

Flow-based density of states for complex actions

Jan M. Pawłowski¹ and Julian M. Urban^{1*}

Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany



(Received 2 May 2022; accepted 12 September 2023; published 29 September 2023)

Emerging sampling algorithms based on normalizing flows have the potential to solve ergodicity problems in lattice calculations. Furthermore, it has been noted that flows can be used to compute thermodynamic quantities which are difficult to access with traditional methods. This suggests that they are also applicable to the density-of-states approach to complex action problems. In particular, flow-based sampling may be used to compute the density directly, in contradistinction to the conventional strategy of reconstructing it via measuring and integrating the derivative of its logarithm. By circumventing this procedure, the accumulation of errors from the numerical integration is avoided completely and the overall normalization factor can be determined explicitly. In this proof-of-principle study, we demonstrate our method in the context of two-component scalar field theory where the internal $O(2)$ symmetry is explicitly broken by an imaginary external field. First, we concentrate on the zero-dimensional case which can be solved exactly. We show that with our method, the Lee-Yang zeroes of the associated partition function can be successfully located. Subsequently, we confirm that the flow-based approach correctly reproduces the density computed with conventional methods in one- and two-dimensional models.

DOI: [10.1103/PhysRevD.108.054511](https://doi.org/10.1103/PhysRevD.108.054511)

I. INTRODUCTION

Lattice calculations are a powerful approach to study quantum field theories nonperturbatively by applying Markov Chain Monte Carlo (MCMC) sampling [1], see [2–9] for recent reviews. However, for many physically interesting cases, the associated Euclidean lattice action is complex valued. This prohibits the application of standard importance sampling, the most prominent example being the notorious sign problem in lattice QCD at finite chemical potential [10,11]. In this context, it has been shown that with the density-of-states (DoS) approach [12–22], certain complex action problems can be successfully treated [23–29]. However, directly computing the DoS is generally not possible due to the intrinsically high variance of the associated observables. Instead, the usual approach is to measure the derivative of its logarithm with restricted MCMC calculations, followed by reconstructing the DoS itself via numerical integration. The high precision required to control the accumulation of errors from the approximation of the integral can be computationally expensive.

Recently, it has been noted that similar thermodynamic quantities in lattice field theory can be computed directly

using generative machine learning models with tractable probability densities [30–32], thereby completely avoiding the aforementioned numerical reconstruction of the quantity of interest. Normalizing flows are one such class of probabilistic models for which both efficient sampling and density estimation are made possible using a change of variables formula [33–37]. They have been successfully applied to model real scalar field theory [38–42], $U(1)$ and $SU(N)$ gauge theories [43–47], as well as theories with dynamical fermions [48–50]. For their application to sign problems, flows have been studied in the context of contour deformations [51,52].

In this work, we apply flow-based sampling to the direct computation of the DoS for lattice field theories with complex actions. Specifically, we investigate scalar ϕ^4 -theory with two real-valued components where one is coupled to an imaginary external field, thereby explicitly breaking the internal $O(2)$ symmetry. We first consider the exactly solvable, zero-dimensional case as a toy model for a proof-of-principle demonstration, showing that the DoS as well as the partition function and magnetization as functions of the external field are computed correctly with our approach. In particular, we can locate Lee-Yang zeros [53] of the partition function together with the associated discontinuities in the magnetization. We then apply the approach to actual lattice models in one and two dimensions, accurately reproducing the densities obtained with conventional MCMC methods.

This paper is organized as follows. In Sec. II, we briefly review the DoS approach pertinent to the type of complex

*Corresponding author: urban@thphys.uni-heidelberg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

action problem considered here. Section III serves to introduce the basic concepts of normalizing flows that are relevant to this work. We explain our approach in Sec. IV and present numerical results in Sec. V. We summarize our contributions and provide an outlook in Sec. VI.

II. DENSITY OF STATES

We consider lattice field theories with complex-valued actions of the form

$$S(\phi) = S_r(\phi) + ihX(\phi), \quad (1)$$

where $S_r, X, h \in \mathbb{R}$. The partition function and expectation values of observables are defined as

$$Z = \int \mathcal{D}\phi e^{-S_r(\phi) - ihX(\phi)}, \quad (2)$$

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}\phi e^{-S_r(\phi) - ihX(\phi)} \mathcal{O}(\phi). \quad (3)$$

Since the action is complex, standard importance sampling is not directly applicable and reweighting often becomes prohibitively expensive when increasing h due to the average phase factor being close to zero.

One ansatz to make the computation more tractable is to consider the DoS as a function of the quantity that generates the imaginary part of the action, i.e.

$$\rho(c) = \int \mathcal{D}\phi e^{-S_r(\phi)} \delta(X(\phi) - c). \quad (4)$$

Essentially, $\rho(c)$ corresponds to slices of the partition function for the real part of the action, with the configuration space restricted to hypersurfaces of constant $X(\phi) = c$. In MCMC calculations, this restriction can be achieved e.g. by confining the dynamics through additional rejections, or by replacing the δ -distribution with a Gaussian of finite width, which is the approach used in the present work (see Sec. IV for details).

If $\rho(c)$ is known, the partition function for the full action as well as expectation values of observables (that are functions of c only) can be computed in terms of one-dimensional integrals with a residual phase,

$$Z = \int dc \rho(c) e^{-ihc}, \quad (5)$$

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int dc \rho(c) e^{-ihc} \mathcal{O}(c). \quad (6)$$

However, similar to partition functions themselves and thermodynamic quantities in general, a direct computation of $\rho(c)$ is often infeasible with conventional MCMC

algorithms due to the high variance associated with the required observables. Instead, it is usually reconstructed from measurements of $\partial_c \log \rho(c)$, as detailed below.

III. NORMALIZING FLOWS

Starting with a prior distribution over a continuous space \mathcal{X} with a known probability density $r(\xi)$, an invertible transport map (flow) $f: \mathcal{X} \rightarrow \mathcal{X}, \xi \mapsto \phi$ can be used to redistribute samples under r to samples under a new density $q(\phi)$. We only require that the map be diffeomorphic, i.e. that both f and its inverse are differentiable. The resulting density $q(\phi)$ is fixed by the choice of prior distribution and map, and it can be evaluated explicitly as

$$q(\phi) = r(\xi) \left| \det \left(\frac{\partial f}{\partial \xi} \right) \right|^{-1}, \quad (7)$$

where $\phi = f(\xi)$ and $\det(\partial f / \partial \xi)$ is the Jacobian determinant of f . Because the density after the transformation can be computed explicitly, flows provide a mechanism for both sampling and density estimation.

By choosing a sufficiently expressive parametrization of f , the space of associated transformations—corresponding to a large variational family of model densities q —can be explored through numerical optimization in order to find an instance that best approximates some target density p . In particular, the parameters of f may be optimized by performing stochastic gradient descent on a measure of the discrepancy between the two densities q and p , i.e. an appropriate loss function.

A common choice is the Kullback-Leibler (KL) divergence, which is a measure of the relative entropy between distributions. It is defined as

$$\begin{aligned} D_{\text{KL}}(q||p) &= \int \mathcal{D}\phi q(\phi) (\log q(\phi) - \log p(\phi)) \\ &= \langle \log q(\phi) - \log p(\phi) \rangle_{\phi \sim q(\phi)} \geq 0, \end{aligned} \quad (8)$$

and takes the minimum value $D_{\text{KL}}(q||p) = 0$ if and only if $q = p$. With expectation values measured using samples from the model distribution q , $D_{\text{KL}}(q||p)$ can be stochastically estimated without requiring samples from the target distribution p .

For targets of the form $p(\phi) = e^{-S(\phi)} / Z$, the KL divergence takes the form

$$D_{\text{KL}}(q||p) = \langle \log q(\phi) + S(\phi) \rangle_{\phi \sim q(\phi)} + \log Z. \quad (9)$$

Since the partition function Z is usually not known *a priori*, D_{KL} can only be estimated up to the constant $\log Z$. However, this does not affect gradients and one may freely use $(D_{\text{KL}} - \log Z)$ as the loss function for optimization, which then provides a bound on $\log Z$. Writing the partition function as

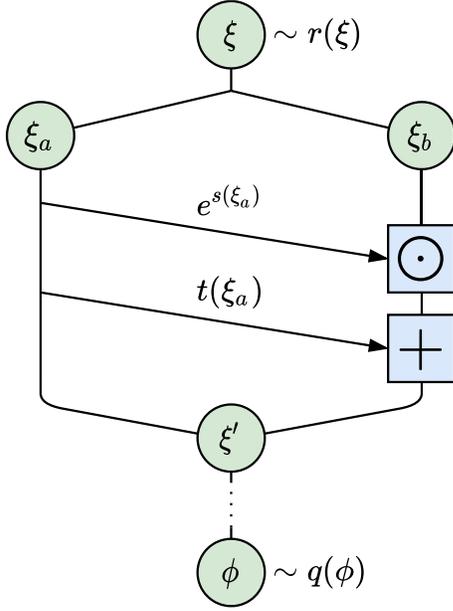


FIG. 1. Illustration of the affine coupling layer defined in Eq. (11). The blue boxes depict elementwise operations.

$$Z = \int \mathcal{D}\phi q(\phi) \frac{e^{-S(\phi)}}{q(\phi)} = \langle e^{-S(\phi) - \log q(\phi)} \rangle_{\phi \sim q(\phi)}, \quad (10)$$

it follows that any model giving good agreement to the target distribution necessarily provides a precise, unbiased estimate of Z at low cost through model samples alone, without the need to ever sample from the target.

A common building block for the construction of invertible flow transformations is the affine coupling layer. In each such layer, the input ξ is split into two equal-sized subsets ξ_a, ξ_b which are transformed according to

$$\begin{aligned} \xi'_a &= \xi_a, \\ \xi'_b &= \xi_b \odot e^{s(\xi_a)} + t(\xi_a), \end{aligned} \quad (11)$$

i.e. ξ_a remains unchanged (frozen) while ξ_b is updated (active); see Fig. 1 for an illustration. Here, the symbol \odot denotes elementwise multiplication. Each affine coupling layer is trivially invertible and has a triangular Jacobian matrix, thereby making the computation of its determinant and thus the model density q tractable. The context functions s, t that take the frozen variables as inputs and are used to update the active ones can be arbitrary functions. They are commonly parametrized by deep, feedforward neural networks, with are then trained during optimization. Importantly, they are not required to be invertible themselves, thereby providing much freedom in choosing a particular parametrization. Expressive flow transformations are built by chaining together many such affine coupling layers with alternating frozen and active subsets.

IV. FLOW-BASED DENSITY OF STATES

As already mentioned in Sec. II, we employ a formulation of the DoS approach where the δ -distribution in Eq. (4) is replaced by a Gaussian of finite width, following e.g. [18,22]. This enables the straightforward application of both standard sampling algorithms like hybrid/Hamiltonian Monte Carlo (HMC) as well as our flow-based approach. Exactness of all expressions can be retained at the cost of a residual sign problem (which is tractable for sufficiently small width) or by extrapolating to the limit of vanishing width.

First, we note that the result of the Gaussian integral

$$\int dc e^{-\frac{P}{2}(c-a)^2} = \sqrt{\frac{2\pi}{P}} \equiv \mathcal{N} \quad (12)$$

is independent of a . Hence, we can rewrite Eq. (2) as

$$Z = \int \mathcal{D}\phi \int dc e^{-\frac{P}{2}(c-X(\phi))^2 - \log \mathcal{N} - S_r(\phi) - ihX(\phi)}. \quad (13)$$

We then define the P -dependent DoS as

$$\rho_P(c) = \int \mathcal{D}\phi e^{-S_{c,P}(\phi)}, \quad (14)$$

where

$$S_{c,P}(\phi) = S_r(\phi) + \frac{P}{2}(c - X(\phi))^2 + \log \mathcal{N}. \quad (15)$$

The “true” DoS as defined in Eq. (4) is recovered in the limit $P \rightarrow \infty$.

Assuming continuity and convergence of the integrals, the partition function can be expressed in terms of ρ_P as

$$\begin{aligned} Z &= \int dc \int \mathcal{D}\phi e^{-S_{c,P}(\phi) - ihX(\phi)} \\ &= \int dc \rho_P(c) \frac{\int \mathcal{D}\phi e^{-S_{c,P}(\phi) - ihX(\phi)}}{\int \mathcal{D}\phi e^{-S_{c,P}(\phi)}} \\ &= \int dc \rho_P(c) \langle e^{-ihX(\phi)} \rangle_{\phi \sim e^{-S_{c,P}(\phi)}}, \end{aligned} \quad (16)$$

Hence, in this formulation, the partition function is still a one-dimensional integral over the P -dependent DoS, but with an additional average phase factor computed on ensembles sampled with $S_{c,P}(\phi)$. The fluctuations of this phase factor are tractable as long as the parameter P is large enough, such that $X(\phi)$ does not deviate too strongly from c . Accordingly, expectation values of observables can be written as

$$\langle \mathcal{O} \rangle = \frac{\int dc \rho_P(c) \langle e^{-ihX(\phi)} \mathcal{O}(\phi) \rangle_{\phi \sim e^{-S_{c,P}(\phi)}}}{\int dc \rho_P(c) \langle e^{-ihX(\phi)} \rangle_{\phi \sim e^{-S_{c,P}(\phi)}}}. \quad (17)$$

As mentioned previously, a direct computation of $\rho_P(c)$ with traditional MCMC methods is often infeasible for problems of interest. Instead, the usual strategy is to compute

$$\begin{aligned} \frac{\partial \log \rho_P(c)}{\partial c} &= \frac{\int \mathcal{D}\phi e^{-S_{c,P}(\phi)} (-P(c - X(\phi)))}{\int \mathcal{D}\phi e^{-S_{c,P}(\phi)}} \\ &= \langle -P(c - X(\phi)) \rangle_{\phi \sim e^{-S_{c,P}(\phi)}}, \end{aligned} \quad (18)$$

and then to reconstruct $\log(\rho_P(c)/\rho_P(0))$ by numerical integration, e.g. with the trapezoidal rule. In contrast, normalizing flows trained with $S_{c,P}(\phi)$ as the target action allow for a direct computation of $\rho_P(c)$ [including the overall factor $\rho_P(0)$] using configurations sampled from $q(\phi)$, as long as the overlap of the target and model distributions is sufficient. This can be seen by rewriting Eq. (14) as

$$\begin{aligned} \rho_P(c) &= \int \mathcal{D}\phi q(\phi) \frac{e^{-S_{c,P}(\phi)}}{q(\phi)} \\ &= \langle e^{-S_{c,P}(\phi) - \log q(\phi)} \rangle_{\phi \sim q(\phi)}, \end{aligned} \quad (19)$$

similar to Eq. (10). A successfully trained flow minimizes the fluctuations of the exponent in the last expression, such that the variance of the expectation value remains tractable. This is precisely the crucial advantage of flow-based sampling over conventional MCMC methods that allows the computation of thermodynamic quantities via variationally optimized reweighting [31,32].

In order to compute ρ_P across a wide range, one could train independent flows for each value of c . Alternatively, a more efficient approach would be to start by training one flow at some given point (e.g. $c = 0$) and then perform retraining for each additional point. However, since high precision in c is desired, these strategies seem impractical. Apart from such a training procedure already being computationally expensive, a large number of different parameter sets for all the individual flow transformations would then have to be stored and loaded into memory for evaluation. Instead, we propose to encode the full information about ρ_P for all c in a single-flow model. This is achieved by promoting the transport map $f(\xi)$ to a conditional transformation $f_c(\xi)$, which additionally depends on c . In particular, the context functions s , t of all affine couplings as defined in Eq. (11) are modified to take c as an additional input. This only marginally increases the computational effort of evaluating the transformation, although it may be necessary to make the flow more expressive overall in order to properly model the dependence on c .

Furthermore, we introduce an additional c -dependent offset at the last layer, such that the conditional generation of field configurations ϕ from prior samples ξ takes the form

$$\phi(\xi|c) = f_c(\xi) + \bar{\phi}(c), \quad (20)$$

with $\bar{\phi}(c)$ chosen such that $X(\bar{\phi}(c)) = c$. This offset already provides the correct mean field configuration for each c and thereby greatly simplifies training from the start, because the flow only has to model the distribution around the given $\bar{\phi}(c)$. Since this amounts to just a constant shift, the Jacobian of the transformation remains unchanged.

It should be stressed that for the simple architecture utilized in the present work, outputs of intermediate layers are not necessarily meaningful since we do not specify a particular flow trajectory in the space of probability distributions. Only the field variables at the prior and target endpoints of the transformation carry definite meaning. Nevertheless, enforcing such a particular sequence of intermediate densities by appropriate loss functions—e.g., by interpolating the coupling constants of nonquadratic action terms as well as the parameter P from 0 to their target values, thereby resulting in a Gaussian prior on one side and the desired target density on the other side—could have several advantages, and will be considered in future work. Most importantly, such an approach makes the relation of the present approach to thermodynamic integration and the Jarzynski equality [54] more concrete. This, in turn, enables the application of stochastic normalizing flow approaches [41,42,55], which may be particularly effective at determining the types of thermodynamic quantities considered here. Furthermore, explicitly determining the precise P -dependence of the DoS in this manner may allow a direct extrapolation to the vanishing-width limit without the need to optimize several flow instances for different target parameters.

Taken together, the full transformation defined in Eq. (20) consisting of a conditional transport map and an additional offset induces a conditional model distribution q_c for each c , such that the P -dependent DoS may be computed as

$$\rho_P(c) = \langle e^{-S_{c,P}(\phi) - \log q_c(\phi)} \rangle_{\phi \sim q_c(\phi)}. \quad (21)$$

During training, one may estimate the conditional KL divergence $D_{\text{KL}}(q_c \| e^{-S_{c,P}})$ at random points c , distributed uniformly across a sufficiently large interval, in order to enforce optimal generalization for arbitrary c .

At this point we emphasize again that in order to evaluate the above expression for ρ_P , only samples from the model distribution are required. Importantly, this implies that once the flow has been trained, the remaining computations can be performed extremely efficiently in a manner that has been fittingly described as “embarrassingly parallel”.

In particular, field configurations do not need to be arranged in a Markov chain and no accept/reject steps are necessary. This constitutes a further potential advantage of our proposed approach over conventional MCMC calculations.

V. NUMERICAL RESULTS

A. Zero-dimensional model

For a first demonstration of our approach, we consider a zero-dimensional model of a single two-component scalar field $\phi = (\phi_1, \phi_2)$ with quartic self-interaction in an imaginary external field. The simplicity of this model facilitates a comparison to exact results. The action is

$$S(\phi) = \frac{m^2}{2} (\phi_1^2 + \phi_2^2) + \frac{\lambda}{4} (\phi_1^2 + \phi_2^2)^2 + ih\phi_1, \quad (22)$$

where $\phi_1, \phi_2, m^2, \lambda, h \in \mathbb{R}$. Accordingly, we can identify $X(\phi) \equiv \phi_1$ and

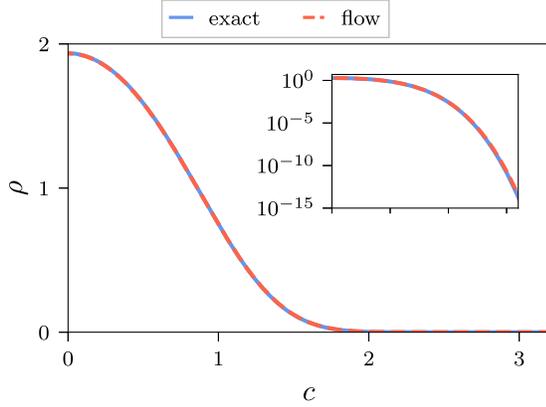


FIG. 2. Comparison of the DoS computed with flow-based sampling to the exact solution for the zero-dimensional model.

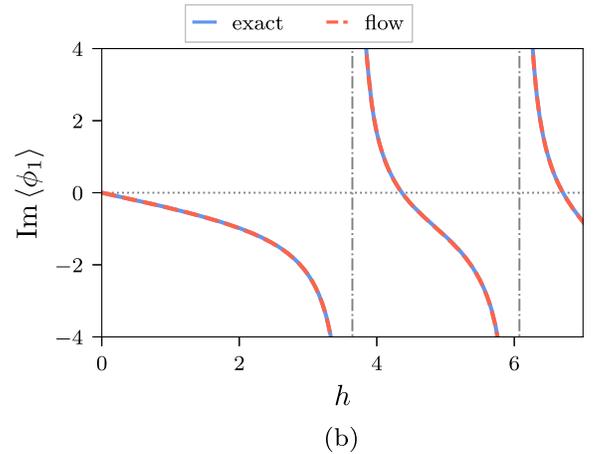
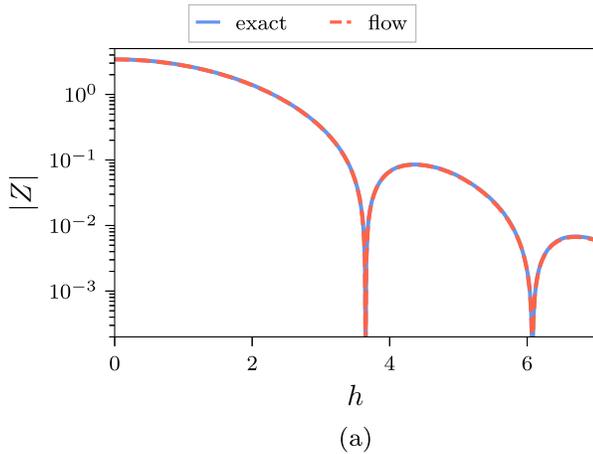


FIG. 3. Comparison of the flow results to the exact solution for the zero-dimensional model: partition function (a) and average imaginary part of ϕ_1 ; (b) as functions of h . The locations of the Lee-Yang zeroes in (a) are associated with the discontinuities of the observable in (b).

$$S_{c,P}(\phi) \equiv \frac{m^2}{2} (\phi_1^2 + \phi_2^2) + \frac{\lambda}{4} (\phi_1^2 + \phi_2^2)^2 + \frac{P}{2} (c - \phi_1)^2 + \log \mathcal{N}. \quad (23)$$

We train a normalizing flow using this target action with $m^2 = 1$, $\lambda = 1$, $P = 1000$. The context functions s , t are implemented as fully connected networks where the condition c is provided as an additional input dimension; see Appendix A 1 for further model and training details. The offset in this case is simply $\bar{\phi}(c) = (c, 0)$ such that $X(\bar{\phi}(c)) \equiv \bar{\phi}_1(c) = c$, as required by the construction of Sec. IV.

For the purpose of this proof-of-principle study, we simply assume for the remainder of this work that $\rho_P \approx \rho$ for sufficiently large P and use Eqs. (5) and (6) instead of Eqs. (16) and (17). While the accuracy in reproducing the exact results in this case completely justifies this assumption, we emphasize that this is an approximation and one should generally extrapolate $P \rightarrow \infty$ more carefully. The particular value of P used here has been adopted from [22] for simplicity, and is observed to lead to reasonable results in the present setting. However, in general, it should be determined from a careful analysis of the trade-off between the accuracy of the DoS and the required computational effort, as larger values can lead to increasing autocorrelations in the case of HMC by forcing a reduction of the step size, as well as a degradation of the estimator variance in the case of our flow-based approach.

The DoS computed with flow-based sampling is compared against the exact result in Fig. 2, conclusively demonstrating the correctness of our approach across several orders of magnitude. Furthermore, the partition function $Z(h)$ is accurately reproduced, as shown in Fig. 3(a). In particular, the locations of the first two Lee-Yang zeroes can be clearly identified. They are associated with discontinuities

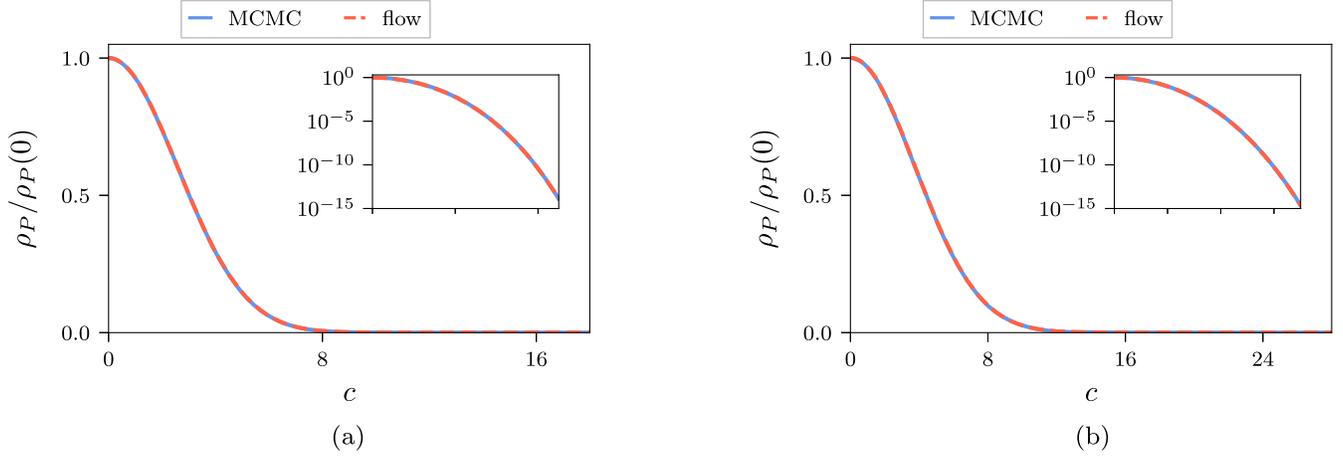


FIG. 4. Comparison of the flow results for the normalized DoS to the reconstructions from MCMC calculations for the one- (a) and two-dimensional (b) models with $P = 1000$.

in the average imaginary part of ϕ_1 , which is also accurately determined with our method as shown in Fig. 3(b).

B. One- and two-dimensional models

In order to verify that our approach also works beyond the rather trivial zero-dimensional setting, we consider actual low-dimensional lattice models of the two-component scalar field theory described above. Working in lattice units, the associated action in d dimensions is defined as

$$S(\phi) = \sum_{n \in \Lambda} \left(\frac{1}{2} \sum_{\mu=1}^d |\phi(n) - \phi(n + \hat{\mu})|^2 + \frac{m^2}{2} |\phi(n)|^2 + \frac{\lambda}{4} |\phi(n)|^4 + ih\phi_1(n) \right), \quad (24)$$

where $\phi(n) = (\phi_1(n), \phi_2(n))$, Λ is the set of all lattice sites given by integer-valued coordinates n , $\hat{\mu}$ denotes the unit vector in direction μ , and we assume periodic boundary conditions in all directions. The action for each individual site is essentially equivalent to the zero-dimensional model, differing only in the additional kinetic term. Accordingly, we can identify

$$X(\phi) \equiv \sum_{n \in \Lambda} \phi_1(n) \quad (25)$$

as well as

$$S_{c,P}(\phi) \equiv \sum_{n \in \Lambda} \left(\frac{1}{2} \sum_{\mu=1}^d |\phi(n) - \phi(n + \hat{\mu})|^2 + \frac{m^2}{2} |\phi(n)|^2 + \frac{\lambda}{4} |\phi(n)|^4 \right) + \frac{P}{2} \left(c - \sum_{n \in \Lambda} \phi_1(n) \right)^2 + \log \mathcal{N}. \quad (26)$$

We train flows using this target action with $m^2 = 1, \lambda = 1, P = 1000$ for one- and two-dimensional lattices of size 8 and 4×4 , respectively. To enforce equivariance under translations, the context functions s, t are implemented as convolutional neural networks with the condition c provided as an additional input channel. The offset in this case is $\bar{\phi}(c) = (c/|\Lambda|, 0)$, where $|\Lambda|$ denotes the total number of lattice sites, such that $X(\bar{\phi}(c)) = c$. In order to provide conventional baseline results, we employ HMC with the same target action and value for P , subsequently reconstructing ρ_P up to an overall factor using the numerical integration method described in Sec. IV; see Appendix A 2 for further details on model, training, and simulation.

Figure 4 compares $\rho_P(c)/\rho_P(0)$ (i.e. normalized to 1 at $c = 0$) obtained with flow-based sampling to the MCMC baseline. Similar to the zero-dimensional case, the results accurately reproduce the conventional computation, thereby confirming that our approach also works here as intended. We emphasize again that for the MCMC baseline, the reconstruction of the DoS at some point $c \neq 0$ by numerical integration requires precise knowledge of $\partial_c \log \rho_P(c)$ from 0 to c . In contrast, with the flow-based approach, the DoS can be independently probed at arbitrary points because it is computed directly.

VI. SUMMARY AND OUTLOOK

In this paper, we apply flow-based sampling to the DoS approach to complex action problems. Specifically, we show that normalizing flows can be used to compute the DoS directly, thereby disposing of the need to reconstruct it from measurements of a derivative quantity through MCMC calculations. We demonstrate this method in the context of simple models with imaginary external fields, confirming the correctness and accuracy of our approach.

Due to the conceptual and practical differences between the flow-based and conventional strategies, an in-depth comparison of the computational cost is not straightforward

and beyond the scope of this proof-of-principle study. We note that reaching the same level of accuracy in the final result using the flow-based approach is significantly cheaper in practice on the same hardware for the particular problem setting investigated in this work. This is likely due to the embarrassingly parallel sampling instead of the sequential evaluation of Markov chains. Furthermore, because of the numerical integration, the conventional ansatz may generally require higher precision in order to achieve an accurate reconstruction, whereas with the flow-based method the DoS can be directly probed at arbitrary points. However, it is unclear *a priori* how the upfront cost of training the flow compares against thermalizing the Markov chains. In general, the overall scaling of the computational cost required to achieve constant performance at different lattice volumes and action parameters depends strongly on the particular target and flow architecture, as well as the anticipated number of generated field configurations. Disentangling the various contributions to the total cost and defining thresholds for performance gains is a highly nontrivial task; see also [56] for a detailed discussion of this issue. For DoS calculations as considered in the present work, the strong localization of the target density with increasing values of the parameter P is expected to further increase the difficulty of the modeling task, thereby compounding the typically already poor volume scaling. Nevertheless—independently of how the cost actually scales in practice—depending on the particular scientific goals of such a calculation, the intrinsic advantages of the flow-based approach described in this work may well justify any additional expenses, and motivate further exploration in this direction.

In the future, we are interested in extending the present approach to higher dimensions, larger volumes, and fields with more components. This may be informative for the study of an approximate model of QCD near the second order phase transition where the external field plays the role of the quark mass [57]. In particular, computing the DoS could help to constrain the location of the Lee-Yang edge singularity. In this context, it may also be worthwhile to implement equivariance of the flow under the residual $O(N - 1)$ symmetry in order to better match the symmetries of the target distribution. Further interesting avenues include the relativistic Bose gas at finite chemical potential [58–60] as well as the application to gauge theories via gauge-equivariant flows [43,44], such as e.g. $U(1)$ gauge theory with a topological term [25] or QCD in the heavy-dense limit [27].

Beyond the aim to resolve complex action problems, the DoS can of course also be employed to compute observables for theories with purely real actions. While (unless one is interested in certain thermodynamic quantities) there is typically no reason to replace a standard MCMC algorithm with a generically much more expensive method, such an approach may be particularly useful for the treatment of ergodicity problems [61], since the target

distribution can be mapped out explicitly in regions of configuration space that are otherwise pathologically undersampled. Hence, the approach presented in this work also constitutes a promising ansatz for circumventing issues such as topological freezing and critical slowing down via flow-based methods, complementary to the more commonly investigated independence as well as hybrid sampling strategies.

ACKNOWLEDGMENTS

This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2181/1–390900948 (the Heidelberg STRUCTURES Excellence Cluster) and the Collaborative Research Centre SFB 1225 (ISOQUANT). J. M. U. is supported in part by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under grant Contract No. DE-SC0011090. This work is funded by the U.S. National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, [62]). The code used for the results in this paper is based on [63].

APPENDIX: IMPLEMENTATION DETAILS

1. Zero-dimensional model

The flow for the zero-dimensional model consists of 16 affine coupling layers with context functions being fully connected networks featuring three hidden layers with 64 neurons each. As activation functions we choose LeakyReLU [64] between layers together with a tanh activation after the final layer. Each network has two input neurons, one for the frozen variable (either ϕ_1 or ϕ_2 depending on the layer) and the other one for the condition c ; as well as two output neurons providing the values for s , t in Eq. (11). For the training, we apply the Adam optimizer with a learning rate of $1e-3$ and a batch size of 10k, with a total of 5k gradient updates. In order to compute $\rho(c)$, 10k samples are drawn for each value of c , with a spacing of $\Delta c = 0.01$.

2. One- and two-dimensional models

For the one- and two-dimensional models, the flow also consists of 16 affine coupling layers. Here, the context functions are convolutional neural networks featuring two hidden layers with eight channels each, using intermediate LeakyReLU activations and a final tanh activation as in the zero-dimensional case. Each network has three input channels, two for the frozen subsets of ϕ_1 , ϕ_2 (determined by alternating checkerboard masking, see e.g. [63]) and one for the condition c ; as well as two output channels providing the values for s , t in Eq. (11). The conditional input is constructed with the same dimensions as a single component of ϕ with c evenly distributed across all sites, i.e. with values of $c/|\Lambda|$ on each site where $|\Lambda|$ is the total

number of lattice points. For the training, we apply the Adam optimizer with a learning rate of $1e-3$ and a batch size of 1k, with a total of 50k gradient updates. In order to compute $\rho_P(c)$, 1×10^7 samples are drawn for each value of c , again with a spacing of $\Delta c = 0.01$.

For the conventional MCMC calculations, we use HMC with a step size of 0.02 and 50 steps per trajectory for the one-dimensional as well as a step size of 0.01 and 100 steps for the two-dimensional case. This results in acceptance rates of roughly 60–90%, with the highest values generally

observed around $c = 0$ and decreasing rates for larger c . For each c , we run 10k Markov chains in parallel, where in each chain the first 1k steps are discarded for equilibration. Subsequently, the chains are evaluated for 100k steps and every 10th configuration is recorded, resulting in a total of 1×10^8 configurations for each value of c , using a spacing of $\Delta c = 0.01$ as before. As described in the main text, $\rho_P(c)/\rho_P(0)$ is reconstructed from $\partial_c \log(\rho_P(c)/\rho_P(0))$ using the trapezoidal rule and exponentiating the resulting values for $\log(\rho_P(c)/\rho_P(0))$.

-
- [1] C. Morningstar, in *Proceedings of the 21st Annual Hampton University Graduate Studies Program (HUGS 2006)* (2007), [arXiv:hep-lat/0702020](#).
- [2] R. C. Brower, A. Hasenfratz, E. T. Neil, S. Catterall, G. Fleming, J. Giedt, E. Rinaldi, D. Schaich, E. Weinberg, and O. Witzel (USQCD Collaboration), *Eur. Phys. J. A* **55**, 198 (2019).
- [3] C. Lehner *et al.* (USQCD Collaboration), *Eur. Phys. J. A* **55**, 195 (2019).
- [4] A. S. Kronfeld, D. G. Richards, W. Detmold, R. Gupta, H.-W. Lin, K.-F. Liu, A. S. Meyer, R. Sufian, and S. Syritsyn (USQCD Collaboration), *Eur. Phys. J. A* **55**, 196 (2019).
- [5] V. Cirigliano, Z. Davoudi, T. Bhattacharya, T. Izubuchi, P. E. Shanahan, S. Syritsyn, and M. L. Wagman (USQCD Collaboration), *Eur. Phys. J. A* **55**, 197 (2019).
- [6] W. Detmold, R. G. Edwards, J. J. Dudek, M. Engelhardt, H.-W. Lin, S. Meinel, K. Orginos, and P. Shanahan (USQCD Collaboration), *Eur. Phys. J. A* **55**, 193 (2019).
- [7] A. Bazavov, F. Karsch, S. Mukherjee, and P. Petreczky (USQCD Collaboration), *Eur. Phys. J. A* **55**, 194 (2019).
- [8] M. Mathur and T. P. Sreeraj, *Phys. Rev. D* **94**, 085029 (2016).
- [9] B. Joó, C. Jung, N. H. Christ, W. Detmold, R. Edwards, M. Savage, and P. Shanahan (USQCD Collaboration), *Eur. Phys. J. A* **55**, 199 (2019).
- [10] G. Aarts, *J. Phys. Conf. Ser.* **706**, 022004 (2016).
- [11] C. Gattringer and K. Langfeld, *Int. J. Mod. Phys. A* **31**, 1643007 (2016).
- [12] G. Bhanot, S. Black, P. Carter, and R. Salvador, *Phys. Lett. B* **183**, 331 (1987).
- [13] G. Bhanot, K. Bitar, and R. Salvador, *Phys. Lett. B* **188**, 246 (1987).
- [14] G. Bhanot, K. Bitar, S. Black, P. Carter, and R. Salvador, *Phys. Lett. B* **187**, 381 (1987).
- [15] G. Bhanot, R. Salvador, S. Black, P. Carter, and R. Toral, *Phys. Rev. Lett.* **59**, 803 (1987).
- [16] G. Bhanot, A. Gocksch, and P. Rossi, *Phys. Lett. B* **199**, 101 (1987).
- [17] F. Wang and D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [18] Z. Fodor, S. D. Katz, and C. Schmidt, *J. High Energy Phys.* **03** (2007) 121.
- [19] K. Langfeld, B. Lucini, and A. Rago, *Phys. Rev. Lett.* **109**, 111601 (2012).
- [20] K. Langfeld and J. M. Pawłowski, *Phys. Rev. D* **88**, 071502 (R) (2013).
- [21] K. Langfeld, B. Lucini, R. Pellegrini, and A. Rago, *Eur. Phys. J. C* **76**, 306 (2016).
- [22] S. Borsanyi and D. Sexty, *Phys. Lett. B* **815**, 136148 (2021).
- [23] K. Langfeld and B. Lucini, *Phys. Rev. D* **90**, 094502 (2014).
- [24] C. Gattringer and P. Törek, *Phys. Lett. B* **747**, 545 (2015).
- [25] C. Gattringer, M. Giuliani, A. Lehmann, and P. Törek, *Proc. Sci. LATTICE2015* (2016) 194.
- [26] K. Langfeld, *Proc. Sci. LATTICE2016* (2017) 010.
- [27] N. Garron and K. Langfeld, *Eur. Phys. J. C* **76**, 569 (2016).
- [28] M. Körner, K. Langfeld, D. Smith, and L. von Smekal, *Phys. Rev. D* **102**, 054502 (2020).
- [29] B. Lucini, O. Francesconi, M. Holzmann, D. Lancaster, and A. Rago, *J. Phys. Conf. Ser.* **2207**, 012052 (2022).
- [30] K. A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller, and P. Kessel, *Phys. Rev. E* **101**, 023304 (2020).
- [31] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, *Phys. Rev. Lett.* **126**, 032001 (2021).
- [32] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, *Proc. Sci. LATTICE2021* (2022) 338.
- [33] E. G. Tabak and E. Vanden-Eijnden, *Commun. Math. Sci.* **8**, 217 (2010).
- [34] E. G. Tabak and C. V. Turner, *Commun. Pure Appl. Math.* **66**, 145 (2013).
- [35] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](#).
- [36] L. Dinh, J. Sohl-Dickstein, and S. Bengio, [arXiv:1605.08803](#).
- [37] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [38] M. S. Albergo, G. Kanwar, and P. E. Shanahan, *Phys. Rev. D* **100**, 034515 (2019).
- [39] D. C. Hackett, C.-C. Hsieh, M. S. Albergo, D. Boyda, J.-W. Chen, K.-F. Chen, K. Cranmer, G. Kanwar, and P. E. Shanahan, [arXiv:2107.00734](#).
- [40] L. Del Debbio, J. Marsh Rossney, and M. Wilson, *Phys. Rev. D* **104**, 094507 (2021).
- [41] M. Caselle, E. Cellini, A. Nada, and M. Panero, *J. High Energy Phys.* **07** (2022) 015.

- [42] A. G. D. G. Matthews, M. Arbel, D. J. Rezende, and A. Doucet, [arXiv:2201.13117](https://arxiv.org/abs/2201.13117).
- [43] G. Kanwar, M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, S. Racanière, D. J. Rezende, and P. E. Shanahan, *Phys. Rev. Lett.* **125**, 121601 (2020).
- [44] D. Boyda, G. Kanwar, S. Racanière, D. J. Rezende, M. S. Albergo, K. Cranmer, D. C. Hackett, and P. E. Shanahan, *Phys. Rev. D* **103**, 074504 (2021).
- [45] S. Foreman, X.-Y. Jin, and J. C. Osborn, in *Proceedings of the 9th International Conference on Learning Representations* (2021), [arXiv:2105.03418](https://arxiv.org/abs/2105.03418).
- [46] S. Foreman, X.-Y. Jin, and J. C. Osborn, *Proc. Sci. LATTICE2021* (**2022**) 508.
- [47] S. Foreman, T. Izubuchi, L. Jin, X.-Y. Jin, J. C. Osborn, and A. Tomiya, *Proc. Sci. LATTICE2021* (**2022**) 073.
- [48] M. S. Albergo, G. Kanwar, S. Racanière, D. J. Rezende, J. M. Urban, D. Boyda, K. Cranmer, D. C. Hackett, and P. E. Shanahan, *Phys. Rev. D* **104**, 114507 (2021).
- [49] J. Finkenrath, [arXiv:2201.02216](https://arxiv.org/abs/2201.02216).
- [50] M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, G. Kanwar, S. Racanière, D. J. Rezende, F. Romero-López, P. E. Shanahan, and J. M. Urban, *Phys. Rev. D* **106**, 014514 (2022).
- [51] S. Lawrence and Y. Yamauchi, *Phys. Rev. D* **103**, 114509 (2021).
- [52] M. Rodekamp, E. Berkowitz, C. Gäntgen, S. Krieg, T. Luu, and J. Ostmeyer, *Phys. Rev. B* **106**, 125139 (2022).
- [53] T. D. Lee and C.-N. Yang, *Phys. Rev.* **87**, 410 (1952).
- [54] C. Jarzynski, *Phys. Rev. Lett.* **78**, 2690 (1997).
- [55] H. Wu, J. Köhler, and F. Noé, *Adv. Neural Inf. Process. Syst.* **33**, 5933 (2020).
- [56] R. Abbott *et al.*, [arXiv:2211.07541](https://arxiv.org/abs/2211.07541).
- [57] F. Attanasio, M. Bauer, L. Kades, and J. M. Pawłowski, *Proc. Sci. LATTICE2021* (**2022**) 223.
- [58] G. Aarts, *Phys. Rev. Lett.* **102**, 131601 (2009).
- [59] L. Bongiovanni, K. Langfeld, B. Lucini, R. Pellegrini, and A. Rago, [arXiv:1601.02929](https://arxiv.org/abs/1601.02929).
- [60] O. Francesconi, M. Holzmann, B. Lucini, and A. Rago, *Phys. Rev. D* **101**, 014504 (2020).
- [61] K. Langfeld, P. Buividovich, P. E. L. Rakow, and J. Roscoe, *Phys. Rev. E* **106**, 054139 (2022).
- [62] <http://iaifi.org/>.
- [63] M. S. Albergo, D. Boyda, D. C. Hackett, G. Kanwar, K. Cranmer, S. Racanière, D. J. Rezende, and P. E. Shanahan, [arXiv:2101.08176](https://arxiv.org/abs/2101.08176).
- [64] B. Xu, N. Wang, T. Chen, and M. Li, [arXiv:1505.00853](https://arxiv.org/abs/1505.00853).