Particle Monte Carlo Methods for Lattice Field Theory [1]

Frontiers in Probabilistic Inference @ NeurIPS 2025: Sampling Meets Learning

David Yallup

Kavli Institute for Cosmology & Institute of Astronomy, University of Cambridge

Motivation & contributions

- Neural samplers learn flexible *trivialising maps*, highly relevant for efficient Monte Carlo sampling in Lattice Field Theory problems [2, 3].
- Many of the strengths of neural samplers overlap with existing strengths of Particle MC methods [4, 5].
- Careful evaluation is needed to establish what advances are truly brought by neural methods [6].
- How well do lessons learnt from successful application of neural samplers in LFT transfer more broadly to other challenging sampling problems?
- **Key result: Black-box particle Monte Carlo methods can outperform black-box neural samplers.** The regime in which neural samplers are performant in LFT is one where Particle MC methods are already highly effective.

Challenging sampling targets, scalar field theory on a lattice

LFT problems are characterised by strong couplings, multi-modal distributions, and high-dimensional state spaces.

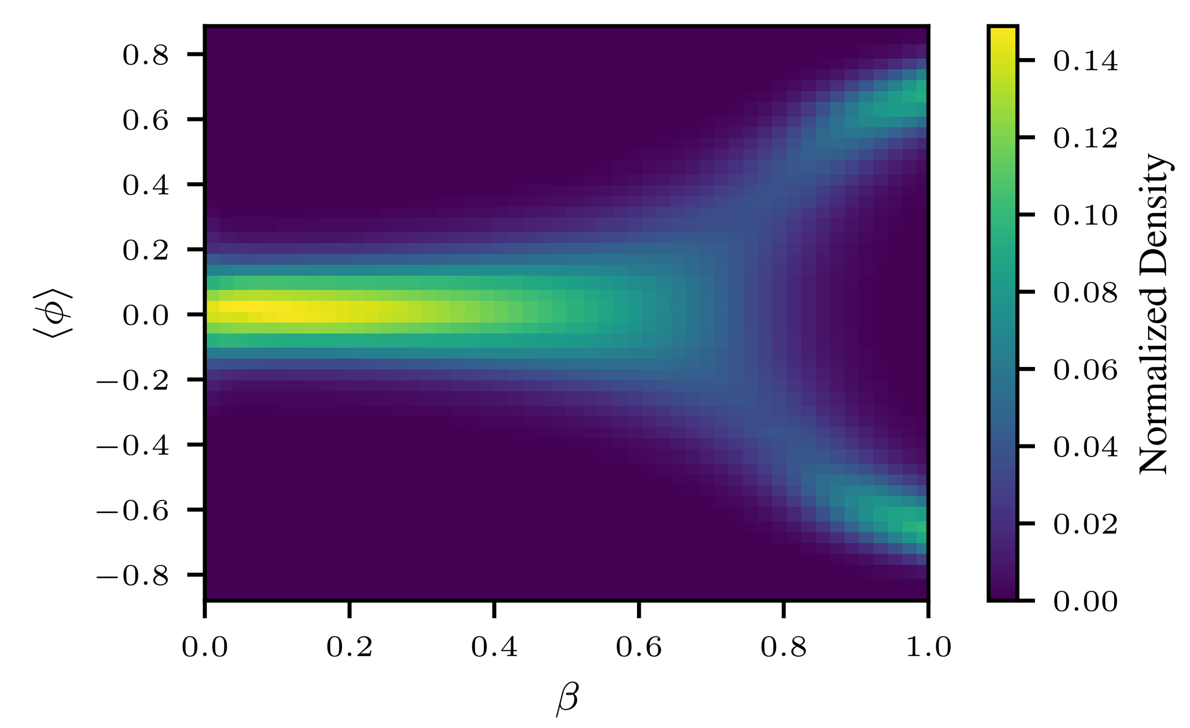

Fig. 1: Magnetization across inverse temperature β, showing density over $(\beta, \langle\phi\rangle)$ across the full temperature range.

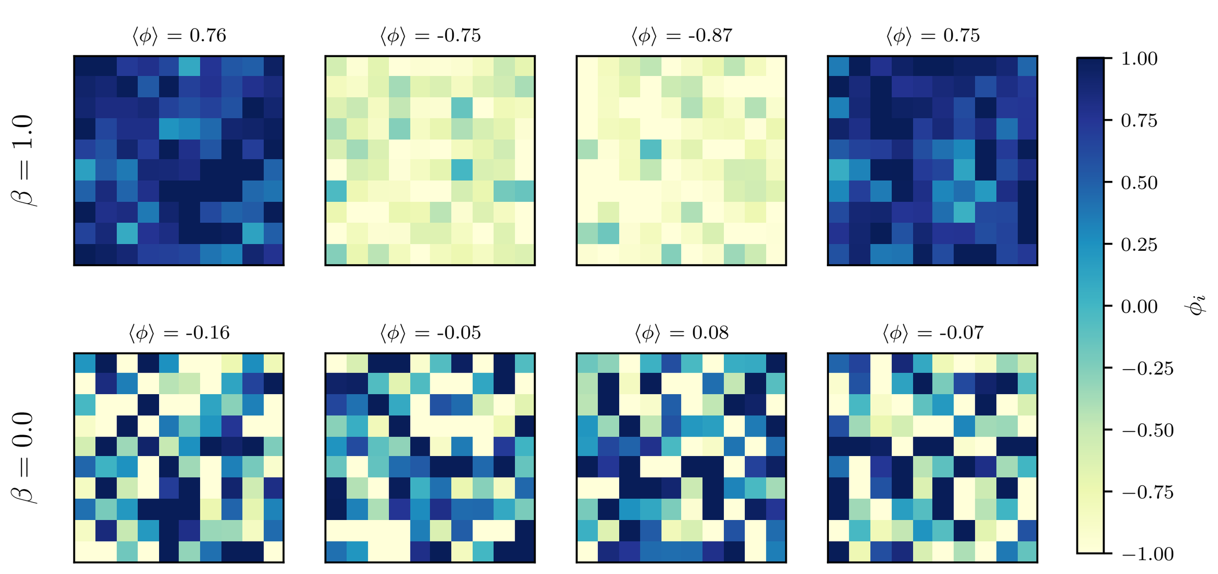

Fig. 2: Eight 10×10 field configurations sampled by NS at $\beta = 1.0$ (top) and $\beta = 0.0$ (bottom); colors denote site-wise fields ϕ_i with a shared colorbar and text annotations give mean magnetization $\langle\phi\rangle$.

We target the two-dimensional scalar ϕ^4 theory with action

$$S(\phi) = \sum_x \left[\frac{1}{2} \sum_{\nu=1}^{2} (\phi_x - \phi_{x+\hat{\nu}})^2 + \frac{1}{2} m_0^2 \phi_x^2 + \frac{\lambda}{4} \phi_x^4 \right],$$

evaluated at parameters that produce a pronounced multimodal target, $m_0^2 = -4$, $\lambda = 1$.

Observables: magnetization $m(\phi) = V^{-1} \sum_x \phi_x$, correlation length, and $\log Z$.
Lattices: $L \times L$ with $L \in \{10, 15, 18\}$ (state dimension 100–324).

Particle Monte Carlo toolkit

- **Hardware accelerated sampling**: `blackjax` implementation [7] leverages GPU parallelism over $O(10^3)$ particles, massively parallel regime.
- **Sequential Monte Carlo**: temperature ladder $\{\beta_t\}$ chosen adaptively [8]; ESS-based resampling plus rejuvenation using random walk (RW), short HMC trajectories, or independence MH (IRMH).
- **Nested sampling**: multivariate slice moves evolve live points [9].
- **Black box partition function estimation**: tune solely with particle covariance, sample field values directly. Estimate full path of bridging distribution giving $\log Z$ and estimates of critical exponents.

Results

Taking the $L = 10$, $\beta = 1$ case as a representative baseline:

Table 1: Sampling Quality and Performance Comparison. We report the mean and standard deviation of quality metrics computed against 10 reference sample sets from a long AHMC chain. MMD and W_2 measure discrepancy from ground truth (lower is better). The AHMC row establishes a baseline for inherent variance between different sets of true samples. All runtimes are in seconds on an NVIDIA L4 GPU (*AHMC takes around 200s to run on this GPU, the time listed is on CPU).

Method	MMD×1000	$W_2 \times 100$	$\log Z$	Runtime (s)
AHMC (control)	2.96 ± 0.29	418.45 ± 1.09	—	7.4*
CNF [10]	6.04 ± 0.44	418.13 ± 1.07	—	1028.5
Black-box methods				
NS	3.70 ± 0.28	412.90 ± 0.71	-65.20	36.85
SMC-RW	7.45 ± 0.60	421.76 ± 0.61	-65.26	**7.74**
SMC-HMC	**3.43 ± 0.42**	416.80 ± 0.53	-65.64	8.15
SMC-IRMH	9.55 ± 0.85	425.77 ± 0.38	-66.17	34.95
CNF MLP	9.64 ± 0.31	**411.77 ± 1.43**	—	2450.5

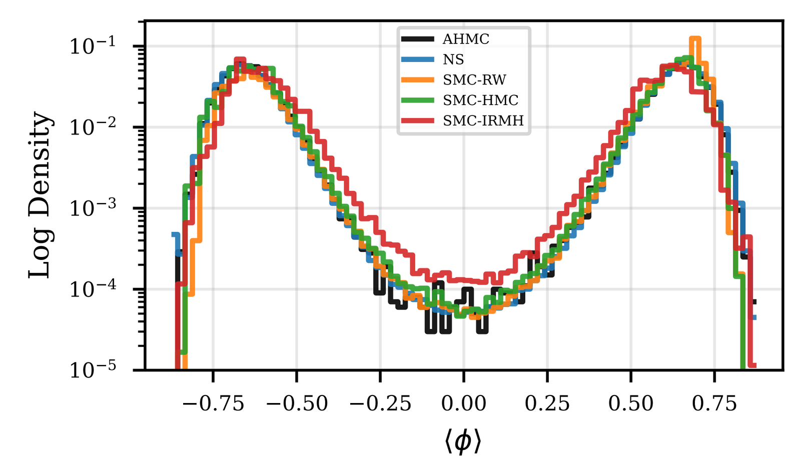

Fig. 3: Log-density magnetization histograms ($L = 10$, $\beta = 1$) comparing AHMC reference samples against the four black-box particle samplers.

- We compare against a reverse-KL trained CNF with a fixed-step Euler solver [10]. This is a well-tuned baseline for this problem that encodes the relevant symmetries (chiefly the $\mathbb{Z}_2$ symmetry). We also include a black-box variant built from three dense linear layers.
- Using MMD and W_2 metrics, all particle methods are competitive with the tuned AHMC baseline, sampling efficiently from both modes.
- Additional experiments show these results are robust to modest lattice scaling, and simple hyperparameter tuning can give further accuracy gains.

Takeaways & Further work

- Flow-based samplers amortize the cost into training and can deliver impressively efficient proposals at inference time. Learning these proposals incurs a significant cost, and seems to rely on careful engineering of relevant symmetries.
- Simulation-free, black-box training of neural samplers remains an open challenge. Particle MC methods provide a strong off-the-shelf baseline in this regime.
- For ϕ^4 theory, the ***a priori*** known symmetry is already an effective trivialising map. It is less clear how effectively neural methods apply to general inverse problems where ***a priori*** knowledge is weak.
- Combining stochastic samplers with learned proposals offers potential for the best of both worlds.
- **Still plenty of room to develop improved, fast, well-tuned particle MC methods!**

References

[1] D. Yallup, *Particle Monte Carlo methods for Lattice Field Theory*, 11, 2025, `2511.15196`.

[2] M. S. Albergo, G. Kanwar and P. E. Shanahan, *Flow-based generative models for Markov chain Monte Carlo in lattice field theory*, *Phys. Rev. D* **100** (2019) 034515 [`1904.12072`].

[3] L. Del Debbio, J. Marsh Rossney and M. Wilson, *Efficient modeling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability*, *Phys. Rev. D* **104** (2021) 094507.

[4] P. Del Moral, A. Doucet and A. Jasra, *Sequential monte carlo samplers, Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68** (2006) 411.

[5] J. Skilling, *Nested sampling for general Bayesian computation, Bayesian Analysis* **1** (2006) 833 .

[6] J. He, Y. Du, F. Vargas, D. Zhang, S. Padhy, R. OuYang et al., *No trick, no treat: Pursuits and challenges towards simulation-free training of neural samplers*, 2025.

[7] A. Cabezas, A. Corenflos, J. Lao and R. Louf, *Blackjax: Composable Bayesian inference in JAX*, 2024.

[8] P. Fearnhead and B. M. Taylor, *An adaptive sequential monte carlo sampler*, 2010.

[9] D. Yallup, N. Kroupa and W. Handley, *Nested slice sampling*, in *Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025, https://openreview.net/forum?id=ekbkMSuPo4.

[10] M. Gerdes, P. de Haan, C. Rainone, R. Bondesan and M. C. N. Cheng, *Learning lattice quantum field theories with equivariant continuous flows*, *SciPost Phys.* **15** (2023) 238 [`2207.00283`].

Contact: `dy297@cam.ac.uk` | GitHub: `yallup` | LinkedIn: `linkedin.com/in/dyallup`