9] Jasche J., Lavaux G., 2019, *Astronomy & Astrophysics*, 625, A64
- [10] Cranmer K., Brehmer J., Louppe G., 2020, *Proceedings of the National Academy of Sciences*, 117, 30055
- [11] Ho M., et al., 2024, *The Open Journal of Astrophysics*, 7, 001c.120559
- [12] Mishra-Sharma S., Cranmer K., 2022, *Physical Review D*, 105, 063017
- [13] Anau Montel N., Coogan A., Correa C., Karchev K., Weniger C., 2022, *Monthly Notices of the Royal Astronomical Society*, 518, 2746
- [14] Villaescusa-Navarro F., et al., 2021 ([arXiv:2109.10360](https://arxiv.org/abs/2109.10360))
- [15] Brehmer J., Mishra-Sharma S., Hermans J., Louppe G., Cranmer K., 2019, *The Astrophysical Journal*, 886, 49
- [16] Zhao X., Mao Y., Cheng C., Wandelt B. D., 2022, *The Astrophysical Journal*, 926, 151
- [17] Cole A., Miller B. K., Witte S. J., Cai M. X., Grootes M. W., Nattino F., Weniger C., 2022, *Journal of Cosmology and Astroparticle Physics*, 09, 004

- [18] Coogan A., Anau Montel N., Karchev K., Grootes M. W., Nattino F., Weniger C., 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 66
- [19] Alvey J., Gerdes M., Weniger C., 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 3662
- [20] Lemos P., et al., 2024, *Physical Review D*, 109, 083536
- [21] Gatti M., et al., 2024, *Physical Review D*, 109, 063534
- [22] Tucci B., Schmidt F., 2024, *Journal of Cosmology and Astroparticle Physics*, 05, 063
- [23] Abellán G. F., Cañas Herrera G., Martinelli M., Savchenko O., Sciotti D., Weniger C., 2024 ([arXiv:2403.14750](https://arxiv.org/abs/2403.14750))
- [24] von Wietersheim-Kramsta M., Lin K., Tessore N., Joachimi B., Loureiro A., Reischke R., Wright A. H., 2024 ([arXiv:2404.15402](https://arxiv.org/abs/2404.15402))
- [25] Jeffrey N., Wandelt B. D., 2020 ([arXiv:2011.05991](https://arxiv.org/abs/2011.05991))
- [26] Makinen T. L., Charnock T., Alsing J., Wandelt B. D., 2021, *Journal of Cosmology and Astroparticle Physics*, 11, 049
- [27] Modi C., Pandey S., Ho M., Hahn C., Blancard B. R.-S., Wandelt B., 2023 ([arXiv:2309.15071](https://arxiv.org/abs/2309.15071))
- [28] Weinberg D. H., 1992, *Monthly Notices of the Royal Astronomical Society*, 254, 315
- [29] Gramann M., 1993, *The Astrophysical Journal*, 405, 449
- [30] Croft R. A., Gaztanaga E., 1997, *Monthly Notices of the Royal Astronomical Society*, 285, 793
- [31] Brenier Y., Frisch U., Hénon M., Loeper G., Matarrese S., Mohayaee R., Sobolevskii A., 2003, *Monthly Notices of the Royal Astronomical Society*, 346, 501
- [32] Schmittfull M., Baldauf T., Zaldarriaga M., 2017, *Physical Review D*, 96, 023505
- [33] Feng Y., Seljak U., Zaldarriaga M., 2018, *Journal of Cosmology and Astroparticle Physics*, 2018, 043–043
- [34] Jindal V., Liang A., Singh A., Ho S., Jamieson D., 2023 ([arXiv:2303.13056](https://arxiv.org/abs/2303.13056))
- [35] Shallue C. J., Eisenstein D. J., 2023, *Monthly Notices of the Royal Astronomical Society*, 520, 6256
- [36] Flöss T., Meerburg P. D., 2024, *Journal of Cosmology and Astroparticle Physics*, 2024, 031
- [37] Wang Z., Shi F., Yang X., Li Q., Liu Y., Li X., 2024, *Science China Physics, Mechanics, and Astronomy*, 67, 219513
- [38] Doerer L., Jamieson D., Stopyra S., Lavaux G., Leclercq F., Jasche J., 2023 ([arXiv:2312.09271](https://arxiv.org/abs/2312.09271))
- [39] Modi C., Lanusse F., Seljak U., Spergel D. N., Perreault-Levasseur L., 2021 ([arXiv:2104.12864](https://arxiv.org/abs/2104.12864))
- [40] Modi C., Li Y., Blei D., 2023, *Journal of Cosmology and Astroparticle Physics*, 2023, 059
- [41] Legin R., Ho M., Lemos P., Perreault-Levasseur L., Ho S., Hezaveh Y., Wandelt B., 2023 ([arXiv:2304.03788](https://arxiv.org/abs/2304.03788))
- [42] Park C. F., Ono V., Mudur N., Ni Y., Cuesta-Lazaro C., 2023 ([arXiv:2311.08558](https://arxiv.org/abs/2311.08558))
- [43] Ono V., Park C. F., Mudur N., Ni Y., Cuesta-Lazaro C., Villaescusa-Navarro F., 2024, *The Astrophysical Journal*, 970, 174

- [44] List F., Anau Montel N., Weniger C., 2023, in 37th Conference on Neural Information Processing Systems. ([arXiv:2310.19910](https://arxiv.org/abs/2310.19910))
- [45] Villaescusa-Navarro F., et al., 2020, *The Astrophysical Journal Supplement Series*, 250, 2
- [46] Aghanim N., et al., 2020, *Astronomy & Astrophysics*, 641, A6
- [47] Springel V., 2005, *Monthly Notices of the Royal Astronomical Society*, 364, 1105
- [48] Bracewell R. N., 1983, *Journal of the Optical Society of America*, 73, 1832
- [49] Ronneberger O., Fischer P., Brox T., 2015, in Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. pp 234–241, [doi:10.48550/arXiv.1505.04597](https://doi.org/10.48550/arXiv.1505.04597)
- [50] Jamieson D., Li Y., de Oliveira R. A., Villaescusa-Navarro F., Ho S., Spergel D. N., 2023, *The Astrophysical Journal*, 952, 145
- [51] Yeomans J. M., 1992, Statistical mechanics of phase transitions. Clarendon Press, [doi:10.1093/oso/9780198517290.001.0001](https://doi.org/10.1093/oso/9780198517290.001.0001)
- [52] Blei D. M., Kucukelbir A., McAuliffe J. D., 2017, *Journal of the American Statistical Association*, 112, 859
- [53] Nix D., Weigend A., 1994, in Tesauro G., Touretzky D., Leen T., eds, Vol. 7, Advances in Neural Information Processing Systems. MIT Press, https://proceedings.neurips.cc/paper_files/paper/1994/file/061412e4a03c02f9902576ec55ebbe77-Paper.pdf
- [54] Falcon W., The PyTorch Lightning team 2019, [doi:10.5281/zenodo.3828935](https://doi.org/10.5281/zenodo.3828935), <https://github.com/Lightning-AI/lightning>
- [55] Miller B. K., Cole A., Forré P., Louppe G., Weniger C., 2021, in 35th Conference on Neural Information Processing Systems. ([arXiv:2107.01214](https://arxiv.org/abs/2107.01214))
- [56] Miller B. K., Cole A., Weniger C., Nattino F., Ku O., Grootes M. W., 2022, *Journal of Open Source Software*, 7, 4205
- [57] Villaescusa-Navarro F., 2018, Astrophysics Source Code Library, [record ascl:1811.008](https://www.aanda.org/record/ascl:1811.008)




Figure 2: *Top row*: slices of the initial density field z_{truth} of the target simulation and the corresponding MAP estimate $\hat{\mu}_\theta(\mathbf{x}_{\text{obs}})$. *Bottom row left*: diagonal values D_θ of the trained posterior precision matrix $Q_\theta = Q^P + Q_\theta^L$ as a function of the wavenumber k . *Bottom row right*: the standard deviation computed from 1000 samples.

A Appendix: MAP estimation and standard deviation

In Fig. 2 we plot the MAP estimate $\hat{\mu}_\theta(\mathbf{x}_{\text{obs}})$, the standard deviation computed from 1000 samples, and the diagonal values D_θ of the trained posterior precision matrix $Q_\theta = Q^P + Q_\theta^L$ as a function of the wavenumber k . The latter quantity allows estimating the precision in the reconstruction at different scales. As expected, we find that D_θ is larger for the smallest k values (where evolution is linear and easier to reconstruct), and then it sharply drops to a plateau that extends up to k_{Nyq} . In addition, we see that D_θ shows little variation with the direction of k and can be approximately described as a function of k . Hence, knowledge of this function $D_\theta(k)$ allows to transform any IC point estimator into a fast sampler. Because of this property, the obtained samples for a given \mathbf{x}_{obs} have the form of a MAP estimate with a Gaussian random field added on top of it, which causes the standard deviation of samples to be nearly constant (as can be seen in Fig. 2).

In this work, we found that a very simple Fourier-diagonal \mathbf{x} -independent precision matrix is enough to get consistent posterior samples. However, it might be possible to improve our results by also modeling the Q_θ matrix as \mathbf{x} -dependent, just as for the $\hat{\mu}_\theta$ estimator. For other scenarios, one could also consider alternative structures of the Q_θ matrix. As an example, we experimented with the convolutional matrix $Q_\theta^L = C_\theta^\dagger D_\theta^L C_\theta$, where D_θ^L is diagonal and C_θ denotes the convolution operation of a certain kernel size, achieving altogether similar reconstruction results. Instead of convolutions, one can also consider wavelets, scattering transform, etc. In those cases, the prior and the likelihood precision matrices are diagonal in different bases, so sampling is no longer as straightforward and fast, and one needs to resort to techniques such as data augmentation approaches for sampling (see Appendix B of [44]).

B Appendix: Summary statistics

Given a (dimensionless) 3D matter overdensity field $\delta(\mathbf{x}) = \rho(\mathbf{x})/\bar{\rho} - 1$ at some fixed redshift, its power spectrum $P(k)$ is defined as follows

$$\langle \tilde{\delta}(\mathbf{k}) \tilde{\delta}^*(\mathbf{k}') \rangle = (2\pi)^3 P(k) \delta^{(3)}(\mathbf{k} - \mathbf{k}'), \quad (6)$$

where $\tilde{\delta}(\mathbf{k}) = \mathcal{F}\{\delta(\mathbf{x})\}$ denotes the Fourier transform of the field, the brackets $\langle \dots \rangle$ indicate the average over many realizations, and the Dirac delta function follows from translational invariance. The power spectrum is the Fourier transform of the matter field two-point correlation function, and thus describes the density fluctuations of the Universe as a function of scale k .

Transfer function and cross-correlation between the two fields $\delta_a(\mathbf{x})$ and $\delta_b(\mathbf{x})$ are then defined as:

$$T_{ab}(k) = \sqrt{\frac{P_a(k)}{P_b(k)}}; \quad C_{ab}(k) = \frac{P_{ab}(k)}{\sqrt{P_a(k) \times P_b(k)}}, \quad (7)$$

where the cross-power spectrum $P_{ab}(k)$ is defined analogously to (6) from $\langle \tilde{\delta}_a(\mathbf{k}) \tilde{\delta}_b^*(\mathbf{k}') \rangle$. So, if the two fields are identical, $T_{ab}(k)$ and $C_{ab}(k)$ are equal to 1 for all k , and the deviation of these quantities from 1 provides a measure of discrepancy in terms of their amplitudes and phases, respectively, between the two fields at different scales k .

C Appendix: Bayesian coverage test

We perform a coverage test to evaluate whether the IC samples generated by our method are statistically consistent with the true posterior distribution, individually for different k -bins. Specifically, we compute the difference between the Fourier transformed true ICs, $\tilde{z}_{\text{truth}} := \mathcal{F}\{z_{\text{truth}}\}$, and the means of the Fourier transformed posterior samples, $\tilde{z}_{\text{samples}} := \mathbb{E}[\tilde{z}_{\text{samples}}]$, then normalize this difference by the standard deviation of the Fourier transforms of the posterior samples, $\tilde{\sigma}_{\text{samples}} := \sqrt{\mathbb{E}[(\tilde{z}_{\text{samples}})^2] - (\mathbb{E}[\tilde{z}_{\text{samples}}])^2}$, with $z_{\text{samples}} \sim p(z|\mathbf{x})$ as modeled in Eq. (3). Mathematically, this can be expressed via the test statistic

$$\Delta(\mathbf{k}) := \frac{\tilde{z}_{\text{truth}}(\mathbf{k}) - \tilde{z}_{\text{samples}}(\mathbf{k})}{\tilde{\sigma}_{\text{samples}}(\mathbf{k})}. \quad (8)$$

Since our likelihood precision matrix is highly rotationally symmetric (as expected from the isotropy of the Universe), we focus on the radial dimension in this analysis and consider the consistency of our posterior within different spherical shells in k -space.

In Fig. 3, we present histograms of these normalized differences $\Delta(\mathbf{k})$ across various k -ranges. The fact that all histograms closely follow the univariate normal distribution with unit variance and zero mean indicates that the posterior samples accurately capture the first two moments of the true posterior distribution. Note that this is a more challenging test than the one shown in [41, Fig. 5], where the coverage is assessed for all modes at once, since we demonstrate the statistical consistency of our posterior samples *within each individual k -bin*.




Figure 3: *Coverage test*. We plot the histograms of the quantity $\Delta(\mathbf{k})$ defined in (8) for four different spherical shells in k -space. The red lines represent the univariate normal distribution $\mathcal{N}(0, 1)$. The fact that all the histograms closely follow the $\mathcal{N}(0, 1)$ distribution validates our posterior coverage test for different ranges of k .