



Citation for published version:

Nicoli, KA, Anders, CJ, Funcke, L, Hartung, T, Jansen, K, Kessel, P, Nakajima, S & Stornati, P 2022, 'Machine Learning of Thermodynamic Observables in the Presence of Mode Collapse', *Proceedings of Science*, vol. 396, 338. <<https://arxiv.org/abs/2111.11303>>

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link to publication](#)

Publisher Rights
CC BY-NC-ND

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Machine Learning of Thermodynamic Observables in the Presence of Mode Collapse

Kim A. Nicoli,^{a,b,*} Christopher J. Anders,^{a,b} Lena Funcke,^{c,d} Tobias Hartung,^{e,f} Karl Jansen,^g Pan Kessel,^{a,b} Shinichi Nakajima^{a,b,h} and Paolo Stornatiⁱ

^aTechnische Universität Berlin, Machine Learning Group, Marchstrasse 23, Berlin, 10587, Germany

^bBerlin Institute for the Foundation of Learning and Data (BIFOLD), Technische Universität Berlin, Berlin, Germany

^cCenter for Theoretical Physics, Co-Design Center for Quantum Advantage, and NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

^dPerimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada

^eComputation-Based Science and Technology Research Center, The Cyprus Institute, 20 Kavafi Street, 2121 Nicosia, Cyprus

^fDepartment of Mathematical Sciences University of Bath, Bath, United Kingdom

^gDeutsches Elektronen-Synchrotron DESY, Platanenallee 6, 15738 Zeuthen, Germany

^hRIKEN Center for AIP, 1-4-1 Nihonbashi, Chuo-ku, Tokyo, Japan

ⁱICFO, The Barcelona Institute of Science and Technology, Av. Carl Friedrich Gauss 3, 08860 Castelldefels (Barcelona), Spain

E-mail: kim.a.nicoli@tu-berlin.de, anders@tu-berlin.de, lfunccke@mit.edu, tobias.hartung@desy.de, karl.jansen@desy.de, pan.kessel@tu-berlin.de, nakajima@tu-berlin.de, paolo.stornati@icfo.eu

Estimating the free energy, as well as other thermodynamic observables, is a key task in lattice field theories. Recently, it has been pointed out that deep generative models can be used in this context [1]. Crucially, these models allow for the direct estimation of the free energy at a given point in parameter space. This is in contrast to existing methods based on Markov chains which generically require integration through parameter space. In this contribution, we will review this novel machine-learning-based estimation method. We will in detail discuss the issue of mode collapse and outline mitigation techniques which are particularly suited for applications at finite temperature.

Preprint number: MIT-CTP/5353

The 38th International Symposium on Lattice Field Theory, LATTICE2021 26th-30th July, 2021
Zoom/Gather@Massachusetts Institute of Technology

*Speaker

1. Introduction

Recent works have investigated the use of a particular class of deep generative machine learning models, called normalizing flows, in lattice field theory [1–7] following similar approaches in quantum chemistry [8–10] and statistical physics [11–13] (see also [14–17]). These works are proof-of-principle demonstrations for simple two-dimensional field theories and aim to reduce the integrated autocorrelation time for systems close to criticality by using the flow to generate decorrelated field samples.

Another important application of flows was recently pointed out in [12]: they can directly estimate the free energy of a lattice field theory (which can also be accomplished with methods such as tensor networks, see, e.g., [18] and references therein). The free energy is important as it allows to compute the entropy, pressure and the equation of state of the considered physical system. In the case of Quantum Chromodynamics, such thermodynamic observables are of the utmost importance in the physics of the early universe and are probed by heavy ion experiments [19].

In the following, we will review this deep-learning-based estimation technique of the free energy and discuss the important issue of mode collapse. To illustrate the mode collapse of the flow in a concrete example, consider a target density with two modes, as is the case for a quantum mechanical particle in a double well potential or scalar ϕ^4 -theory in the broken phase. The latter example will be discussed in detail in Section 3.2. For both systems, the theory has a spontaneously broken \mathbb{Z}_2 symmetry and thus two modes corresponding to the vacuum expectation values $\pm\Phi$. As we will discuss in Section 2.1, the training process can however lead to a flow that only approximates one mode of the target density of the lattice field theory and assigns (almost) vanishing probability mass to the other [1, 6]. This will lead to systematic errors of the free energy estimate which can be difficult to detect. In this contribution, we report on both mitigation and detection techniques for such a mode collapse and demonstrate their effectiveness for the example of two-dimensional scalar ϕ^4 theory.

2. Normalizing Flows

Let $f : \mathcal{M} \rightarrow \mathcal{N}$ be an orientation-preserving diffeomorphism between two orientable n -dimensional Riemannian manifolds \mathcal{M} and \mathcal{N} . We assume that there is a probability measure $q \, dV$ defined on \mathcal{M} where dV is the measure associated with volume form on \mathcal{M} and $q : \mathcal{M} \rightarrow \mathbb{R}_+$ is a positive smooth map. In particular, it holds that $\int_{\mathcal{M}} q \, dV = 1$. The push-forward measure $f_*(q \, dV)$ is then a probability measure on \mathcal{N} .

In coordinates (x^i, U) on \mathcal{N} , the push-forward $f_*(q \, dV)$ takes the form

$$q(f^{-1}(x)) |f'(f^{-1}(x))|^{-1} \, dV(f^{-1}(x)),$$

where $|f'|$ is the determinant of the Jacobian. In the machine learning literature, one therefore often refers to

$$(f_*q)(x) \equiv q(f^{-1}(x)) |f'(f^{-1}(x))|^{-1} \tag{1}$$

as the push-forward density of q .

We will be interested in the case $\mathcal{M} = \mathcal{N} = \mathbb{R}^n$ since we will consider real-valued scalar fields. The basic idea of a normalizing flow is to define a family of diffeomorphisms f_θ with parameters θ . We then adjust these parameters θ such that the push-forward density $(f_\theta)_*q$ closely approximates a certain target density p .

In practice, the diffeomorphism f_θ is parameterized by a deep neural network. Neural networks are composite functions of the form

$$\begin{aligned} m_\theta: \mathbb{R}^n &\rightarrow \mathbb{R}^n \\ z &\mapsto m_\theta(z), \end{aligned} \quad (2)$$

where $m_\theta(z) = m^L \circ \dots \circ m^1(z)$ is a composition of layers m^i defined by

$$m^i(z) = \sigma(W^i z + b^i),$$

with weights $W^i \in \mathbb{R}^{n,n}$ and biases $b^i \in \mathbb{R}^n$ being the free parameters of the neural network, i.e. $\theta = \{(W^i, b^i)\}_{i=1}^K$.¹ Furthermore, $\sigma(z)$ is a non-linear function, such as $\sigma(z) = \tanh(z)$, which is applied element-wise to each component of the vector $W^i z + b^i \in \mathbb{R}^n$. A neural network is called deep if the number of layers L is large (although there is no clearly defined threshold).

There are various approaches for parameterizing diffeomorphisms by neural networks. We will restrict to a particularly straightforward approach, called *Non-linear Independent Component Estimation* (NICE), which splits the input $z = (z_u, z_d) \in \mathbb{R}^n$ in two parts $z_u \in \mathbb{R}^{n-k}$ and $z_d \in \mathbb{R}^k$ for given $k \in \{1, n-1\}$. A diffeomorphism is then given by

$$f_\theta(z) = \begin{bmatrix} f_u(z) \\ f_d(z) \end{bmatrix} = \begin{bmatrix} z_u \\ z_d + m_\theta(z_u) \end{bmatrix}, \quad (3)$$

where m_θ is a (not necessarily invertible) neural network of the form (2). Due to the splitting of the input z , this can be easily inverted by

$$\begin{bmatrix} z_u \\ z_d \end{bmatrix} = \begin{bmatrix} f_u \\ f_d - m_\theta(f_u) \end{bmatrix}.$$

For the NICE architecture, the determinant of the Jacobian is given by

$$\left| \frac{\partial f_\theta}{\partial x} \right| = \begin{vmatrix} \frac{\partial f_u}{\partial z_u} & \frac{\partial f_u}{\partial z_d} \\ \frac{\partial f_d}{\partial z_u} & \frac{\partial f_d}{\partial z_d} \end{vmatrix} = \begin{vmatrix} \mathbb{I} & \mathbf{0} \\ * & \mathbb{I} \end{vmatrix} = 1.$$

As a result, the diffeomorphism f_θ is volume-preserving, i.e. $\left| \frac{df_\theta}{dz} \right| = 1$. In practice, we compose several of these volume-preserving diffeomorphisms. This combination is again a volume-preserving diffeomorphism because these maps form a group under composition.

A normalizing flow $f_{\theta,*}q$ is typically chosen to be a push-forward of a simple base density, e.g. $q = \mathcal{N}(0, 1)$. This allows for efficient sampling by first drawing $z \sim \mathcal{N}(0, 1)$ and then applying the diffeomorphism f_θ to the sample z , i.e.

$$f_\theta(z) \sim f_{\theta,*}q, \quad (4)$$

where the push-forward density $f_{\theta,*}q$ is given by (1).

¹We restrict to all weights and biases being of the same dimensionality since we will be interested in networks that can be used to model invertible maps.

2.1 Training of the Flow

A lattice field theory can be described by a probability density of the form

$$p(\phi) = \frac{1}{Z} \exp(-S(\phi)), \quad (5)$$

where ϕ , S and Z denote the field, its action, and the partition function respectively.

A similarity measure between two densities p and $f_{\theta^*}q$ is given by the Kullback–Leibler (KL) divergence

$$\text{KL}(f_{\theta^*}q||p) = \int \mathcal{D}[\phi] f_{\theta^*}q(\phi) \log \left(\frac{f_{\theta^*}q(\phi)}{p(\phi)} \right). \quad (6)$$

The KL divergence is non-negative and vanishes if and only if both densities are equal, i.e. $p = q$.² We can therefore train the flow to approximate the target density q by minimizing this KL divergence using gradient descent, i.e. $\theta \leftarrow \theta - \nabla_{\theta} \text{KL}(f_{\theta^*}q||p)$. For this, we observe that the KL divergence can be rewritten as

$$\text{KL}(f_{\theta^*}q||p) = \mathbb{E}_{\phi \sim f_{\theta^*}q} [S(\phi) + \log (f_{\theta^*}q(\phi))] + \text{const.},$$

where the last summand contains terms independent of θ and can thus be ignored for gradient descent. We now sample from the flow to obtain its Monte-Carlo estimator, i.e.

$$\text{KL}(f_{\theta^*}q||p) \approx \frac{1}{N} \sum_{i=1}^N [S(\phi_i) + \log (f_{\theta^*}q(\phi_i))] + \text{const.}, \quad \phi_i \sim f_{\theta^*}q.$$

The log probability can efficiently be calculated by (1). Additionally, we can very efficiently sample from the flow by pushing forward samples from the base density, see (4). However, the training of the flow may yield poor results for a multi-modal target density. This is because the training relies on self-sampling. During training, self-sampling may lead to a collapse of almost all the flow’s probability mass to a subset of the modes of the target density p . The KL divergence does not penalize this behaviour since the flow does no longer produce samples from the other modes of the target density p . We will discuss both detection and mitigation of mode collapse in the next section.

3. Flow-based Estimation of Free Energy

A promising application of normalizing flows is estimating the free energy of a lattice field theory at temperature T defined by

$$F = -T \log Z, \quad (7)$$

where Z is the partition function. The temperature is given by $T = \frac{1}{N_T a}$ with lattice spacing a and N_T denoting the number of lattice points along the temporal direction of the lattice.

²We restrict to continuous densities here. Otherwise, the densities can have different values on a set of zero measure.

3.1 MCMC-based Estimates of Free Energy

Estimating the free energy with MCMC is challenging. To illustrate this fact, we discuss a reweighting procedure [20, 21] which starts from the observation that the difference in free energies $\Delta F_{e b} = F_e - F_b = -T \log\left(\frac{Z_e}{Z_b}\right)$ between two different points e and b in parameter space can be calculated by

$$\mathbb{E}_{p_b} \left[\frac{\exp(-S_e)}{\exp(-S_b)} \right] = \frac{1}{Z_b} \int \mathcal{D}[\phi] e^{-S_b(\phi)} \frac{e^{-S_e(\phi)}}{e^{-S_b(\phi)}} = \frac{Z_e}{Z_b}. \quad (8)$$

This expectation value can be estimated by MCMC. If we choose the point b in parameter space such that the free energy F_b can be calculated exactly or approximately, we can obtain the value of the free energy at the point e by $F_e = \Delta F_{e b} + F_b$.

In practice, the variance of the estimator (8) will become prohibitively large if the two distributions p_b and p_e have a small overlap. This can be avoided by choosing intermediate distributions p_{i_1}, \dots, p_{i_K} such that neighbouring distributions p_{i_k} and $p_{i_{k+1}}$ overlap sufficiently. The free energy difference can then be obtained by

$$\Delta F_{e b} = \Delta F_{e i_K} + \Delta F_{i_K i_{K-1}} + \dots + \Delta F_{i_1 b}. \quad (9)$$

This comes at the price of an accumulated error of all free energy differences $\Delta F_{i_k i_{k+1}}$. The error therefore crucially depends on all points of the (discretized) trajectory connecting the points e and b in parameter space.

3.2 Example: Two-dimensional ϕ^4 Theory

This dependence on the trajectory can lead to serious problems, as we illustrate in a concrete example of the ϕ^4 theory in two dimensions with the action

$$S = \sum_{x \in \Lambda} (-2\kappa \sum_{\hat{\mu}=1}^2 \varphi(x) \varphi(x + \hat{\mu}) + (1 - 2\lambda) \varphi(x)^2 + \lambda \varphi(x)^4), \quad (10)$$

where κ is the hopping parameter and λ denotes the bare coupling. For vanishing hopping parameter κ , the free energy can be calculated analytically [1] and is given by

$$F(\lambda) = -|\Lambda| T \ln z(\lambda),$$

where $|\Lambda|$ denotes the number of sites of the lattice Λ and

$$z(\lambda) = \sqrt{\frac{1 - 2\lambda}{4\lambda}} \exp\left(\frac{(1 - 2\lambda)^2}{8\lambda}\right) K_{\frac{1}{4}}\left(\frac{(1 - 2\lambda)^2}{8\lambda}\right),$$

with K_n being the Bessel function of the second kind.

As the hopping parameter κ is increased, spontaneous breaking of the \mathbb{Z}_2 -symmetry $\phi \rightarrow -\phi$ is observed. This is illustrated in Figure 1. Now, suppose we want to calculate the free energy with MCMC for parameters in the broken phase, e.g. $\lambda_e = 0.022$ and $\kappa_e = 0.5$. We can then choose a trajectory through parameter space for which the bare coupling is kept constant, i.e. $\lambda_{i_k} = 0.022$, and the initial hopping parameter is $\kappa_b = 0$. We then increase the hopping parameter by a step size




Figure 1: Left: absolute magnetization density $\frac{\langle |\phi| \rangle}{|\Lambda|}$, where $|\Lambda|$ denotes the number of lattice points, as a function of the hopping parameter κ for the bare coupling $\lambda = 0.022$ on a 16×8 lattice. The values were estimated with an overrelaxed HMC [22–25]. **Right:** contributions to the error from the HMC free energy estimate at hopping parameter $\kappa = 0.6$ and bare coupling $\lambda = 0.022$ for a 16×8 lattice along the discretized trajectory through parameter space. The width of a bar corresponds to the step size $\Delta\kappa$. The area of a bar shows the error of the corresponding free energy difference ΔF (calculated as described in [26]), i.e. the total area of all bars is the total error of the free energy estimate F_e at the target value $\kappa_e = 0.6$. Every chain is four hundred thousands steps long with an overrelaxation every ten steps. Even for this relatively small lattice, a significant contribution to the overall error comes from the region around the critical value of κ .

$\Delta\kappa = 0.05$ up to $\kappa = 0.2$ and then use a smaller step size $\Delta\kappa = 0.01$ in order to ensure sufficient overlap. Crucially, the estimate of the free energy in the broken phase will now suffer from critical slowing down as the corresponding trajectory has to cross the phase transition in order to reach the initial hopping parameter $\kappa_a = 0$. This will lead to a significant increase in the statistical error, see Figure 1.

3.3 Flow-based Estimators of the Free Energy

Normalizing flows allow us to directly estimate the free energy $F = -T \log(Z)$ at a given point in parameter space and therefore allow us to avoid critical slowing down in the specific situations discussed in the previous section. This can be seen by observing that we can estimate the partition function Z using a trained flow in two different ways. Firstly, we can use samples from the flow

$$Z = \mathbb{E}_{\phi \sim f_{\theta^*} q} \left[\frac{e^{-S(\phi)}}{f_{\theta^*} q(\phi)} \right] \approx \frac{1}{N} \sum_{i=1}^N \frac{e^{-S(\phi_i)}}{f_{\theta^*} q(\phi_i)} \equiv \hat{Z}_q, \quad \phi_i \sim f_{\theta^*} q.$$

Using this definition, we obtain the *reverse estimator* of the free energy by

$$\hat{F}_q = -T \log(\hat{Z}_q). \quad (11)$$

Secondly, one can use samples from the target density p

$$Z^{-1} = \mathbb{E}_{\phi \sim p} \left[\frac{f_{\theta^*} q(\phi)}{e^{-S(\phi)}} \right] \approx \frac{1}{N} \sum_{i=1}^N \frac{f_{\theta^*} q(\phi_i)}{e^{-S(\phi_i)}} \equiv \hat{Z}_p^{-1}, \quad \phi_i \sim p.$$

to obtain the *forward estimator*

$$\hat{F}_p = T \log(\hat{Z}_p^{-1}). \quad (12)$$

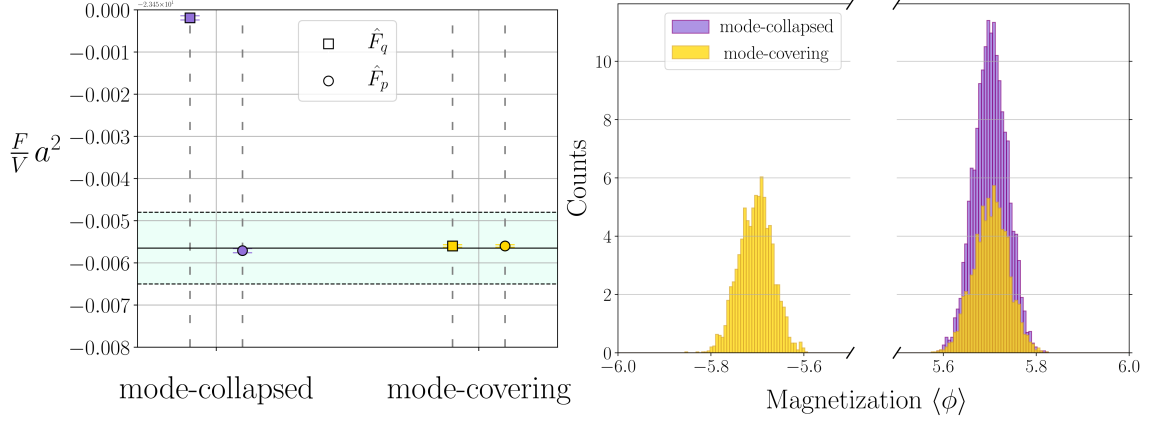


Figure 2: **Left:** estimation of the free energy using both the forward and reverse estimators. One of the flows has mode-collapsed (purple) while the other is mode-covering (yellow). For the former case, this results in a discrepancy in the free energy prediction of the two estimators while for the latter we obtain consistent results. The MCMC reference values are estimated by the method described in Section 3.1. The mean of MCMC estimate is shown by the solid line while the color band denotes the confidence interval of one standard deviation. As the MCMC algorithm, we again use an HMC with overrelaxation. **Right:** histogram of the magnetization obtained by direct sampling from both a mode-collapsed and mode-covering flow.

Both estimators have relative strengths and weaknesses. If we are confident that the flow closely approximates the target density p , it is advisable to use the reverse estimator (12) because sampling from the flow is more efficient. However, this estimator may lead to incorrect results if the flow is mode-dropping. In contrast, the forward estimator \hat{F}_p uses samples from p and thus cannot neglect any mode of the target density p . If mode-dropping is a risk (for example in the broken phase of the ϕ^4 -theory), one should therefore also use the forward estimator (12) as a consistency check.

4. Numerical Experiments

In the following, we will illustrate the difference in using the forward and the reverse variants for the free energy estimation in the presence of mode-dropping. To this end, we consider two normalizing flows trained for the two-dimensional scalar ϕ^4 -theory for a hopping parameter of $\kappa = 0.6$ and a bare coupling of $\lambda = 0.022$ on a $N_T \times N_L = 16 \times 8$ lattice. The theory is thus considered in its broken phase, see Figure 2. One of the flows is mode-collapsed on a single mode of the target density p , while the other flow covers both modes, as can be seen on the right of Figure 2.

For both flows, we then use the forward estimator \hat{F}_p and reverse estimator \hat{F}_q as defined in (12) and (11) respectively. The estimated values for the free energy are visualized in Figure 2. For the mode-collapsed flow, we see a clear discrepancy in the prediction while the mode-covering flow leads to consistent values of the forward and reverse estimators. The fact that the forward estimator \hat{F}_q gives the correct result for the mode-collapsed model can heuristically be understood by assuming that the flow is approximately $f_{\theta^*}q(x) \approx 2p(x)$ for the covered mode \mathcal{M}_1 and $f_{\theta^*}q(x) \approx 0$ for the

other mode \mathcal{M}_2 . This implies that

$$\mathbb{E}_{\phi \sim p} \left[\frac{f_{\theta^*} q(\phi)}{e^{-S(\phi)}} \right] \approx \underbrace{\int_{\mathcal{M}_1} p(\phi) \frac{f_{\theta^*} q(\phi)}{e^{-S(\phi)}}}_{\approx \frac{2}{Z}} + \underbrace{\int_{\mathcal{M}_2} p(\phi) \frac{f_{\theta^*} q(\phi)}{e^{-S(\phi)}}}_{\approx 0} \approx \frac{2}{Z} \underbrace{\int_{\mathcal{M}_1} p(\phi)}_{\approx \frac{1}{2}} \approx Z^{-1}. \quad (13)$$

In summary, this experiment clearly illustrates that forward estimation of the free energy is crucial in the presence of mode collapse.

5. Conclusion

Deep generative models, in particular normalizing flows, allow for a direct estimation of the free energy. Current normalizing flow architectures are however far from perfect. For example, they are challenging to train in the broken phase (particularly for larger lattices) and can suffer from mode collapse for multi-modal densities. In this contribution, we have briefly outlined how forward estimation of the free energy can help to mitigate this weakness.

6. Acknowledgements

K.A.N. , C.A. , P.K. and S.N. are funded by the German Ministry for Education and Research as BIFOLD - Berlin Institute for the Foundations of Learning and Data (ref. 01IS18025A and ref 01IS18037A). P.S. is supported from Agencia Estatal de Investigación (“Severo Ochoa” Center of Excellence CEX2019-000910-S, Plan National FIDEUA PID2019-106901GB-I00/10.13039 / 501100011033, FPI), Fundació Privada Cellex, Fundació Mir-Puig, and from Generalitat de Catalunya (AGAUR Grant No. 2017 SGR 1341, CERCA program). L.F. is partially supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C²QA) under contract number DE-SC0012704, by the DOE QuantiSED Consortium under subcontract number 675352, by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>), and by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under grant contract numbers DE-SC0011090 and DE-SC0021006. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Industry Canada and by the Province of Ontario through the Ministry of Colleges and Universities

References

- [1] K.A. Nicoli, C.J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel et al., *Estimation of thermodynamic observables in lattice field theories with deep generative models*, *Phys. Rev. Lett.* **126** (2021) 032001.
- [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.

- [3] G. Kanwar, M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, S. Racanière et al., *Equivariant flow-based sampling for lattice gauge theory*, *Phys. Rev. Lett.* **125** (2020) 121601.
- [4] D. Boyda, G. Kanwar, S. Racanière, D.J. Rezende, M.S. Albergo, K. Cranmer et al., *Sampling using $su(n)$ gauge equivariant flows*, *arXiv preprint arXiv:2008.05456* (2020) .
- [5] L. Del Debbio, J.M. Rossney and M. Wilson, *Efficient modelling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability*, *arXiv preprint arXiv:2105.12481* (2021) .
- [6] D.C. Hackett, C.-C. Hsieh, M.S. Albergo, D. Boyda, J.-W. Chen, K.-F. Chen et al., *Flow-based sampling for multimodal distributions in lattice field theory*, *arXiv preprint arXiv:2107.00734* (2021) .
- [7] P. de Haan, C. Rainone, M. Cheng and R. Bondesan, *Scaling up machine learning for quantum field theory with equivariant continuous flows*, *arXiv preprint arXiv:2110.02673* (2021) .
- [8] F. Noé, S. Olsson, J. Köhler and H. Wu, *Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning*, *Science* **365** (2019) .
- [9] P. Wirnsberger, A.J. Ballard, G. Papamakarios, S. Abercrombie, S. Racanière, A. Pritzel et al., *Targeted free energy estimation via learned mappings*, *The Journal of Chemical Physics* **153** (2020) 144112.
- [10] P. Wirnsberger, G. Papamakarios, B. Ibarz, S. Racanière, A.J. Ballard, A. Pritzel et al., *Normalizing flows for atomic solids*, *arXiv preprint arXiv:2111.08696* (2021) .
- [11] D. Wu, L. Wang and P. Zhang, *Solving statistical mechanics using variational autoregressive networks*, *Physical Review Letters* **122** (2019) 080602.
- [12] K.A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller and P. Kessel, *Asymptotically unbiased estimation of physical observables with neural samplers*, *Physical Review E* **101** (2020) 023304.
- [13] K. Nicoli, P. Kessel, N. Strodthoff, W. Samek, K.-R. Müller and S. Nakajima, *Comment on "solving statistical mechanics using vns": Introducing savant-vns enhanced by importance and mcmc sampling*, *arXiv preprint arXiv:1903.11048* (2019) .
- [14] D. Bachtis, G. Aarts and B. Lucini, *Quantum field-theoretic machine learning*, *Phys. Rev. D* **103** (2021) 074510.
- [15] J.M. Pawłowski and J.M. Urban, *Reducing autocorrelation times in lattice simulations with generative adversarial networks*, *Machine Learning: Science and Technology* **1** (2020) 045011.
- [16] A. Tanaka and A. Tomiya, *Towards reduction of autocorrelation in hmc by machine learning*, *arXiv preprint arXiv:1712.03893* (2017) .

- [17] S. Bulusu, M. Favoni, A. Ipp, D.I. Müller and D. Schuh, *Generalization capabilities of translationally equivariant neural networks*, *Phys. Rev. D* **104** (2021) 074504.
- [18] S. Akiyama, D. Kadoh, Y. Kuramashi, T. Yamashita and Y. Yoshimura, *Tensor renormalization group approach to four-dimensional complex ϕ^4 theory at finite density*, *arXiv preprint arXiv:2005.04645* (2020) .
- [19] W. Busza, K. Rajagopal and W. Van Der Schee, *Heavy ion collisions: the big picture and the big questions*, *Annual Review of Nuclear and Particle Science* **68** (2018) 339.
- [20] P. De Forcrand, M. D'Elia and M. Pepe, *'t hooft loop in $su(2)$ yang-mills theory*, *Phys. Rev. Lett.* **86** (2001) 1438.
- [21] O. Philipsen, *The qcd equation of state from the lattice*, *Progress in Particle and Nuclear Physics* **70** (2013) 55.
- [22] S.L. Adler, *Over-relaxation method for the monte carlo evaluation of the partition function for multiquadratic actions*, *Phys. Rev. D* **23** (1981) 2901.
- [23] C. Whitmer, *Over-relaxation methods for monte carlo simulations of quadratic and multiquadratic actions*, *Phys. Rev. D* **29** (1984) 306.
- [24] D.J. Callaway and A. Rahman, *Lattice gauge theory in the microcanonical ensemble*, *Phys. Rev. D* **28** (1983) 1506.
- [25] Z. Fodor and K. Jansen, *Overrelaxation algorithm for coupled gauge-higgs systems*, *Phys. Rev. B* **331** (1994) 119.
- [26] U. Wolff, A. Collaboration et al., *Monte carlo errors with less errors*, *Computer Physics Communications* **156** (2004) 143.