Flow-based sampling for fermionic lattice field theories

Michael S. Alberg^{1,*}, Gurtej Kanwar^{2,3,†}, Sébastien Racanière^{4,‡}, Danilo J. Rezende^{4,§}, Julian M. Urban^{5,||},
Denis Boyda^{6,2,3}, Kyle Cranmer¹, Daniel C. Hackett^{2,3}, and Phiala E. Shanahan^{2,3}

¹*Center for Cosmology and Particle Physics, New York University, New York, New York 10003, USA*

²*Center for Theoretical Physics, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139, USA*

³*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions,
Cambridge, Massachusetts 02139-4307, USA*

⁴*DeepMind, London NIC 4DJ, United Kingdom*

⁵*Institut für Theoretische Physik, Universität Heidelberg,
Philosophenweg 16, 69120 Heidelberg, Germany*

⁶*Argonne Leadership Computing Facility, Argonne National Laboratory, Lemont, Illinois 60439, USA*

(Received 22 June 2021; accepted 12 November 2021; published 15 December 2021)

Algorithms based on normalizing flows are emerging as promising machine learning approaches to sampling complicated probability distributions in a way that can be made asymptotically exact. In the context of lattice field theory, proof-of-principle studies have demonstrated the effectiveness of this approach for scalar theories, gauge theories, and statistical systems. This work develops approaches that enable flow-based sampling of theories with dynamical fermions, which is necessary for the technique to be applied to lattice field theory studies of the Standard Model of particle physics and many condensed matter systems. As a practical demonstration, these methods are applied to the sampling of field configurations for a two-dimensional theory of massless staggered fermions coupled to a scalar field via a Yukawa interaction.

DOI: [10.1103/PhysRevD.104.114507](https://doi.org/10.1103/PhysRevD.104.114507)

I. INTRODUCTION

Lattice field theory is among the most successful approaches for regularizing and computing path integral expectation values in quantum field theory. In particular, the path integral can be numerically evaluated by formulating a stochastic process weighted by the Euclidean lattice action and applying Markov Chain Monte Carlo (MCMC) sampling [1]. The Euclidean lattice results can then be systematically related to the corresponding continuum Minkowski theory. This procedure enables the investigation of equilibrium properties in the theory of interest and is a powerful and well-established method to study strongly coupled quantum field theories nonperturbatively. Key areas of application include fundamental interactions, most

prominently quantum chromodynamics (QCD), as well as problems in condensed matter theory; see Refs. [2–9] for recent reviews.

The sequential nature of the Markov chain is a potential drawback to the MCMC sampling approach for computing path integrals in lattice field theory. In particular, known Markov chain update schemes for many theories of interest are local or diffusive, which can result in severe autocorrelations between successive elements of the chain. Naturally, the larger the autocorrelation between samples, the more samples must be drawn to achieve a result at fixed statistical precision. Close to criticality, e.g., when approaching the continuum limit of lattice field theories, autocorrelations also diverge rapidly for such local or diffusive Markov chains. This issue, referred to as critical slowing down, can render computations prohibitively expensive [10–12].

These challenges have motivated extensive work to replace local/diffusive MCMC algorithms, such as Hybrid Monte Carlo (HMC) [13], with other sampling procedures. Specialized Markov chain steps have been developed in a number of specific contexts, including cluster updates [14–22], worm algorithms [23–25], sampling in terms of dual variables [26–28], and event-chain algorithms [29–33]. Though these methods have been shown to mitigate critical slowing down in some settings, they cannot be applied to many theories of interest, including lattice QCD. In this light,

*alberg@nyu.edu

†gurtej@mit.edu

‡sracaniere@google.com

§danilor@google.com

||urban@thphys.uni-heidelberg.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). 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.

development of sampling algorithms for lattice field theory based on machine learning techniques is underway, and previous works have applied a variety of tools such as adversarial learning and self-learning Monte Carlo methods [34–47]. Progress has recently been made in substituting the proposal mechanism in MCMC methods with a variational ansatz based on a class of samplers known as normalizing flows [48–50], which can be optimized to approximately sample from the target Boltzmann distribution [51–53]. Within this approach, proposed samples are by construction uncorrelated, and asymptotic exactness can be guaranteed by implementing a Markov chain with a Metropolis accept/reject step or through reweighting.

Though flow-based models have been extended to exactly incorporate the gauge symmetries inherent in many quantum field theories [54,55], existing applications are so far restricted to purely bosonic theories. For theories involving fermions, the anticommuting nature of the associated operators must be treated by formulating the theory in terms of integrals over Grassmann-valued field variables. These integrals may be evaluated analytically, resulting in a purely bosonic theory described by an effective action which incorporates the dynamics of the fermion fields via “fermion determinant” terms. While flow-based methods may in principle be applied to learn this effective action over bosonic fields, the cost of computing such determinants scales unfavorably with the number of fermionic degrees of freedom, and their exact evaluation is typically intractable at the scale of state-of-the-art calculations.

In this paper, a framework is presented for the application of flow-based sampling algorithms to lattice quantum field theories with dynamical fermions, such as lattice QCD and theories describing many condensed matter systems. We construct approaches based primarily on the pseudofermion method [56] to avoid an explicit computation of the fermion determinant while guaranteeing asymptotic exactness of the sampling schemes. We investigate the application of a number of flow-based samplers in the context of a simple, two-dimensional Yukawa model with a real scalar field coupled to staggered fermions. Our results demonstrate that lattice field theories with dynamical fermions are amenable to flow-based sampling and provide a starting point for extensions to higher-dimensional settings as well as theories involving gauge fields. The architectures developed here are also applicable to flow-based acceleration of traditional MCMC approaches [57–61]. The primary contributions of this work are these:

- (i) In Sec. III, identifying four distinct sampling schemes based on generative models that capture different decompositions/marginalizations of the target distribution over boson and pseudofermion field configurations.
- (ii) In Sec. IV, constructing and optimizing efficient, expressive flow model architectures that respect the symmetries of the pseudofermion action, in

particular translational symmetry with antiperiodic temporal boundary conditions.

- (iii) In Sec. V, implementing and numerically benchmarking these sampling approaches in the context of a two-dimensional field theory with one pair of mass-degenerate fermions.

The remainder of this paper is organized as follows. In Sec. II, we review the description of fermions in lattice field theory, the use of pseudofermions, and the boundary conditions and translational symmetry of pseudofermion actions. In Sec. III, we outline four exact generative sampling schemes for fermionic theories and subsequently develop suitable flow architectures as the generative models for use in these sampling schemes in Sec. IV. Details and numerical results of the application of our framework to a Yukawa theory in two dimensions are presented in Sec. V. In Sec. VI, we discuss the applicability of our developments to update-based approaches. We summarize our findings and provide an outlook in Sec. VII.

II. FERMIONS ON THE LATTICE

The simulation of dynamical fermion degrees of freedom in lattice field theory is a highly nontrivial task for many theories of physical interest, both conceptually and computationally. In this section, we briefly review the main concepts behind formulations of lattice fermions and their numerical implementation. For a comprehensive treatment we refer the reader to one of the standard textbooks; see e.g., Refs. [62–64].

A. Path integrals with fermions

We consider field theories of interacting fermionic and bosonic degrees of freedom discretized on a d -dimensional Euclidean hypercubic lattice with periodic boundary conditions. The action of such a theory can be expressed as

$$S(\psi, \bar{\psi}, \phi) = S_B(\phi) + S_F(\psi, \bar{\psi}, \phi), \quad (1)$$

where the subscripts B and F denote the bosonic and fermionic contributions to the action, the discretized boson field variables are collectively denoted by ϕ , and the discretized fermion field variables are denoted by $\psi, \bar{\psi}$. For the present work, we assume that the fermionic action is bilinear in N_f flavors of Dirac fermions $\psi_f, \bar{\psi}_f$ and is given by

$$S_F(\psi, \bar{\psi}, \phi) = \sum_{f=1}^{N_f} \bar{\psi}_f D_f(\phi) \psi_f, \quad (2)$$

where the Dirac operator $D_f(\phi)$ includes the kinetic terms, mass terms, and any coupling to bosonic fields for each fermion flavor f . The precise form of $D_f(\phi)$ is determined by the theory of interest and the choice of discretization.

Expectation values of observables $\mathcal{O}$ are computed via path integrals of the form

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}[\phi] \mathcal{D}[\psi, \bar{\psi}] e^{-S_F(\psi, \bar{\psi}, \phi)} e^{-S_B(\phi)} \mathcal{O}(\psi, \bar{\psi}, \phi), \quad (3)$$

where

$$Z = \int \mathcal{D}[\phi] \mathcal{D}[\psi, \bar{\psi}] e^{-S_F(\psi, \bar{\psi}, \phi)} e^{-S_B(\phi)}, \quad (4)$$

and the fermion fields ψ and $\bar{\psi}$ are defined in terms of anticommuting Grassmann numbers. For bilinear actions of the form given in Eq. (2), integration over the Grassmann-valued fermion fields can be performed explicitly, giving

$$\int \mathcal{D}[\psi, \bar{\psi}] e^{-S_F(\psi, \bar{\psi}, \phi)} = \prod_{f=1}^{N_f} \det D_f(\phi). \quad (5)$$

By applying Wick's theorem, the dependence of the observable on the fermions can be integrated out. Path integral expectation values can then be written in terms of purely bosonic degrees of freedom as

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}[\phi] \left[\prod_{f=1}^{N_f} \det D_f(\phi) \right] e^{-S_B(\phi)} \mathcal{O}(\phi). \quad (6)$$

This expectation value may be estimated via MCMC sampling by computing an average over a statistical ensemble of configurations ϕ ,

$$\langle \mathcal{O} \rangle \approx \frac{1}{N} \sum_{\phi \sim p} \mathcal{O}(\phi), \quad (7)$$

where $\sum_{\phi \sim p}$ denotes a sum over N configurations sampled from the probability distribution

$$p(\phi) = \frac{1}{Z} e^{-S_B(\phi)} \prod_{f=1}^{N_f} \det D_f(\phi). \quad (8)$$

Direct sampling schemes for high-dimensional lattice distributions are typically not known, even for theories without fermions. Nevertheless, the distribution $p(\phi)$ can be sampled via MCMC methods with guaranteed asymptotic exactness under certain ergodicity and balance constraints [65]. Among these, the HMC algorithm [13] has been established as the *de facto* standard method for producing configurations in lattice field theory and is routinely employed in state-of-the-art QCD studies and beyond. This algorithm is based on the numerical treatment of Hamiltonian equations of motion in a fictitious time dimension, where quantum fluctuations are encoded by the random sampling of the associated canonical momenta. Given a field configuration and a set of momenta, this evolution is computed with a symplectic integrator such as the leapfrog algorithm. A Metropolis-Hastings accept/reject step [66,67]

results in an algorithm satisfying detailed balance, despite the accumulation of numerical errors along the discretized integration trajectory. At scale, the fermion determinants are treated stochastically via the pseudofermion method introduced in the following section.

B. Pseudofermions

The fermion determinants in Eq. (8) cannot be calculated directly at scale because the Dirac matrices D_f are high-dimensional. For d -dimensional field configurations with L sites per spatial dimension and L_t sites in the temporal dimension, the total number of fermionic degrees of freedom scales as the total number of lattice sites, $V = L_t L^{d-1}$. Each Dirac matrix D_f then has dimensions $\mathcal{O}(V \times V)$, and an exact computation of the determinants becomes intractable at the scales of many theories of interest.

Instead, Gaussian integrals over auxiliary bosonic fields can be used to replace the direct evaluation of determinant factors based on the identity

$$\det \mathcal{M} = \frac{1}{Z_N} \int \mathcal{D}[\varphi_R, \varphi_I] e^{-\varphi^\dagger \mathcal{M}^{-1} \varphi}, \quad (9)$$

where the normalization constant Z_N is defined as

$$Z_N = \int \mathcal{D}[\varphi_R, \varphi_I] e^{-\varphi^\dagger \varphi}. \quad (10)$$

Here, φ_R, φ_I denote the real and imaginary components of the auxiliary complex field φ , and the matrix $\mathcal{M}$ must be positive definite. Since the Dirac matrices D_f are typically not positive definite, one cannot directly apply this identity to each factor of $\det D_f$. However, for fermion flavors f_1 and f_2 appearing as degenerate pairs, based on γ_5 Hermiticity one can instead use the equality

$$\det D_{f_1} \det D_{f_2} = \det D_{f_1} D_{f_1}^\dagger, \quad (11)$$

and then apply Eq. (9) to the positive-definite matrix $\mathcal{M} = D_{f_1} D_{f_1}^\dagger$. For fermion flavors f not included in any degenerate pair, one can apply one-flavor algorithms [68–70] to replace D_f with a positive-definite matrix $\mathcal{M}$ capturing identical dynamics.

Using the pseudofermion approach, a path integral as in Eq. (6) can thus be rewritten in terms of an action involving the auxiliary pseudofermion fields φ ,

$$S(\phi, \varphi) = S_B(\phi) + S_{PF}(\phi, \varphi) \quad \text{with}$$

$$S_{PF}(\phi, \varphi) = \varphi^\dagger \mathcal{M}^{-1}(\phi) \varphi \equiv \sum_{k=1}^{N_{pf}} \varphi_k^\dagger \mathcal{M}_k^{-1}(\phi) \varphi_k, \quad (12)$$

after replacing the fermion determinants in the given lattice theory by the determinants of N_{pf} positive-definite matrices $\mathcal{M}_k$ as

$$\prod_{f=1}^{N_f} \det D_f(\phi) = \prod_{k=1}^{N_{\text{pf}}} \det \mathcal{M}_k(\phi). \quad (13)$$

Each term $\varphi_k^\dagger \mathcal{M}_k^{-1} \varphi_k$ in the pseudofermion action can be efficiently computed using iterative solvers such as the conjugate gradient method. Having formulated the theory using pseudofermions in Eq. (12), evaluation of the path integral via MCMC methods can then be performed in this augmented space by sampling from the joint distribution

$$p(\phi, \varphi) = \frac{1}{Z} e^{-S_B(\phi) - S_{PF}(\phi, \varphi)}. \quad (14)$$

The distribution $p(\phi)$ given in Eq. (8) is obtained from this joint distribution by marginalizing over the auxiliary pseudofermion fields φ .

While the pseudofermion method renders the treatment of fermion determinants tractable in principle, the joint action may strongly fluctuate in certain limits. This feature can slow down MCMC sampling of the joint distribution and can lead to an unfavorable volume scaling of the associated computational effort, especially when many components of the bosonic field are updated simultaneously. Accordingly, numerous modifications of the pseudofermion formulation have been developed to improve the structure of the action; see e.g., Refs. [71–76]. These developments are complementary to the application of generative models for sampling the pseudofermion distribution and could be combined with any of the approaches presented here. For example, in this work we improve the efficiency of modeling and sampling the distributions under study by applying even-odd preconditioning [77], as discussed in Sec. VA and Appendix B.

C. Boundary conditions and translational symmetry

In lattice studies of purely bosonic theories, it is common to choose periodic boundary conditions in all directions of the lattice, allowing the incorporation of an exact discrete translational symmetry in lattice actions for such theories.

For fermion fields, one needs to impose antiperiodic boundary conditions in the time direction in order to obtain a consistent definition of the trace for the Euclidean partition function. Actions for theories involving fermionic fields are then invariant under simultaneous spatiotemporal translations of ϕ , ψ , and $\bar{\psi}$, with appropriate boundary conditions applied for each field.

To be consistent with the boundary conditions for ψ and $\bar{\psi}$, each Dirac matrix $D_f(\phi)$ must include appropriate signs for any terms coupling fields across the temporal boundary. As a result, these boundary conditions affect the pseudofermion formulation of the theory as well, and the pseudofermion action $S_{PF}(\phi, \varphi)$ is invariant under simultaneous translations of ϕ and φ with antiperiodic temporal boundary conditions applied to φ .

In general, the discretization chosen for the Dirac operator determines which particular lattice translations are included in the translational symmetry group. In the staggered formulation [78], for example, the spinor components of each flavor of fermion are distributed over the components of hypercubes with 2^d sites each, and the translational symmetry group includes all translations by an even number of sites. Translations by an odd number of sites in particular directions correspond to more complicated internal symmetry transformations that mix spinor degrees of freedom, and must involve sign flips on specific field components to leave the staggered action invariant [79]. These translational symmetries of the staggered formulation play a role in the application to the staggered-fermion Yukawa model presented in Sec. V.

III. EXACT GENERATIVE SAMPLING SCHEMES FOR FERMIONIC THEORIES

Generating importance-weighted field configurations for a lattice field theory involving fermions can proceed via the marginal distribution $p(\phi)$ defined in Eq. (8), the joint distribution $p(\phi, \varphi)$ defined in Eq. (14), or through other choices of marginalized distributions defined in Table I. In this work, we develop exact sampling schemes based on generative models that directly approximate these

TABLE I. List of possible distributions derived from the joint target density in Eq. (14). The normalizing constant Z is given by Eq. (4) and Z_N is defined in Eq. (10). Notes: (A) only the joint, φ -conditional, and ϕ -conditional densities can be efficiently computed (up to normalization). (B) The φ conditional can be sampled exactly by the method specified in Eq. (16). (C) A closed form for the φ -marginal density is not generally known (even unnormalized).

Name	Probability density	Use case
Joint ^A	$p(\phi, \varphi) = \frac{1}{Z} \exp(-S_B(\phi) - \varphi^\dagger [\mathcal{M}(\phi)]^{-1} \varphi)$	Sec. III D
ϕ marginal	$p(\phi) = \frac{Z_N}{Z} \exp(-S_B(\phi)) \det \mathcal{M}(\phi)$	Secs. III A and III C
φ conditional ^{A,B}	$p(\varphi \phi) = \frac{1}{Z_N \det \mathcal{M}(\phi)} \exp(-\varphi^\dagger [\mathcal{M}(\phi)]^{-1} \varphi)$	Secs. III A, III B, and III C
φ marginal ^C	$p(\varphi) = \frac{1}{Z} \int d\phi \exp(-S_B(\phi) - \varphi^\dagger [\mathcal{M}(\phi)]^{-1} \varphi)$	...
ϕ conditional ^A	$p(\phi \varphi) = \frac{\exp(-S_B(\phi) - \varphi^\dagger [\mathcal{M}(\phi)]^{-1} \varphi)}{\int d\phi \exp(-S_B(\phi) - \varphi^\dagger [\mathcal{M}(\phi)]^{-1} \varphi)}$	Sec. III B

FIG. 1. Diagrams illustrating the four types of sampling schemes described in Sec. III. Blue circles/ellipses depict the current state of the Markov chain. Yellow boxes depict exactly sampleable densities either produced from generative models or by Equation (16). Green boxes correspond to Metropolis accept/reject steps using the acceptance probabilities defined in the text. Dotted lines indicate the Markov chain, whereas solid lines correspond to the internal operations of each Markov chain step.

distributions. In defining these sampling schemes, we assume that the model probability density may be computed for each generative model (this property holds for the flow-based models defined below). This section details four asymptotically exact schemes for constructing Markov chains to draw samples of ϕ , as illustrated in Fig. 1.

A. Modeling and sampling of $p(\phi)$

Since we must ultimately sample only the field ϕ , one could directly model the ϕ -marginal distribution (row 2 of Table I) by constructing a generative sampler providing a distribution $q(\phi)$ approximating $p(\phi)$. Samples drawn from the model distribution $q(\phi)$ can be used in asymptotically exact sampling schemes by either constructing an independence Metropolis Markov chain or applying reweighting/resampling based on reweighting factors $p(\phi)/q(\phi)$. A direct application of either approach requires computing $p(\phi)$ involving the aforementioned determinant factors. The acceptance probability for such a Metropolis Markov chain would be

$$A_M(\phi \rightarrow \phi') = \min \left(1, \frac{e^{-S_B(\phi')} \det \mathcal{M}(\phi') q(\phi)}{e^{-S_B(\phi)} \det \mathcal{M}(\phi) q(\phi')} \right). \quad (15)$$

This sampling scheme is illustrated in Fig. 1(a).

Instead of evaluating the ratio $\det \mathcal{M}(\phi')/\det \mathcal{M}(\phi)$ directly, which becomes prohibitive at scale, it is possible to apply the pseudomarginal method [80] using stochastic approximations of both the numerator and denominator of Eq. (15) in a way that retains asymptotic exactness. In this

stochastic generalization of the Metropolis algorithm, one computes an estimate of $p(\phi)$ using an unbiased stochastic estimator when ϕ is initially proposed. This estimate of $p(\phi)$ is then used in all subsequent accept/reject tests for the next element in the Markov chain. Applied to a theory with fermions, this amounts to computing a stochastic estimate of the fermion determinant for each proposed configuration.¹

For example, we can use an unbiased estimator based on pseudofermions. An (unnormalized) estimate for $p(\phi)$ can be obtained by generating a pseudofermion φ from the conditional $p(\varphi|\phi)$ and measuring the quantity $e^{-\varphi^\dagger (\mathcal{M}^{-1}(\phi) - 1) \varphi} e^{-S_B(\phi)}$. The φ conditional can be directly sampled according to

$$\varphi = \mathcal{A}(\phi) \chi, \quad \text{where } \chi \sim \frac{1}{Z_N} e^{-\chi^\dagger \chi}. \quad (16)$$

The matrix $\mathcal{A}$ is defined by the identity $\mathcal{M}(\phi) \equiv \mathcal{A}(\phi) \mathcal{A}^\dagger(\phi)$ and reduces to the Dirac matrix $D(\phi)$ in a two-flavor example. This estimate can be extremely noisy in practice and may give poor statistical performance. However, it can be improved upon by using multiple pseudofermion draws; see e.g., Refs. [74,76,82,83].

In principle, any unbiased stochastic estimator of the fermion determinant can be applied (whether based on

¹Note that the pseudomarginal algorithm is not equivalent to the sampling scheme based on stochastic estimates of ratios [62,81], for which asymptotic exactness has not been demonstrated.

pseudofermions or entirely distinct). The limit of taking arbitrarily precise estimators recovers the exact acceptance probability in Eq. (15). Acceptance rates obtained by using the exact form can thus be interpreted as an upper bound on sampling performance.

B. Gibbs sampling using $p(\phi|\varphi), p(\varphi|\phi)$

An alternative to modeling $p(\phi)$ directly is to construct samplers for both conditional distributions $p(\phi|\varphi)$ and $p(\varphi|\phi)$ and build an asymptotically exact Gibbs sampler that alternately samples from these distributions to update ϕ and φ . For such a Gibbs sampler to satisfy detailed balance, the update to ϕ must satisfy detailed balance for $p(\phi|\varphi)$ and the update to φ must satisfy detailed balance for $p(\varphi|\phi)$. The φ conditional can be exactly and directly sampled as described in Eq. (16), automatically fulfilling this requirement. On the other hand, the ϕ conditional (row 5 of Table I) can be approximated by a generative model distribution $q(\phi|\varphi) \approx p(\phi|\varphi)$. This model can be incorporated into an exact Markov chain transition for the ϕ -conditional distribution as follows. Start with a state ϕ , sample φ' from $p(\varphi|\phi)$ using Eq. (16), conditionally propose ϕ' from $q(\phi|\varphi')$, and then apply a Metropolis-Hastings accept/reject step with the acceptance probability given by

$$\begin{aligned} A_G(\phi \rightarrow \phi'|\varphi') &= \min \left(1, \frac{p(\phi'|\varphi') q(\phi|\varphi')}{p(\phi|\varphi') q(\phi'|\varphi')} \right), \\ &= \min \left(1, \frac{p(\phi', \varphi') q(\phi|\varphi')}{p(\phi, \varphi') q(\phi'|\varphi')} \right). \end{aligned} \quad (17)$$

This step satisfies detailed balance for the ϕ -conditional distribution $p(\phi|\varphi)$ as required and guarantees asymptotic exactness. Note that in contrast to the ϕ -marginal sampler described in Sec. III A, computing the acceptance probability at scale for the sampling scheme described here and in Secs. III C and III D does not rely on unbiased stochastic determinant estimators.

In this approach, the field φ' is independently resampled conditioned on ϕ at each step of the Markov chain, and therefore does not need to be stored. This Gibbs sampler can thus be interpreted as an exact Markov chain over ϕ alone, with the sampling of φ' contained inside each Markov chain step as depicted in Fig. 1(b). The approach closely mirrors the typical sampling strategy employed in HMC, in which pseudofermions φ' are sampled according to the exact conditional distribution $p(\varphi|\phi)$ and Hamiltonian evolution is used to construct an update step that satisfies detailed balance for the conditional distribution $p(\phi|\varphi')$. The generative model proposal and Metropolis-Hastings step for $p(\phi|\varphi')$ can thus be considered an optimizable replacement of the molecular dynamics trajectory utilized in HMC, with the difference that the mechanism of generating a proposal configuration ϕ' does not directly depend on ϕ (as is the case

for a symplectic integrator), but only indirectly through φ' . However, this also means that in contrast to all other schemes described here, the Gibbs sampler is not an independence sampler. Drawing configurations from the model and constructing the Markov chain cannot be done asynchronously, since the generation of a proposal explicitly depends on the previous element of the chain.

C. Autoregressive modeling and sampling of $p(\phi, \varphi)$

The joint distribution $p(\phi, \varphi)$ can be autoregressively decomposed as the product $p(\phi, \varphi) = p(\phi)p(\varphi|\phi)$ in terms of the ϕ marginal and φ conditional (rows 2 and 3 of Table I). A generative model for the joint distribution could therefore be produced by approximating both components independently, i.e., $q(\phi, \varphi) = q(\phi)q(\varphi|\phi)$. This autoregressive decomposition allows the joint distribution to be reproduced in terms of two potentially simpler distributions. Note that although the exact sampling procedure described in Eq. (16) can be applied to draw samples from $p(\varphi|\phi)$, computing the normalizing constant of this φ -conditional distribution is not tractable. This is not an obstacle when one is only interested in conditionally sampling φ , as is the case for HMC or the approaches of Secs. III A and III B, but motivates modeling the distribution in the case where an approximation with a tractable density is required.

Exactness can be straightforwardly enforced in this approach by employing Markov chain steps in which joint samples (ϕ', φ') are proposed independently from $q(\phi, \varphi)$, and a Metropolis-Hastings accept/reject step is applied for the proposed transition $(\phi, \varphi) \rightarrow (\phi', \varphi')$ according to the acceptance probability

$$A_A(\phi, \varphi \rightarrow \phi', \varphi') = \min \left(1, \frac{p(\phi', \varphi') q(\phi)q(\varphi|\phi)}{p(\phi, \varphi) q(\phi')q(\varphi'|\phi')} \right). \quad (18)$$

This sampling scheme is illustrated in Fig. 1(c). Furthermore, unique reweighting factors can be tractably computed for each configuration ϕ as $p(\phi, \varphi)/q(\phi)q(\varphi|\phi)$, thus reweighting approaches may also be used as alternatives to MCMC methods in order to guarantee exactness here.

D. Fully joint modeling and sampling of $p(\phi, \varphi)$

Rather than modeling the factors $p(\phi)$ and $p(\varphi|\phi)$, one could instead apply generative models to jointly sample the fields ϕ and φ according to a distribution $q(\phi, \varphi)$ that directly approximates the joint distribution (row 1 of Table I). This results in joint samples and density estimates analogous to the autoregressive case above, but is a qualitatively distinct approach to modeling this distribution. Exactness can be enforced using a similar Metropolis-Hastings Markov chain transition with acceptance probability

$$A_J(\phi, \varphi \rightarrow \phi', \varphi') = \min \left(1, \frac{p(\phi', \varphi')}{p(\phi, \varphi)} \frac{q(\phi, \varphi)}{q(\phi', \varphi')} \right), \quad (19)$$

or by applying reweighting or direct resampling techniques. This approach is illustrated in Fig. 1(d).

IV. FERMIONIC FLOWS VIA PSEUDOFERMIONS

The sampling approaches discussed above for theories involving fermions can in principle use any generative models that enable both efficient sampling and density estimation for the relevant model distributions. Normalizing flows are one such class of probabilistic models for which these operations are made possible using a change-of-variables formula [48,84,85]. Flow-based methods have been successfully implemented to model the unnormalized Boltzmann distributions of ϕ^4 theory as well as $U(1)$ and $SU(N)$ gauge theories [51,52,54,55], opening up a wealth of potential applications in high energy and nuclear physics as well as condensed matter theory. For an in-depth introduction to normalizing flows for lattice field theory with further implementation details and explanations, we refer the interested reader to Ref. [86]. For a general introduction to machine learning for physicists, we recommend Ref. [87].

Below, we outline the key concepts behind flows relevant for our approach to modeling the distributions listed in Table I:

- (1) A flow-based model consists of an invertible “flow” f and a prior distribution with probability density $r(\cdot)$, which can together be used to produce samples by first drawing z according to $r(z)$, then returning $x = f(z)$.²
- (2) The Jacobian $J = \frac{\partial f}{\partial z}$, combined with the prior density, allows the output probability density q to be evaluated as

$$q(x) = r(z) \left| \det \frac{\partial f}{\partial z} \right|^{-1}. \quad (20)$$

- (3) The flow f is parametrized by free model parameters which may be optimized by minimizing a suitable “loss function” that quantifies the difference between the model distribution $q(x)$ and the target distribution. A common choice also employed in the present work is the Kullback-Leibler (KL) divergence [88].

The remainder of this section describes how models for each of the distributions required for sampling may be

constructed. First, a common training procedure for all such models is described in Sec. IV A based on the idea that each distribution aims to approximate some marginalization of the same joint distribution $p(\phi, \varphi)$. This common training procedure motivates some of the architectural decisions for the construction of the models described in Secs. IV B and IV C. In the following, we label the model densities according to their corresponding target densities in Table I as $q(\phi, \varphi)$, $q(\phi)$, $q(\varphi|\phi)$, and $q(\phi|\varphi)$. In each sampling approach, using model distributions that better approximate the associated target will generally result in higher acceptance rates with potentially lower autocorrelations.

A. Optimizing flow-based models

We first detail a procedure to optimize the model density $q(\phi, \varphi)$ to directly approximate $p(\phi, \varphi)$. The KL divergence between these distributions is defined as

$$\begin{aligned} D_{\text{KL}}(q(\phi, \varphi) \| p(\phi, \varphi)) \\ &= \mathbb{E}_{\phi, \varphi \sim q} [\log(q(\phi, \varphi)/p(\phi, \varphi))], \\ &= \mathbb{E}_{\phi, \varphi \sim q} [\log q(\phi, \varphi) + S_B(\phi) + S_{PF}(\phi, \varphi) + \log Z]. \end{aligned} \quad (21)$$

It is minimized if and only if the target and model density match, for which $D_{\text{KL}} = 0$. In practice, a loss function based on this divergence is computed stochastically as

$$L = \frac{1}{N} \sum_{k=1}^N \log q(\phi_k, \varphi_k) + S_B(\phi_k) + S_{PF}(\phi_k, \varphi_k), \quad (22)$$

in terms of a minibatch of N samples $\{(\phi_k, \varphi_k)\}_{k=1}^N$ drawn from $q(\phi, \varphi)$. The unknown normalizing constant $\log Z$ has been removed in the definition of Eq. (22), since it is just an overall constant shift and does not affect the relevant structure of the loss function.

If the model probability density $q(\phi, \varphi)$ can be directly computed, we can evaluate the gradient of Eq. (22) with respect to the model parameters defining this probability density. Gradient-based optimization methods can then be applied to minimize L . This training procedure is immediately applicable to the models required for the joint sampling approaches derived in Secs. III C and III D. In the former, the distribution $q(\phi, \varphi)$ is defined by $q(\phi)q(\varphi|\phi)$, and this pair of model distributions is simultaneously optimized by minimizing the loss function in Eq. (22). In the latter, a model for $q(\phi, \varphi)$ is directly constructed and optimized.

The remaining distributions required in Secs. III A and III B, namely $q(\phi)$ and $q(\phi|\varphi)$, do not naturally define a joint model probability density. To optimize these distributions using the loss function above, we extend the model architectures by pairing $q(\phi|\varphi)$ with a sampler $q(\varphi)$ and

²We use the generic notation z and x here to stand for potentially high-dimensional variables acted on by the flow. In our applications these may include boson and/or pseudofermion field variables. Flows are described in field-theoretic notation wherever we work with flows particular to lattice field theory sampling.

FIG. 2. Schematic overview of a coupling-layer-based flow model architecture. Masks m_i and their complements $\bar{m}_i = 1 - m_i$ split the degrees of freedom into a subset to be updated and a subset that is frozen and used as input to the context function (red box), which provide the parameters for the invertible transformation f_i applied within each coupling layer g_i . Here, neural networks (NN) are used to implement the context functions. The generic variables z , x_i , and x may generally be high-dimensional field configurations in a lattice field theory setting. Particular implementations of translationally equivariant convolutional networks are defined in Secs. IV B 1 and IV B 2. Constructions of coupling layers with affine transformations based on these convolutional networks are defined in Sec. IV B 3.

pairing $q(\phi)$ with a sampler $q(\phi|\phi)$. The resulting joint models can be optimized as above, and the auxiliary components can be discarded after training. In this work, these auxiliary models are constructed as follows.

We first consider extending the ϕ -conditional model $q(\phi|\phi)$ to a joint model, which requires a marginal distribution $q(\phi)$. None of the sampling approaches presented in Sec. III directly require this marginal distribution; however, as we discuss further in Sec. IV C, we choose to model $q(\phi|\phi)$ by a restricted form of a joint sampler which simultaneously models a marginal distribution $q(\phi)$. In this extended model, both $q(\phi|\phi)$ and $q(\phi)$ are described by parameters that are optimized.

A ϕ -marginal model $q(\phi)$ can be extended to a joint sampler by pairing it with a conditional distribution $q(\phi|\phi)$. In principle such an auxiliary model could be constructed solely for the purposes of training. However, in this case we are free to instead use the exact conditional distribution $p(\phi|\phi)$, which can be exactly and efficiently sampled. The result is a joint distribution defined by first sampling ϕ from the ϕ -marginal model and then sampling the ϕ conditional using Eq. (16), resulting in the joint density $q(\phi)p(\phi|\phi)$. Evaluating the joint KL divergence between this model distribution and the target joint distribution $p(\phi, \phi)$ requires the evaluation of the normalized density $p(\phi|\phi)$, which unfortunately includes a normalizing factor of $\det \mathcal{M}(\phi)$. However, we only require an unbiased stochastic estimator of the gradients of Eq. (22) for optimization. Appendix C details a particular stochastic estimator for these gradients which can be used to avoid the costly determinant evaluation; this estimator is used for all ϕ -marginal models trained in this work.

B. Building blocks

Flow-based models are generally constructed by composing several simple, invertible transformation layers, each described by a number of free parameters. This composition produces an expressive overall transformation

that is nevertheless invertible and has a tractable Jacobian determinant. Coupling layers [49] are one common choice of simple transformation in which the degrees of freedom of each sample are divided into two subsets and one subset is updated conditioned on the other, “frozen” subset, as shown in Fig. 2. A “masking pattern” describes the division into subsets. Transformations of the updated subset are parametrized by “context functions” accepting the frozen subset as input, which are typically implemented using neural networks. For example, a simple coupling layer for a real scalar field $\phi(x) \in \mathbb{R}$ could be constructed based on a checkerboard division into even/odd sites, where the field at even sites is (invertibly) transformed by an element-wise rescaling operation plus an additional offset. The scaling factors and offsets are given by the output of an arbitrary context function, which may be parametrized by a neural network acting on the odd sites. The transformation is applied alternatingly between even and odd sites; see Ref. [86] for a concrete implementation of such coupling layers. Symmetries may be incorporated in such models using appropriate choices of masking patterns, context functions, and transformations. Other choices of layers are also possible (see Sec. IV B 5 below) and are similarly encoded using generic neural networks.

The target densities defined in Table I are all invariant under translations with appropriate boundary conditions, as discussed in Sec. II C. Previous works have shown that exactly incorporating known symmetries into machine learning models can accelerate their training and improve their final quality [89–94]. In the context of normalizing flows, ensuring that the model density is invariant under a symmetry group is achieved by choosing an invariant prior distribution and building transformation layers that are equivariant under the symmetry. Below, we introduce several “building blocks” which are designed to handle these symmetries. These building blocks are used in the implementation of various layers and flow-based models constructed in this work.

1. Translation-equivariant convolutions via P fields and AP fields

The joint distribution $p(\phi, \varphi)$ given in Eq. (14) is invariant under simultaneous field translations given by

$$\phi(\vec{x}, t) \rightarrow \phi'(\vec{x}, t) = \phi(\vec{x} - \delta\vec{x}, t - \delta t) \quad (23)$$

and

$$\begin{aligned} \varphi(\vec{x}, t) &\rightarrow \varphi'(\vec{x}, t) \\ &= \begin{cases} \varphi(\vec{x} - \delta\vec{x}, t - \delta t) & (t - \delta t) \bmod 2L_t < L_t \\ -\varphi(\vec{x} - \delta\vec{x}, t - \delta t) & L_t \leq (t - \delta t) \bmod 2L_t \end{cases} \quad (24) \end{aligned}$$

for any translations $(\delta\vec{x}, \delta t)$ in the translational symmetry group of the discretized theory. In this work, we label fields transforming as Eq. (23) as P fields, and we label fields transforming as Eq. (24) as AP fields.³

In previous applications of flow models to sampling configurations in lattice field theory, translational symmetry has been implemented for bosonic fields by applying convolutional neural networks [95] with circular padding (periodic boundary conditions) to generate parameters for transformations implemented in each flow layer [51,52,59]. All input, intermediate, and output fields in these applications were P fields. As a building block for translation-equivariant coupling layers acting on both bosonic and pseudofermionic fields, we extend this approach to define translation-equivariant convolutions that act on a generic set of input P fields and AP fields, producing output fields with a desired set of transformation properties, i.e., a specification of whether each channel of the output should be a P field or AP field. To implement such convolutional neural networks, we exploit the fact that P fields and AP fields form an algebra under pointwise addition and multiplication and restrict the operations appropriately to satisfy the desired output transformation properties; see Appendix D 1 for explicit implementation details.

2. Translation-equivariant convolutions via group averages

As an alternative to defining equivariant convolutional neural networks, one can symmetrize a nonequivariant architecture by explicitly averaging over the whole symmetry group [89]. Convolutional layers with periodic padding in all dimensions are already equivariant under translations of P fields and under all spatial translations of AP fields, thus only the subgroup of temporal translations needs to be averaged over to ensure equivariance for AP fields. The result is a generic method to produce convolutional neural networks with prescribed P-field and AP-field transformation properties of each output

channel. Compared with standard convolutions or the restricted equivariant architecture given in Sec. IV B 1, this method requires a greater computational effort by a factor proportional to the temporal extent of the lattice, L_t . However, it allows us the use of unrestricted convolutional architectures, including arbitrary activation functions and learned biases; see Appendix D 2 for further details.

3. Affine coupling layers

Translation-equivariant networks constructed by either of the methods discussed in Secs. IV B 1 and IV B 2 can immediately be applied in the construction of translation-equivariant affine coupling layers suitable for transforming real-valued scalar fields. An affine coupling layer transforms a field x to $ax + b$ (multiplication and addition are applied pointwise), where a and b are fields produced by context functions acting on the frozen components of the field x . Coupling layers, context functions, and masking patterns are illustrated in Fig. 2; see also Ref. [86]. Using translation-equivariant convolutional neural networks to produce a and b , either a bosonic field or pseudofermionic field can be updated in a translation-equivariant manner as long as

- (i) The parameters a and b are both P fields if x is a bosonic field; or
- (ii) The parameter a is a P field and b is an AP field if x is a pseudofermionic field.

Such coupling layers can be composed to produce translation-equivariant flows.

4. Equivariant linear operators

The conditional distribution $p(\varphi|\phi)$ is exactly Gaussian, suggesting that it may be efficiently modeled by flows based on architectures other than coupling layers. For example, one may define a linear operator $\mathcal{W} = \mathcal{W}(\phi)$ to transform the pseudofermion fields. The model distribution $q(\varphi|\phi)$ may then be defined by computing $\varphi = \mathcal{W}\chi$, where χ is drawn from the Gaussian distribution $\frac{1}{Z_N} e^{-\chi^\dagger \chi}$, such that

$$\begin{aligned} q(\varphi|\phi) &= \frac{1}{Z_N} e^{-\varphi^\dagger (\mathcal{W}\mathcal{W}^\dagger)^{-1} \varphi} (\det \mathcal{W}\mathcal{W}^\dagger)^{-1}, \\ &= \frac{1}{Z_N} e^{-\chi^\dagger \chi} (\det \mathcal{W}\mathcal{W}^\dagger)^{-1}. \end{aligned} \quad (25)$$

To effectively use this flow model, $\det(\mathcal{W}\mathcal{W}^\dagger)$ must be tractable to compute. In the case of a degenerate pair of fermion flavors, the target distribution is defined by

$$p(\varphi|\phi) = \frac{1}{Z_N \det DD^\dagger} e^{-\varphi^\dagger [D(\phi)D^\dagger(\phi)]^{-1} \varphi}. \quad (26)$$

While it is clearly sufficient for $\mathcal{W}$ to approximate D in this case, it is in fact only necessary that $\mathcal{W}\mathcal{W}^\dagger$ approximates $DD^\dagger$, allowing some freedom in the learned matrix $\mathcal{W}$.

³The fields ϕ and φ in Eqs. (23) and (24), and P fields and AP fields in general, may have multiple components per site.

We build the operator $\mathcal{W}$ as a composition of simple linear operators $\mathcal{W} = \mathcal{W}_n \circ \dots \circ \mathcal{W}_1$, where each $\mathcal{W}_k$ has only local interactions along a fixed dimension, in a fixed direction (that is, with only positive or negative offsets, but not both), allowing the determinant of each matrix to be efficiently computed. We choose to parametrize the components of each operator $\mathcal{W}_k$ by two P fields, produced from learned translation-equivariant functions of ϕ . More specifically, we consider $2d$ types of operators, where each type is defined by a sign $s = \pm 1$ and a choice of one of the d lattice directions. For the two-dimensional application described below, there are thus four distinct operator types. The different types of operators are applied alternatingly in the composition, but the specific order can be chosen arbitrarily. The operator type with couplings in the spatial direction and sign s thus updates a field χ by

$$(\mathcal{W}\chi)_{ij} = a_{ij}\chi_{ij} + b_{ij}\chi_{i+s,j} \quad (27)$$

with periodic boundary conditions along the space dimension: $\chi_{L+1,j} = \chi_{1,j}$ and $\chi_{0,j} = \chi_{L,j}$. An operator with temporal couplings updates a field χ by

$$(\mathcal{W}\chi)_{ij} = a_{ij}\chi_{ij} + b_{ij}\chi_{i,j+s} \quad (28)$$

with antiperiodic boundary conditions along the time dimension: $\chi_{i,L+1} = -\chi_{i,1}$ and $\chi_{i,0} = -\chi_{i,L}$. This construction may be understood as a convolutional layer with appropriate boundary conditions and an additional constraint on the kernel to have nonzero entries only in the center and at one of the $2d$ adjacent sites.

With these definitions, each operator $\mathcal{W}_k$ is block diagonal (for a suitable choice of basis). Each block is of the form

$$\begin{bmatrix} a_1 & & & \pm b_1 \\ b_2 & a_2 & & 0 \\ & \dots & \dots & \\ 0 & & & b_L & a_L \end{bmatrix}, \quad (29)$$

where we have dropped a (spatial or time) index to simplify the notation. The determinant of each block is simply $\Pi_h a_h \pm \Pi_h b_h$, indicating that the Jacobian determinant associated with the full composition can be tractably computed.

5. Convex potential flows

Because of the nonlocal nature of the effective action, we consider an alternative flow architecture to produce a model distribution $q(\phi)$ approximating $p(\phi)$. Convex potential flows (CPFs) are normalizing flows defined via the gradients of a potential that is strongly convex and twice differentiable almost everywhere [96,97]. Strong convexity

of the potential on a convex support $\mathcal{X}$ guarantees the flow to be invertible on $\mathcal{X}$, and this family of normalizing flows can be shown to be a universal density approximator [97]. We can parametrize strongly convex functions by neural networks with mild constraints on their architecture and weights [98].

More specifically, given a convex potential function $u: \mathcal{X} \rightarrow \mathbb{R}$, we define the map

$$[f(z)]_i = \frac{\partial}{\partial z_i} u(z), \quad (30)$$

where the index i specifies how each degree of freedom of z is mapped. Starting from a base density $r(z)$, the resulting probability density produced by mapping through f follows from Eq. (20) as

$$q(x) = r(z) \det H_u(z)^{-1}, \quad (31)$$

where $x = f(z)$ and $H_u(z) = \frac{\partial^2}{\partial z_i \partial z_j} u(z)$ is the Hessian matrix of $u(z)$. Training by minimizing the KL divergence $D_{\text{KL}}(q||p)$ between the model q and a target density p only requires the gradients $\nabla_\theta \log \det H_u(x)$ with respect to the model's parameters θ . Since the Hessian is symmetric and positive definite for strongly convex potentials, we can directly employ a stochastic trace estimator [97,99],

$$\begin{aligned} \nabla \log \det H_u(x) &= \nabla \text{Tr} \log H_u(x), \\ &= \text{Tr}[H_u(x)^{-1} \nabla H_u(x)], \\ &= \mathbb{E}_{\chi \sim e^{-\chi^\dagger \chi}}[(H_u^{-1}(x)\chi)^\dagger \nabla H_u(x)\chi]. \end{aligned} \quad (32)$$

The sample mean over noise vectors χ can be used to estimate this quantity in practice, and the inverse Hessian applied in $H_u(x)^{-1}\chi$ can be efficiently computed by the application of the conjugate-gradient method. Note that this estimator only requires the computation of Hessian-vector products $H\chi$, which is particularly convenient when the Hessian is sparse.

CPFs can be straightforwardly applied to construct flows that sample bosonic fields ϕ . They can also be constrained to be translation equivariant by using appropriate convolutions. In contrast to coupling layers, the CPF potential is a scalar function based on global information, which may result in transformations of the field ϕ that can in general be quite nonlocal. Evaluating the model probability density for use in asymptotically exact Markov chains requires a precise approximation of the log-det Hessian to avoid systematic errors. An exact calculation of the determinant is feasible only for small lattice volumes. For larger field configurations, one could apply a more scalable estimator, such as the estimator based on Lanczos tridiagonalization and the quadrature method described in Ref. [100].

C. Flow models

We next define particular architectures for modeling each of the distributions required for the four sampling approaches introduced in Sec. III. While the space of possible architectures that may be defined from the building blocks of Sec. IV B is large and the present discussion is not exhaustive, the use of each sampling method and each building block is demonstrated at least once. The architectures for each approach detailed in this section are summarized in Fig. 3.

1. Modeling $p(\phi)$ for ϕ -marginal sampling

The ϕ -marginal sampler defined in Sec. III A requires a flow whose model distribution $q(\phi)$ approximates $p(\phi)$.

Such a flow only needs to manipulate P fields. We build this ϕ -marginal model using a composition of CPF layers, where the output of each layer is defined by computing the gradient of a potential $u_i(\cdot)$ (see Sec. IV B 5). These layers act on samples ζ drawn from some base distribution $r_p(\zeta)$. Figure 3(a) depicts this type of ϕ -marginal architecture defined by a composition of CPFs acting on ζ . Each u_i contributes a determinant factor $\det H_{u_i}^{-1}$ to $q(\phi)$ in terms of the Hessian $H_{u_i}^{ab} = \frac{\partial^2 u_i(z)}{\partial z_a \partial z_b}$, such that

$$q(\phi) = r_p(\zeta) \prod_i \det H_{u_i}^{-1}. \quad (33)$$

As discussed in Sec. IV A, this marginal model is extended to the joint density $q(\phi, \varphi) = q(\phi)p(\varphi|\phi)$ for training.

FIG. 3. Architectures for the flow-based models defined in Sec. IV C for each sampling approach. Note that each coupling layer g_k^p or g_k^{ap} employs masking of the updated field as shown in Fig. 2, such that the frozen components of the field are included as input to context functions. Superscripts on coupling layers indicate the translational equivariance structure of coupling layer inputs and outputs (either consistently transforming as P fields or AP fields). Many other choices of architectures are possible to model each distribution; the figure reflects the choices utilized in the numerical study undertaken in Sec. VA. (a) ϕ -Marginal architecture based on convex potential flows (Sec. IV C 1), (b) Fully joint architecture for $q(\phi, \varphi)$ based on coupling layers (Sec. IV C 4), (c) ϕ -Conditional model $q(\phi, \varphi)$ defined via a restricted joint architecture (Sec. IV C 2), and (d) Autoregressive model $q(\phi)q(\varphi|\phi)$ defined via coupling layers and linear flows (Sec. IV C 3).

The density cannot be computed efficiently due to the determinants involved in the definition of $q(\phi)$ as well as in the normalizing constant of $p(\phi|\phi)$, but the flow is nevertheless trainable using stochastic estimates of the gradients. For sampling, the joint density itself may also be estimated using stochastic approximations of the determinant factors.

The architecture of the convex potential network $u(\phi)$ is based on Ref. [97] and is modified appropriately to account for the periodic boundary conditions. It consists of K layers of convolutions of the form

$$\begin{aligned} h_1 &= L_1(\phi), \\ h_{k+1} &= L_k^+(\text{SoftPlus}(\text{ActNorm}(h_k))) + L_k(\phi), \\ u(\phi) &= w_1 \text{Sum}(h_K) + w_2 \frac{\|\phi\|^2}{2}, \end{aligned} \quad (34)$$

where L_j is a convolution layer with periodic boundary conditions and unconstrained weights; L_j^+ is a convolution layer with periodic boundary conditions and positive-only weights; $\text{SoftPlus}(x) = \log(1 + \exp(x))$; $\text{ActNorm}(x) = (x - \mu)/\sigma$ is a layer that normalizes its inputs using a learnable offset μ and scale σ , where μ and σ are initialized as the mean and standard deviation of the inputs of an initialization batch [101]; w_1, w_2 are learnable weights used to control closeness of the flow to the identity map at initialization. The use of periodic boundary conditions for L_j and L_j^+ and the final Sum operation ensures that $u(\phi)$ is invariant to translations.

2. Modeling $p(\phi|\phi)$ for Gibbs sampling

Section III B describes a Gibbs sampling scheme that utilizes the exact conditional $p(\phi|\phi)$ and a modeled conditional density $q(\phi|\phi)$. A ϕ -marginal model $q(\phi)$ is required to extend $q(\phi|\phi)$ to the joint distribution $q(\phi|\phi)q(\phi)$ for training. We achieve this simultaneous modeling of $q(\phi|\phi)$ and $q(\phi)$ by using a fully joint architecture with restricted information flow, as shown in Fig. 3(c). The model consists of a prior distribution over the base configurations ζ, χ denoted by $r_p(\zeta)$ and $r_{ap}(\chi)$, followed by the application of two types of affine coupling layers. First, the layers $g_k^p(\cdot; \chi_k)$ update the P-field configuration conditioned on the AP field, along with the frozen components of ζ_k , to produce $q(\phi|\phi)$ as

$$q(\phi|\phi) = r_p(\zeta) \prod_k \det J_{g_k^p}^{-1}, \quad (35)$$

where $J_{g_k^p}$ is the Jacobian for coupling g_k^p . Second, the couplings $g_k^{\text{ap}}(\cdot)$ transform the AP field χ conditioned solely on its frozen components to obtain $q(\phi)$:

$$q(\phi) = r_{ap}(\chi) \prod_k \det J_{g_k^{\text{ap}}}^{-1}. \quad (36)$$

To conditionally resample ϕ from $q(\phi|\phi)$ while leaving ϕ unchanged, the bosonic prior variable is resampled and the output of the flow is reevaluated while holding the pseudofermionic prior variable χ fixed. When ϕ is resampled from $p(\phi|\phi)$ in the alternate step of the Gibbs sampler, it is important to update the value of χ by passing ϕ through the inverse of the bottom branch of the flow depicted in the figure. This allows future resampling of ϕ as well as the calculation of the conditional probability density defined by the model.

3. Autoregressive modeling of $p(\phi, \phi) = p(\phi)p(\phi|\phi)$

Section III C defines an independence sampler based on an autoregressive joint model with the probability density given by $q(\phi, \phi) = q(\phi)q(\phi|\phi)$. We implement $q(\phi)$ using masked affine coupling layers, as was done in Refs. [51, 54, 55]. The parameters of the affine coupling layers are given by convolutional networks satisfying translational equivariance through standard periodic boundaries. We implement $q(\phi|\phi)$ using a deep linear flow consisting of learned linear operators $\mathcal{W}_k(\phi)$ as detailed in Sec. IV B 4. The parameters of these linear operators are all P fields obtained by similar periodic convolutional networks. The full joint model is given by the autoregressive combination of these two models, i.e., drawing ϕ from the affine model with distribution $q(\phi)$, then drawing ϕ from the conditional deep linear flow with distribution $q(\phi|\phi)$, as shown in Fig. 3(d). The marginal model is defined by sampling ζ from the prior distribution $r_p(\zeta)$, then applying the sequence of coupling layers $g_k^p(\cdot)$ such that the marginal model probability density $q(\phi)$ is given by

$$q(\phi) = r_p(\zeta) \prod_k \det J_{g_k^p}^{-1}. \quad (37)$$

The conditional linear flow is defined by sampling χ from the prior distribution $r_{ap}(\chi)$ and applying the linear operators $\mathcal{W}_k(\phi)$ to obtain the model density

$$q(\phi|\phi) = r_{ap}(\chi) \prod_k \frac{1}{\det \mathcal{W}_k \mathcal{W}_k^\dagger}. \quad (38)$$

We define $r_{ap}(\chi) = \frac{1}{Z_N} e^{-\chi^\dagger \chi}$ to match the choice for the linear operator flow in Eq. (25).

Note that the learned components in this approach may also be combined in novel ways. For example, it is possible to discard the conditional flow with distribution $q(\phi|\phi)$ after training and simply use $q(\phi)$ for ϕ -marginal sampling as described in Secs. III A and IV C 1. This may be advantageous in situations where gradients from an exactly sampleable distribution are not available and training must be fully variational. On the other hand, the conditional deep

linear flow may be used by itself as a determinant estimator for given configurations ϕ .

4. Fully joint modeling of $p(\phi, \varphi)$

Finally, we construct a model that simultaneously samples ϕ and φ in a fully joint approach which can be employed in the exact sampler defined in Sec. III D. The joint model implemented in this work is constructed from affine coupling layers that alternatingly transform the bosonic fields conditioned on the pseudofermionic fields, and vice versa, as shown in Fig. 3(b). The model is defined to sample ζ and χ from the prior distributions $r_p(\zeta)$, $r_{ap}(\chi)$ and subsequently apply alternating layers. Coupling layers $g_k^p(\cdot, \chi_k)$ transform the P-field base configuration ζ conditioned on its frozen components and the AP-field configuration χ_k , while couplings $g_k^{ap}(\cdot, \zeta_k)$ update the AP-field base configuration χ_k conditioned on its frozen components and ζ_k . This gives rise to the joint density

$$q(\phi, \varphi) = r_p(\zeta) r_{ap}(\chi) \prod_k \det J_{g_k^p}^{-1} \det J_{g_k^{ap}}^{-1}. \quad (39)$$

V. APPLICATION TO A SCALAR YUKAWA THEORY IN TWO DIMENSIONS

As a demonstration, we implement the sampling algorithms defined above for a two-dimensional model of a scalar field coupled to fermions via a Yukawa interaction. This model provides a testbed which features fermionic fields, but without the additional complications brought on by gauge symmetry. Apart from providing a suitable test case for the development of flows capable of modeling fermions, studying Yukawa interactions is also interesting in its own right, e.g., for Higgs physics [102] or the quark-meson model [103].

A. Yukawa theory on the lattice

We consider a real, scalar field ϕ coupled to one mass-degenerate pair of Kogut-Susskind staggered fermions [78]. We emphasize that our method can in principle be straightforwardly applied to other discretization schemes, such as Wilson fermions, without additional conceptual difficulties.

The action for this theory is defined as $S(\psi, \bar{\psi}, \phi) = S_B(\phi) + S_F(\psi, \bar{\psi}, \phi)$ [see Eq. (1)]. For the scalar field we choose the usual discretized Klein-Gordon action with quartic self-interaction, defined by

$$S_B(\phi) = \sum_{x \in \Lambda} \left[-2 \sum_{\mu=1}^d \phi(x) \phi(x + \hat{\mu}) + (m^2 + 2d) \phi(x)^2 + \lambda \phi(x)^4 \right], \quad (40)$$

where Λ denotes the set of lattice sites, m the bare scalar mass parameter, λ the coupling, and d the dimension. The fermionic action S_F is defined by the bilinear form in Eq. (2) with $N_f = 2$ and both fermion flavors are defined by the discretized Dirac operator

$$D_{xy} = \sum_{\mu=1}^d \eta_{\mu}(x) \frac{\delta(x - y + \hat{\mu}) - \delta(x - y - \hat{\mu})}{2} + \delta(x - y)(m_f + g\phi(x)), \quad (41)$$

where m_f is the bare mass of the fermion and g the Yukawa coupling. The staggered factor η_{μ} is obtained from the Dirac γ matrices after the staggered transformation and is defined as

$$\eta_1(x) = 1 \quad \text{and} \quad \eta_l(x) = (-1)^{x_1} \dots (-1)^{x_{l-1}}. \quad (42)$$

The Kronecker δ is defined to have antiperiodic boundary conditions in the time direction (conventionally taken to be $\mu = d$) and periodic boundary conditions in the spatial directions, i.e.,

$$\delta(x) = \prod_{\mu=1}^d \delta_{\mu}(x_{\mu}) \quad (43)$$

where

$$\delta_{\mu \neq d}(x_{\mu}) = \begin{cases} 1 & \text{if } x_{\mu} = 0, \pm L \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and } \delta_d(x_d) = \begin{cases} 1 & \text{if } x_d = 0 \\ -1 & \text{if } x_d = \pm L_t \\ 0 & \text{otherwise} \end{cases}. \quad (44)$$

We employ an even-odd preconditioning scheme for the Dirac operator for all models except for the autoregressive model using linear operators. In contrast to the default lexicographic ordering, sorting lattice sites into even and odd allows to bring the matrix into a form that is amenable to an explicit block factorization of the determinant, which leads to improvements in the conditioning and solver performance. This reduces the variance and cost of computing the pseudofermion action required for optimizing models and sampling. Most previous work on improved orderings has focused on techniques for Wilson fermions in the context of gauge theory [71,73,104], but the same insights can be applied to the staggered fermion formulation used in this work, as detailed in Appendix B.

All results reported in this work are computed on a 16×16 lattice geometry using the two choices of action parameters in the symmetric phase given in Table II. For this theory and lattice discretization, there is no additive renormalization to the bare fermion mass m_f . Accordingly,

we directly probe the case of vanishing mass by setting $m_f = 0$. The first set of parameters, for which the Yukawa coupling is chosen to be $g = 0.1$, already provides a realistic test scenario in the sense that the average ratio of fermionic to scalar force magnitudes is around 3%, which is similar to the ratio of fermionic to gauge forces reported in the literature for some lattice QCD computations; see e.g., Refs. [105,106]. The second choice with $g = 0.3$ features a much larger force ratio amounting to about 39%, which thus provides a testbed for theories with more prominent fermionic effects. For simplicity, we will refer to these two parameter choices by the associated value of the Yukawa coupling g .

To evaluate the performance of the sampling approaches presented here, we consider the following quantities and observables. The magnetization of a scalar field configuration is defined as

$$M = \frac{1}{V} \sum_{x \in \Lambda} \phi(x). \quad (45)$$

We measure the average absolute value $\langle |M| \rangle$, which provides a nonzero order parameter that is large in the broken symmetry phase and exponentially suppressed in the symmetric phase. The connected two-point correlation function of ϕ is defined as

$$C_\phi(x, y) = \langle \phi(x)\phi(y) \rangle - \langle \phi(x) \rangle \langle \phi(y) \rangle, \quad (46)$$

where we fix $\langle \phi(x) \rangle = \langle M \rangle = 0$ analytically. We evaluate the source-averaged correlator projected to zero momentum defined by

$$C_\phi(t) = \frac{1}{V} \sum_x \sum_{\vec{y}} C(x, x + (\vec{y}, t)). \quad (47)$$

Fermionic observables can be computed from the matrix elements of the inverse Dirac operator. The chiral condensate of the fermion field is defined by

$$\langle \bar{\psi}\psi \rangle = \left\langle \frac{1}{V} \text{Tr} D^{-1} \right\rangle, \quad (48)$$

and we measure $\langle |\bar{\psi}\psi| \rangle$ as for the magnetization. Using the off-diagonal matrix elements, we also evaluate the average fermionic two-point correlator in the time direction,

$$C_\psi(t) = \langle \psi(y)\bar{\psi}(0) \rangle = \langle D_{y,0}^{-1} \rangle, \quad (49)$$

where $y = (\vec{0}, t)$ with t odd. The particular choices of offsets y select staggered spinor components at the sink $\psi(y)$ that result in a nonzero average correlation function originating from the source $\bar{\psi}(0)$.

TABLE II. The two parameter choices for our reported numerical studies and the associated average absolute magnetization and chiral condensate computed with HMC. All uncertainties reported in this work are obtained using data blocking to account for autocorrelations and applying the statistical jackknife method. Force ratios are determined by dividing the average L^2 norms of fermionic and bosonic force vectors.

V	m^2	λ	g	m_f	$\langle M \rangle$	$\langle \bar{\psi}\psi \rangle$	Force ratio
16^2	-4.00	6.0	0.1	0	0.0733(1)	0.0159(1)	3%
16^2	-1.55	2.4	0.3	0	0.0791(1)	0.0490(1)	39%

B. Model architectures

For each of the four sampling approaches outlined in Sec. III and corresponding model architectures detailed in Sec. IV C, we create specific models for both choices of target action parameters given in Table II:

- (i) For the sampling scheme described in Sec. III A, we construct a CPF model defining a ϕ -marginal distribution $q(\phi)$ approximating the corresponding $p(\phi)$. The model architecture and training follow the generic procedure outlined in Sec. IV C 1.
- (ii) To build a conditional model $q(\phi|\varphi)$ for the Gibbs sampler described in Sec. III B, we implement a restricted affine coupling layer flow as described in Sec. IV C 2. To achieve translational equivariance, we employ the method of group averages described in Sec. IV B 2 for each convolutional network.
- (iii) To produce an autoregressive joint model density $q(\phi, \varphi) = q(\phi)q(\varphi|\phi)$ for the sampling scheme described in Sec. III C, we use a model consisting of affine coupling layers followed by learned equivariant linear transformations as described in Sec. IV C 3.
- (iv) For the fully joint sampling scheme described in Sec. III D, we implement a model with unrestricted affine coupling layers acting on both ϕ and φ , using translation-equivariant convolutions as described in Sec. IV B 1. This results in a fully joint model distribution $q(\phi, \varphi)$ as detailed in Sec. IV C 4.

The models are optimized for each approach based on the joint KL divergence discussed in Sec. IV A. Prior distributions for the initial P-field ζ and AP-field χ , where they are used according to Fig. 3, are Gaussians of the form

$$r_p(\zeta) = \frac{1}{Z_\zeta} e^{-\zeta^\dagger \zeta / (\sigma_\zeta)^2},$$

$$\text{and } r_{\text{ap}}(\chi) = \frac{1}{Z_\chi} e^{-\chi^\dagger \chi / (\sigma_\chi)^2}, \quad (50)$$

with specific values of σ_ζ and σ_χ for each model chosen to enhance the training stability.

Details of the model hyperparameters and training procedure for each of the models can be found in

Appendix A. An exhaustive search over the available parameter space is beyond the scope of the present work, and as such it can be expected that tuning the various model hyperparameters may further improve the reported performance metrics.

C. Discussion and comparison of sampling schemes

After optimization, we use each of the models to construct asymptotically exact samplers for their respective target distributions according to the four schemes given in Sec. III. For each case, we produce 100 distinct Markov chains consisting of 10k steps each, of which the first 1k steps are discarded for thermalization. These Markov chains are used for observable measurements and to investigate and compare metrics of the efficiency of sampling via each of these methods.

First, we confirm that each of the observables described above are measured to be consistent across sampling schemes and with HMC baseline results. Calculations of $\langle |M| \rangle$ and $\langle |\bar{\psi}\psi| \rangle$ using each of the generated ensembles are detailed in Table III and are all consistent with the results obtained through HMC. The scalar and fermionic two-point correlators produced by the four exact Monte Carlo sampling schemes models are also consistent with the HMC baseline, as shown in Figs. 4 and 5.

For the various sampling approaches, Table III compares the autocorrelations of the magnetization and chiral condensate, as well as the Markov chain acceptance rates. Given a chain of N measurements for some real-valued observable X , its autocorrelation function is defined as

$$\Gamma_X(\tau) = \frac{1}{N-\tau} \sum_{i=1}^{N-\tau} X_i X_{i+\tau} - \langle X \rangle^2, \quad (51)$$

where τ denotes the number of Markov chain steps separating the pair of measurements considered. The integrated autocorrelation time for observable X is given as

$$\tau_X^{\text{int}} = \frac{1}{2} + \lim_{\tau^{\text{max}} \rightarrow \infty} \sum_{\tau=1}^{\tau^{\text{max}}} \frac{\Gamma_X(\tau)}{\Gamma_X(0)}. \quad (52)$$

The sum can be truncated at a sufficiently large τ^{max} due to the exponential suppression of $\Gamma_X(\tau)$; $1 \ll \tau^{\text{max}} \ll N$ should be satisfied to ensure that the values of $\Gamma_X(\tau)$ are reliable. For the τ^{int} values reported in this work, we use the Madras-Sokal windowing procedure [108] to choose a suitable τ^{max} by identifying the earliest point where $c\tau^{\text{int}} \leq \tau^{\text{max}}$, with $c = 10$. The integrated autocorrelation times τ_M^{int} and $\tau_{\bar{\psi}\psi}^{\text{int}}$ are given in Table III together with the mean acceptance rates for all four sampling procedures.

To understand the relative performance of the four sampling approaches detailed in Table III, we note that the dependence of the acceptance rate and autocorrelations on model quality is quite distinct in several of these approaches. For one, the ϕ -marginal sampler involves an exact determinant measurement in the sampling step used for the numerical study above, which is not expected to scale efficiently. If replaced with the pseudomarginal estimator discussed in Sec. III A, the variance of the noisy estimates of each determinant would degrade the statistical performance achieved by even an optimally trained model and improved estimators, in particular when encountering large condition numbers. This is an obstacle to working with light fermion masses or field configuration geometries with many lattice sites (e.g., near the thermodynamic limit) independent of the challenge of training accurate model approximations to the target distribution. Thus, the relatively higher acceptance rates and lower autocorrelation times achieved by the ϕ -marginal sampler constructed from CPFs in this study must be contrasted against the potentially difficult scaling challenges or requirements for more precise stochastic estimators in this approach. By comparison, there is no nontrivial upper bound on the acceptance rate of the other three sampling approaches, and they will

TABLE III. Sampling performance metrics and observables for all approaches, computed from 100 Markov chains with 10k proposals each, where the first 1k are discarded for thermalization. For each model, the first row shows results obtained for $g = 0.1$ and the second row for $g = 0.3$, respectively. For comparison, the values obtained with HMC listed in Table II are consistent with the measurements from our models. Autocorrelation times τ^{int} are computed for each of the 100 chains and then averaged, and errors are obtained with statistical jackknife. The results are discussed in more detail in Sec. V C. All models except the autoregressive make use of even-odd preconditioning of the action.

MCMC approach	Modeled targets	Flow model	Acceptance rate	$\langle M \rangle$	$\langle \bar{\psi}\psi \rangle$	τ_M^{int}	$\tau_{\bar{\psi}\psi}^{\text{int}}$
ϕ -marginal (III A)	$p(\phi)$	IV C 1	92%	0.0734(1)	0.0159(1)	0.72(1)	0.71(1)
			92%	0.0792(1)	0.0491(1)	0.67(1)	0.67(1)
Gibbs (III B)	$p(\phi \varphi)$	IV C 2	60%	0.0735(1)	0.0160(1)	2.02(4)	2.02(3)
			44%	0.0792(1)	0.0490(1)	2.74(4)	2.73(4)
Autoregressive (III C)	$p(\phi), p(\phi \phi)$	IV C 3	53%	0.0731(1)	0.0159(1)	2.16(3)	2.16(3)
			43%	0.0790(1)	0.0489(1)	3.62(7)	3.60(7)
Fully joint (III D)	$p(\phi, \varphi)$	IV C 4	37%	0.0733(1)	0.0159(1)	4.98(11)	4.98(11)
			31%	0.0791(1)	0.0490(1)	8.73(30)	8.67(30)

FIG. 4. Connected two-point correlation functions of the scalar field (projected to zero spatial momentum in each time slice t) for each model and choice of action parameters, computed from 100 Markov chains with 9k configurations each. Error estimates are obtained using data blocking with a bin size of 100 and applying statistical jackknife. Left: $g = 0.1$, right: $g = 0.3$. Bottom panels show the ratio of each data point to the HMC baseline, where the shaded regions correspond to the 1σ and 2σ uncertainty bands of the HMC results. For both the scalar correlators here and the fermionic ones in Fig. 5, Hotelling's t-squared statistic [107] comparing each flow model result to the HMC baseline finds results to be consistent with correlated statistical fluctuations.

FIG. 5. Average fermionic two-point correlation in the time direction for each model and choice of action parameters using the same configurations and data blocking as for Fig. 4. Left: $g = 0.1$, right: $g = 0.3$. The choice of odd t selects staggered spinor components at the sinks that give a nonzero average correlation with the source at $x = 0$. Shaded regions in the bottom panel again depict the 1σ and 2σ uncertainties of the HMC baseline results.

achieve an acceptance rate of 100% when perfect model distributions are constructed.

Among these three approaches, the Gibbs sampler must also be further contrasted against the autoregressive and fully joint samplers. In particular, the remaining conditional structure of updates to ϕ and φ in the Gibbs sampler results in autocorrelations even if the acceptance rate is 100%. The magnitude of these residual autocorrelations may be small, but nevertheless puts a bound on the performance that is theoretically achievable by a Gibbs sampler, even in the asymptotic limit of perfect models of the involved distributions. Thus only joint models (either autoregressive or fully joint) can completely eliminate autocorrelations in the ideal limit of perfect models. In practice, however, the distinctions between joint models and Gibbs sampling may be minor. For example, the results presented in Table III demonstrate that at the similar acceptance rates of roughly 40%–50% for the Gibbs and autoregressive samplers, the integrated autocorrelation times for the magnetization and condensate are similar, despite the additional autocorrelations introduced by the particular conditional structure of the Gibbs sampling scheme. The fully joint sampler shows a lower acceptance rate and greater autocorrelations, indicating that the differences are largely based on the model approximation qualities.

The particular flow-based models implemented to approximate the various distributions used for the four sampling approaches also have distinct scaling prospects. It has been found in previous work [55] that flows based on coupling layers using convolutional networks may be easily transferred between different lattice volumes and thereby trained efficiently. This generalizability applies to the affine coupling layer implementations used for the Gibbs, autoregressive, and fully joint samplers described in this work. The CPF implementation for ϕ -marginal sampling is also based on convolutional networks for the construction of the convex potentials, thus enabling efficient measurements of these potentials at all lattice volumes. However, in this case computing the Jacobian of the transformation to calculate $q(\phi)$ is potentially expensive, because it requires the evaluation of the Hessian of each u_i . Stochastic estimation of these Hessian factors may introduce additional noise in exact sampling schemes based on these particular flow architectures.

These results numerically demonstrate the effectiveness of our proposed flow models and sampling schemes. The observed performance differences cannot immediately be attributed to inherent advantages of the chosen building blocks, but may also depend strongly on the model implementation details and theory-specific characteristics. The situation may also be quite different for larger volumes and dimensions as well as other types of fields and interactions, and disentangling the effects of implementation details from asymptotic scaling properties will be the subject of future research. Furthermore, there is a large

space of possible combinations of the building blocks introduced here that could be explored in future work to determine models that may have more efficient training, sampling, and scaling prospects. While an exhaustive search over this space is beyond the scope of this exploratory work, the present results serve as a guide for the design of custom flows for lattice simulations with dynamical fermions in other applications.

VI. APPLICABILITY TO UPDATE-BASED APPROACHES

While the sampling schemes presented in this paper are based on the independent proposal of new field configurations (except for the Gibbs sampler; see Sec. V for further discussion), the flow-based models defined here may also be used in methods that instead propose configuration updates, rather than independent samples. A simple way to achieve this with our architectures would be to formulate stochastic processes in the flow prior that guarantee asymptotic exactness under the target distribution, such as partial heatbath resampling, HMC, or Langevin-type algorithms, rather than independently drawing a completely new prior sample in every update step; such partial updates have previously been studied in the context of other generative models [96,109,110] as well as trivializing map approaches [111,112]. In contrast to these update-based methods, direct sampling approaches have the advantage that autocorrelations in the flow-based Markov chain are in principle eliminated for an ideal model. Imperfect models, however, can still result in residual correlations caused by rejections in the Metropolis step. Whether these residual correlations from an imperfect model can outweigh the autocorrelations in corresponding update-based methods is an open question, the answer to which may also depend strongly on the model details and the specific problem under consideration.

Apart from devising modified sampling schemes for the types of flows presented in this work, one may also consider defining flows that directly transform configurations in order to produce proposals for Markov chain updates. Related work on learning improved HMC-like updates includes A-NICE-MC [57], its recent application to the lattice simulation of scalar ϕ^4 theory [59], L2HMC [58], and DLHMC [60], which was demonstrated to successfully mitigate topological freezing in the context of $U(1)$ lattice gauge theory in two dimensions. These approaches require the implementation of flows suitable for transforming the primary fields and conjugate momenta conditioned on each other. The flows over pseudofermion variables developed in this work can therefore be used to extend such methods to the setting of lattice field theory involving dynamical fermion fields. These insights may also inform the design of novel building blocks for the self-learning Monte Carlo method mentioned in the introduction, which

was recently applied to non-Abelian gauge theory with dynamical fermions [47].

VII. SUMMARY AND OUTLOOK

In this paper, we introduce four approaches to applying flow-based sampling to fermionic lattice field theories based on different decompositions of the joint action over bosonic and pseudofermionic fields. All approaches satisfy detailed balance and thus provide asymptotically exact sampling schemes. We further introduce several techniques to model the distributions required in these sampling approaches, including the construction of flow-based models satisfying the more complex translational symmetry group arising from the pseudofermion boundary conditions (which must be antiperiodic in time). All four sampling methods are demonstrated to successfully produce asymptotically exact samplers in a proof-of-principle application to a two-dimensional Yukawa theory.

The flow-based model architectures presented here represent a selection from a large class of possible ways to model the distributions required for the four sampling methods detailed in this work. The observed relative performance of the methods in the application to the Yukawa theory provides a starting point for understanding the distinctions between these sampling approaches and architectures, but should not be considered a definitive indicator of the performance of these approaches or architectures in the context of other theories or at larger scales. Applying these methods to QCD will require future work to understand the scaling of the method with lattice volume, spacetime dimension, and with the involvement of gauge fields.

Importantly, investigating the continuum limit of flow-based samplers is relevant to determine their potential to mitigate critical slowing down. This question arises with or without fermions, and empirical studies are required to understand the scaling of these methods both for fermionic and purely bosonic theories. Nonetheless, the “building blocks” of flows suitable for fields including pseudofermions, and the sampling strategies outlined in this work, provide the basis for developing efficient flow-based samplers for fermionic theories.

Continued work into improved stochastic approximations of determinants [83,113–118] complements the flow-based sampling approach presented here. Such developments may be combined with our proposed flow-based sampling framework. For example, most flow-based samplers constructed in this work were designed to target the action after the application of even-odd preconditioning, a standard improvement technique for the fermionic determinant. Other preconditioners for the action may also be applied to increase the performance of flow-based samplers for these theories. Furthermore, improvements to unbiased stochastic estimates of fermionic determinants may increase the efficiency of the ϕ -marginal sampler at scale.

In summary, this work sets the stage for flow-based sampling of lattice field theories with fermions and paves the way towards an application of generative neural samplers to lattice QCD with dynamical quarks and similar problems in condensed matter theory. The next steps in this endeavor are the transfer of insights gained from the Yukawa model studied in this work to interactions between fermions and gauge fields, and the study of the scalability of the method. If they can be achieved, high quality flow-based models at the scale of state-of-the-art calculations may be able to circumvent the limitations of traditional sampling algorithms, thereby expanding the frontiers of lattice QCD and other lattice field theories.

ACKNOWLEDGMENTS

We thank F. Attanasio, M. Bauer, S. Blücher, C. Corvalán, W. Detmold, J. M. Pawłowski, I.-O. Stamatescu, N. Strodthof, and F. P. G. Ziegler for helpful discussions. G. K., D. B., D. C. H., and P. E. S. are supported in part by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Grant Contract No. DE-SC0011090. P. E. S. is additionally supported by the National Science Foundation under EAGER Grant No. 2035015, by the U.S. DOE Early Career Award No. DE-SC0021006, by a NEC research award, and by the Carl G and Shirley Sontheimer Research Fund. K. C. is supported by the National Science Foundation under the Grants ACI-1450310, OAC1836650, and OAC-1841471 and by the Moore-Sloan data science environment at NYU. M. S. A. thanks the Flatiron Institute and is supported by the Carl Feinberg Fellowship in Theoretical Physics and the James Arthur Fellowship. This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2181/1–390900948 (the Heidelberg STRUCTURES Excellence Cluster), the Collaborative Research Centre SFB 1225 (ISOQUANT), and the U.S. National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, [119]). This work is associated with an Argonne Leadership Computing Facility (ALCF) Aurora Early Science Program project, and used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DEAC02-06CH11357. The authors acknowledge the Massachusetts Institute of Technology (MIT) SuperCloud and Lincoln Laboratory Supercomputing Center [120] for providing HPC resources that have contributed to the research results reported within this paper.

APPENDIX A: MODEL AND TRAINING DETAILS

This section lists all necessary details to reproduce the flow architectures discussed in Sec. V B for the application to the two-dimensional Yukawa theory studied in this work.

TABLE IV. Model and training hyperparameters for all architectures discussed in Sec. V B. Further details and references are provided in Appendix A, in particular the specifics of the linear operators for the joint autoregressive model that are not listed here.

<i>Model parameters</i>	ϕ marginal	Gibbs	Autoregressive	Fully joint
Flow layers	3	16	12	12
Convolutions per layer	4	3	10	6
Number conv channels	16	32	64	64
σ_ζ	0.34	0.1	0.34	0.34
σ_χ	...	0.1	0.15	0.15
Kernel size	3	3	3	3
Activations (inner / final)	SoftPlus/...	LeakyReLU/Tanh	LeakyReLU/...	LeakyReLU/Tanh
<i>Training parameters</i>				
Gradient steps	500k	30k	200k	50k
Batch size	3072	2000	3072	3072
Learning rate schedule	10^{-3} , 10^{-4} after 80k	10^{-3} , 10^{-5} after 20k	10^{-4} , 2×10^{-5} after 60k, 10^{-5} after 120k	3×10^{-4} , 6×10^{-5} after 30k
Gradient clipping (value / norm)	10/32	.../...	10/32	10/1000

In particular, all relevant model and training hyperparameters are given in Table IV. Additional peculiarities of the linear operator and fully joint model implementations not listed in the table are discussed below. Furthermore, we provide references for the elementary machine learning components used in the implementation and optimization.

All models are trained using the well-known Adam optimizer [121] with default settings. In some cases, clipping of the gradient value and norm was employed to stabilize training [122]. The deep neural networks providing the context functions and convex potentials in the flow architectures are implemented exclusively in terms of convolutional networks with several hidden layers and channels [95]. The nonlinear activation functions employed in these networks are the rectified linear unit (ReLU) [123], in particular the LeakyReLU variant [124], as well as the SoftPlus function in the case of the CPF layers, as detailed in Sec. IV C 1. In some cases, an additional Tanh activation is applied to the output of each network, as specified in Table IV.

For all experiments in this work using the CPF architecture, we set $w_1 = 5 \times 10^{-3}$ and $w_2 = 1$ at initialization [cf. Equation (34)] and all convolutional layers use a stride of 1.

As detailed in Sec. V A, even-odd preconditioning of the Dirac matrix is not applied for the autoregressive architecture with linear operators, as we observe that the model gives a better approximation to the nonpreconditioned action with standard lexicographic ordering. The space of possible model adjustments to make the even-odd decomposition compatible with linear operators is large, and modifications could be explored to improve the results. The conditional density $q(\phi|\phi)$ is implemented using the composition of 128 equivariant linear operators $\{\mathcal{W}_k\}_{k=1}^{128}$. The 128 linear operators are jointly defined by the stacking

of a single squeezing layer breaking invariance under odd translations as explained in Appendix D 1, followed by a convolutional network with periodic boundary conditions. This network features 10 hidden layers with 64 channels each and uses intermediate LeakyReLU activations. In total, the network has 256 output channels, with each pair of output channels providing the values of a and b in the definition of one of the 128 linear operators. The a output is additionally transformed using a normalized SoftPlus function. We also find it useful to add an L^2 regularization loss for both outputs, with a weight of 10^{-5} .

For the fully joint model, active components of the scalar field are transformed using its frozen components as well as the pseudofermion field. The updated scalar field together with the frozen pseudofermion components are then used to update the active pseudofermion sites.

APPENDIX B: EVEN-ODD PRECONDITIONING

Ordering lattice sites into even and odd subsets allows writing the Dirac matrix D as a 2×2 block matrix of the form

$$D = \begin{pmatrix} m_f + g\phi_o & D_{oe} \\ D_{eo} & m_f + g\phi_e \end{pmatrix} \equiv \begin{pmatrix} \mathcal{A} & \mathcal{B} \\ \mathcal{C} & \mathcal{D} \end{pmatrix}, \quad (\text{B1})$$

where we denote the blocks as $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$, $\mathcal{D}$ for simplicity. The constant blocks $\mathcal{B} = D_{oe}$ and $\mathcal{C} = D_{eo}$ couple odd to even sites and vice versa, and ϕ_o and ϕ_e indicate the components of ϕ , respectively, associated with odd and even sites of the lattice. This form allows a more efficient stochastic approximation of the determinant by decomposing it into the determinant of either diagonal block and the associated Schur complement as

$$\begin{aligned}\det D &= \det(\mathcal{A}) \det(\mathcal{D} - \mathcal{C}\mathcal{A}^{-1}\mathcal{B}), \\ &= \det(\mathcal{A}\mathcal{C}^{-1}\mathcal{D} - \mathcal{B}) \det(\mathcal{C}),\end{aligned}\quad (\text{B2})$$

or, equivalently,

$$\begin{aligned}\det D &= \det(\mathcal{D}) \det(\mathcal{A} - \mathcal{B}\mathcal{D}^{-1}\mathcal{C}), \\ &= \det(\mathcal{D}\mathcal{B}^{-1}\mathcal{A} - \mathcal{C}) \det(\mathcal{B}).\end{aligned}\quad (\text{B3})$$

Rewriting from the first to the second form in Eqs. (B2) and (B3) ensures that the resulting expression does not involve terms that mix $\mathcal{A}$ and $\mathcal{D}$ with their respective inverses. This may lead to numerical instabilities if $m_f = 0$, which can result in ill-conditioned $\mathcal{A}$ and $\mathcal{D}$. Since $\mathcal{B}$ and $\mathcal{C}$ are constant, the terms $\det(\mathcal{B})$, $\det(\mathcal{C})$ drop out of the path integral and thus do not affect Metropolis-Hastings acceptance rates or gradients for optimization of flow-based models. Hence, they can be ignored for the purpose of training and sampling flow models.⁴

The reduced $V/2 \times V/2$ form of the Dirac operator makes determinant estimation significantly cheaper while keeping the additional computational overhead minimal. Half of the pseudofermion degrees of freedom completely decouple from the scalar field and can be discarded. The reduced operator partially retains the original periodic and antiperiodic boundary conditions when applied to the even or odd sublattices, respectively, reducing to a translation symmetry for even shifts. When utilizing affine coupling layers with a checkerboard mask, it is exactly this subset of the translational symmetry group that is preserved, and the translationally equivariant architecture is directly applicable to learning a distribution over the reduced subset of pseudofermions.

The improvement can be pushed to higher order by noting that the diagonal matrix elements of the preconditioned Dirac operator are close to unity, which makes it possible to employ an ILU preconditioning scheme [71,73]. It relies on the fact that the preconditioning matrices for the even-odd ordering step described above can be computed explicitly, which is not generally true for other ordering schemes. Though originally designed for the Wilson Dirac operator, the same procedure can be applied to the staggered fermions studied in this work. Since ILU preconditioning breaks the translation symmetry of the pseudofermion action completely, it is not directly compatible with any of our equivariant flow constructions that target the distribution of ϕ . However, in an experiment modeling the even-odd preconditioned ϕ -marginal distribution using an affine coupling layer model, additional ILU preconditioning for the same architecture led to a moderately improved acceptance rate.

⁴If one is interested in overall estimates of $\log Z$, the constant contributions from these terms must then be included.

APPENDIX C: STOCHASTIC ESTIMATOR FOR GRADIENTS OF $\log \det \mathcal{M}$

The calculation of loss gradients required for the optimization of some of the models introduced in this work requires the evaluation of gradients

$$\nabla_{\phi} \log \det \mathcal{M}(\phi) \quad (\text{C1})$$

taken with respect to the field ϕ . In general, $\mathcal{M}(\phi)$ is a positive definite matrix either arising from Dirac matrices of a pair of mass degenerate fermions as $\mathcal{M} = DD^{\dagger}$, or from one-flavor methods (see Sec. II B). The calculation of the exact determinant $\det \mathcal{M}(\phi)$ may be intractable because of the scaling with the number of lattice degrees of freedom, thus a stochastic estimator is instead defined in this section to tractably evaluate Eq. (C1).

By assumption, $\mathcal{M}$ is a positive-definite matrix and thus the following stochastic trace estimator is applicable:

$$\begin{aligned}\nabla \log \det \mathcal{M}(\phi) &= \nabla \text{Tr} \log \mathcal{M}(\phi), \\ &= \text{Tr}[\mathcal{M}(\phi)^{-1} \nabla \mathcal{M}(\phi)], \\ &= \mathbb{E}_{\chi \sim e^{-\chi^{\dagger} \chi}}[(\mathcal{M}^{-1}(\phi) \chi)^{\dagger} \nabla \mathcal{M}(\phi) \chi].\end{aligned}\quad (\text{C2})$$

Here, the noise vector χ is assumed to be drawn from the unit-variance isotropic distribution with an appropriate number of degrees of freedom to match the dimensions of $\mathcal{M}$.

In the case of two degenerate fermionic flavors, an interesting connection can also be made to the gradient of the negative pseudofermion action [where the sign is chosen to match the positive sign of $\log \det \mathcal{M}(\phi)$]. This gradient can be evaluated to be

$$\begin{aligned}\nabla(-\phi^{\dagger}(DD^{\dagger})^{-1}\phi) &= \phi^{\dagger}(DD^{\dagger})^{-1}(\nabla DD^{\dagger})(DD^{\dagger})^{-1}\phi, \\ &= \eta^{\dagger}(\nabla DD^{\dagger})\eta,\end{aligned}\quad (\text{C3})$$

where $\eta \equiv (DD^{\dagger})^{-1}\phi = (D^{\dagger})^{-1}\chi$, in terms of the noise vector $\chi \sim e^{-\chi^{\dagger} \chi}$ used to generate the pseudofermion field. A short derivation shows that this is equivalent to the stochastic estimator of the two-flavor determinant,

$$\begin{aligned}\text{Tr}[(DD^{\dagger})^{-1}(\nabla DD^{\dagger})] &= \text{Tr}[D^{-1}(\nabla DD^{\dagger})(D^{\dagger})^{-1}], \\ &= \mathbb{E}_{\chi \sim e^{-\chi^{\dagger} \chi}}[(D^{\dagger})^{-1}\chi)^{\dagger}(\nabla DD^{\dagger})(D^{\dagger})^{-1}\chi], \\ &= \mathbb{E}_{\chi \sim e^{-\chi^{\dagger} \chi}}[\eta^{\dagger}(\nabla DD^{\dagger})\eta].\end{aligned}\quad (\text{C4})$$

This relation allows the gradient estimator to be computed using the same tools utilized for the evaluation of HMC forces with respect to the pseudofermion action.

APPENDIX D: TRANSLATION-EQUIVARIANT NETWORKS

1. Convolutional networks with P fields and AP fields

In Sec. IV B 1, we introduced P fields and AP fields; that they form an algebra can be seen as follows. The set of P fields is stable under linear combinations and pointwise multiplications. On the other hand, the set of AP fields is only stable under linear combinations as the product of two AP fields is a P field, while the product of a P field with an AP field is an AP field. In other words, the set of P fields and AP fields forms a superalgebra [125] under pointwise addition and multiplication. Pointwise application of a function to a P field results in a new P field. For AP fields, more care is required as not all functions can be applied pointwise. Application of an odd function to an AP field results in another AP field, while pointwise application of an even function results in a P field. In this appendix, we explain in more detail how properties of those fields allow us to build expressive neural networks which are equivariant under translations. This means that if $\mathcal{T} \in \mathbb{Z}^d$ is an arbitrary spacetime translation and $f(\phi, \varphi) = \phi', \varphi'$ is one of our neural networks, then $f(\mathcal{T} \cdot \phi, \mathcal{T} \cdot \varphi) = \mathcal{T} \cdot \phi', \mathcal{T} \cdot \varphi'$. This discussion is not specific to two-dimensional spacetime, and applies for any dimension d .

First, convolutions can be built for both types of fields. For P fields, this is achieved by first padding the field using periodic padding, then applying a normal convolution. For a convolution with kernel shape $2k + 1$, all fields must be padded by k sites in each direction. As a concrete example, assume a one-dimensional lattice of size 5, a P field with values [1,2,3,4,5] and a convolution kernel [1,1,1]. The padded P field would be [5,1,2,3,4,5,1]. Applying the convolution would result in a new P field with values [8,6,9,12,10]. For AP fields, the only necessary change is to use antiperiodic padding along the time dimension, and periodic along the space dimension. Consider the two-dimensional example

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \quad (\text{D1})$$

and a 3×3 kernel with all weights equal to 1. The AP field should be padded by 1 site in each direction with signs applied to the temporal padding, giving

$$\begin{pmatrix} -9 & 7 & 8 & 9 & -7 \\ -3 & 1 & 2 & 3 & -1 \\ -6 & 4 & 5 & 6 & -4 \\ -9 & 7 & 8 & 9 & -7 \\ -3 & 1 & 2 & 3 & -1 \end{pmatrix}. \quad (\text{D2})$$

Applying the convolution gives the transformed AP field

$$\begin{pmatrix} 9 & 45 & 21 \\ 9 & 45 & 21 \\ 9 & 45 & 21 \end{pmatrix}. \quad (\text{D3})$$

The above construction of P and AP convolutions was illustrated with only one channel, but the extension to multiple channels is straightforward.

Any nonlinearity can be applied between convolutions for a P field without spoiling translational equivariance. As mentioned above, the LeakyReLU activation function is used throughout this work. For convolutions applied to an AP field, nonlinearities used as activation functions must be restricted to odd functions, for which we choose

$$\text{sign}(\varphi) \log(1 + |\varphi|). \quad (\text{D4})$$

With a P convolution, a bias can be applied along with a convolution at each step, since a bias is constant across all sites and thus transforms as a P field. However, a traditional bias cannot be applied to the convolution of an AP field without spoiling translational equivariance. For an AP field φ , a biaslike operation $\varphi \rightarrow \varphi + b \text{sign}(\varphi)$ in terms of a constant b can be applied instead. To avoid potential issues with the nondifferentiability of the sign function, we used a differentiable approximation given by applying $\varphi \rightarrow \varphi + b \tanh(\varphi/4)$.

All of the above constructions (P and AP convolutions, nonlinearities and biases) are equivariant with respect to translations on the lattice. By stacking them, we create expressive translation-equivariant neural networks. Our neural networks can also jointly transform pairs of P and AP fields. For example, we found the following transformation to work well:

Input: P-field P and AP-field A

$P' = \text{conv}(P)$

$A' = \text{conv}(A)$

$P'' = \text{concatenate}(P', |A'|)$

$A'' = \text{concatenate}(A', P'A')$

Output: P-field P'' and AP-field A''

The group of translational symmetries of a staggered action described in Sec. II C only includes translations by even numbers of lattice sites. Implementing symmetries in a network that are not symmetries of the target function restricts the expressivity and may make it difficult or impossible to represent an effective approximation of the target function by the network. For the models targeting the staggered fermion action in the study described in the main text, we avoided encoding translational symmetry by an odd number of sites by explicitly breaking equivariance with respect to odd translations. For example, to break the symmetry by odd translations along the first dimension of a field x , we fold its even and odd indices along the channel dimension; this doubles its number of channels while halving the number of points along the first dimension.

We then apply a convolution with stride 1 and kernel size 1, which mixes all the channels. Finally, we split the channels in two and fold them back along the first dimension to get a new field with the same lattice size as x . This approach mirrors the “squeezing” operation applied in Real NVP flows [49].

2. Explicit symmetrization by group averages

Equivariance of neural networks under discrete symmetry groups can also be achieved by explicitly averaging over the group. For discrete translations with antiperiodic boundary conditions in particular, this approach avoids the restriction to odd activation functions and vanishing biases, but requires an increase of the computational effort proportional to the timelike extension of the lattice.

Let $\mathcal{T}_{\vec{x},t}^a \in \mathbb{Z}^d$ denote a translation by $(\vec{x}, t)$ where antiperiodic boundary conditions are applied in time and periodic boundary conditions are applied in space, and let $\mathcal{T}_{\vec{x},t}^p$ denote a translation by $(\vec{x}, t)$ with periodic boundary conditions for all directions. The action of $\mathcal{T}^a$ and $\mathcal{T}^p$ is the same along all spatial dimensions, and we define $\mathcal{T}_{\vec{x}} = \mathcal{T}_{\vec{x},0}^p = \mathcal{T}_{\vec{x},0}^a$. For simplicity we now restrict to working with a two-dimensional $L \times L$ lattice with coordinates $(\vec{x}, t)$, but the following construction immediately generalizes to higher dimensions and nonsymmetric lattices. Both P fields and AP fields are maps from $\mathbb{Z}_L \times \mathbb{Z}_L$ to $\mathbb{R}^c$, where c is a number of channels. Under lattice translations, P fields are acted upon by $\mathcal{T}^p$, while AP fields are acted upon by $\mathcal{T}^a$.

Consider a function f that maps the pair (ϕ, φ) of a P field and an AP field to another field $f(\phi, \varphi)$. Assume that the output $f(\phi, \varphi)$ transforms with periodic boundary conditions along the space dimension, that is:

$$f(\mathcal{T}_{\vec{x}}\phi, \mathcal{T}_{\vec{x}}\varphi) = \mathcal{T}_{\vec{x}}f(\phi, \varphi). \quad (\text{D5})$$

Using averaging, we will now construct two maps u and v with the transformation properties:

$$u(\mathcal{T}_{\vec{x},t}^p\phi, \mathcal{T}_{\vec{x},t}^a\varphi) = \mathcal{T}_{\vec{x},t}^p u(\phi, \varphi), \quad (\text{D6})$$

$$v(\mathcal{T}_{\vec{x},t}^p\phi, \mathcal{T}_{\vec{x},t}^a\varphi) = \mathcal{T}_{\vec{x},t}^a v(\phi, \varphi). \quad (\text{D7})$$

We define⁵

⁵Note that if f is odd, then u below will be forced to be independent of φ . This can be avoided by either using non-odd nonlinearities, or by having nonzero biases.

$$u(\phi, \varphi) = \frac{1}{2L} \sum_{n=0}^{2L-1} \mathcal{T}_{0,-n}^p f(\mathcal{T}_{0,n}^p \phi, \mathcal{T}_{0,n}^a \varphi) \quad (\text{D8})$$

and

$$v(\phi, \varphi) = \frac{1}{2L} \sum_{n=0}^{2L-1} \mathcal{T}_{0,-n}^a f(\mathcal{T}_{0,n}^p \phi, \mathcal{T}_{0,n}^a \varphi). \quad (\text{D9})$$

We now wish to prove that the transformation properties in Eqs. (D6) and (D7) apply to these definitions. The proof is roughly the same for both cases,⁶ so we will only write it for Eq. (D6):

$$\begin{aligned} u(\mathcal{T}_{\vec{x},t}^p\phi, \mathcal{T}_{\vec{x},t}^a\varphi) &= \frac{1}{2L} \sum_{n=0}^{2L-1} \mathcal{T}_{0,-n}^p f(\mathcal{T}_{0,n}^p \mathcal{T}_{\vec{x},t}^p \phi, \mathcal{T}_{0,n}^a \mathcal{T}_{\vec{x},t}^a \varphi), \\ &= \frac{1}{2L} \sum_{n=0}^{2L-1} \mathcal{T}_{0,-n}^p f(\mathcal{T}_{\vec{x}} \mathcal{T}_{0,n+t}^p \phi, \mathcal{T}_{\vec{x}} \mathcal{T}_{0,n+t}^a \varphi), \\ &= \frac{1}{2L} \sum_{n=0}^{2L-1} \mathcal{T}_{0,-n+t}^p \mathcal{T}_{\vec{x}} f(\mathcal{T}_{0,n}^p \phi, \mathcal{T}_{0,n}^a \varphi), \\ &= \mathcal{T}_{\vec{x},t}^p u(\phi, \varphi). \end{aligned} \quad (\text{D10})$$

Equation (D10) was obtained using the change of variables $n \rightarrow n - t$ and the equivariance property given in Eq. (D5).

The functions u and v may be used to define equivariant affine coupling layers for the construction of equivariant flows. To achieve equivariance, the underlying function f needs to be evaluated $2L$ times instead of once, hence the aforementioned increase in computational cost. Since the masked affine couplings employed in this work already restrict the translational equivariance to multiples of two, one may also consistently use only every second term in the sums defining u and v without breaking the symmetry further, implying a factor L increase of the cost instead of $2L$. Still, the additional computational requirements are significant compared to the approach detailed in Appendix D 1, and for large-scale implementations one may have to partially trade equivariance against efficiency by excluding more terms from the sums.

⁶The proof is actually a particular instance of a more general property: if π_1 and π_2 are representations of a finite group G on vector spaces V and W , and if $f: V \rightarrow W$ is any map between these spaces, then $\frac{1}{|G|} \sum_{g \in G} \pi_2(g)^{-1} f(\pi_1(g))$ is an equivariant map from V to W .

- [1] C. Morningstar, in *21st Annual Hampton University Graduate Studies Program (HUGS 2006)* (2007), 2.
- [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] U. Wolff, *Nucl. Phys. B, Proc. Suppl.* **17**, 93 (1990).
- [11] S. Schaefer, R. Sommer, and F. Virotta, *Proc. Sci., LAT2009* (2009) 032 [[arXiv:0910.1465](https://arxiv.org/abs/0910.1465)].
- [12] S. Schaefer, R. Sommer, and F. Virotta (ALPHA Collaboration), *Nucl. Phys. B* **845**, 93 (2011).
- [13] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
- [14] J. Hoshen and R. Kopelman, *Phys. Rev. B* **14**, 3438 (1976).
- [15] U. Wolff, *Nucl. Phys. B* **300**, 501 (1988).
- [16] R. H. Swendsen and J.-S. Wang, *Phys. Rev. Lett.* **58**, 86 (1987).
- [17] R. G. Edwards and A. D. Sokal, *Phys. Rev. D* **38**, 2009 (1988).
- [18] U. Wolff, *Phys. Rev. Lett.* **62**, 361 (1989).
- [19] R. C. Brower and P. Tamayo, *Phys. Rev. Lett.* **62**, 1087 (1989).
- [20] M. Hasenbusch, *Nucl. Phys. B* **333**, 581 (1990).
- [21] R. Sinclair, *Phys. Rev. D* **45**, 2098 (1992).
- [22] W. Bietenholz, A. Pochinsky, and U.-J. Wiese, *Phys. Rev. Lett.* **75**, 4524 (1995).
- [23] H. G. Evertz, *Adv. Phys.* **52**, 1 (2003).
- [24] N. Prokof'ev and B. Svistunov, *Phys. Rev. Lett.* **87**, 160601 (2001).
- [25] N. Kawashima and K. Harada, *J. Phys. Soc. Jpn.* **73**, 1379 (2004).
- [26] R. Savit, *Rev. Mod. Phys.* **52**, 453 (1980).
- [27] V. Azcoiti, E. Follana, A. Vaquero, and G. Di Carlo, *J. High Energy Phys.* **08** (2009) 008.
- [28] C. Gattringer, T. Kloiber, and M. Müller-Preussker, *Phys. Rev. D* **92**, 114508 (2015).
- [29] E. P. Bernard, W. Krauth, and D. B. Wilson, *Phys. Rev. E* **80**, 056704 (2009).
- [30] M. Michel, J. Mayer, and W. Krauth, *Europhys. Lett.* **112**, 20003 (2015).
- [31] Y. Nishikawa, M. Michel, W. Krauth, and K. Hukushima, *Phys. Rev. E* **92**, 063306 (2015).
- [32] M. Hasenbusch and S. Schaefer, *Phys. Rev. D* **98**, 054502 (2018).
- [33] Z. Lei and W. Krauth, *Europhys. Lett.* **121**, 10008 (2018).
- [34] K. Zhou, G. Endrődi, L.-G. Pang, and H. Stöcker, *Phys. Rev. D* **100**, 011501(R) (2019).
- [35] D. Wu, L. Wang, and P. Zhang, *Phys. Rev. Lett.* **122**, 080602 (2019).
- [36] J. M. Pawłowski and J. M. Urban, *Mach. Learn. Sci. Tech.* **1**, 045011 (2020).
- [37] K. A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller, and P. Kessel, *Phys. Rev. E* **101**, 023304 (2020).
- [38] J. Carrasquilla, *Adv. Phys. X* **5**, 1797528 (2020).
- [39] L. Wang, Y. Jiang, L. He, and K. Zhou, [arXiv:2005.04857](https://arxiv.org/abs/2005.04857).
- [40] J. Vielhaben and N. Strodthoff, *Phys. Rev. E* **103**, 063304 (2021).
- [41] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 041101(R) (2017).
- [42] J. Liu, H. Shen, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 241104(R) (2017).
- [43] H. Shen, J. Liu, and L. Fu, *Phys. Rev. B* **97**, 205140 (2018).
- [44] Y. Nagai, H. Shen, Y. Qi, J. Liu, and L. Fu, *Phys. Rev. B* **96**, 161102(R) (2017).
- [45] C. Chen, X. Y. Xu, J. Liu, G. Batrouni, R. Scalettar, and Z. Y. Meng, *Phys. Rev. B* **98**, 041102(R) (2018).
- [46] Y. Nagai, M. Okumura, K. Kobayashi, and M. Shiga, *Phys. Rev. B* **102**, 041124(R) (2020).
- [47] Y. Nagai, A. Tanaka, and A. Tomiya, [arXiv:2010.11900](https://arxiv.org/abs/2010.11900).
- [48] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [49] L. Dinh, J. Sohl-Dickstein, and S. Bengio, [arXiv:1605.08803](https://arxiv.org/abs/1605.08803).
- [50] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *J. Mach. Learn. Res.* **22**, 1 (2021), <http://jmlr.org/papers/v22/19-1028.html>.
- [51] M. S. Alberg, G. Kanwar, and P. E. Shanahan, *Phys. Rev. D* **100**, 034515 (2019).
- [52] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, [arXiv:2007.07115](https://arxiv.org/abs/2007.07115).
- [53] L. Del Debbio, J. M. Rossney, and M. Wilson, *Phys. Rev. Lett.* **126**, 032001 (2021).
- [54] G. Kanwar, M. S. Alberg, D. Boyda, K. Cranmer, D. C. Hackett, S. Racanière, D. J. Rezende, and P. E. Shanahan, *Phys. Rev. Lett.* **125**, 121601 (2020).
- [55] D. Boyda, G. Kanwar, S. Racanière, D. J. Rezende, M. S. Alberg, K. Cranmer, D. C. Hackett, and P. E. Shanahan, *Phys. Rev. D* **103**, 074504 (2021).
- [56] D. H. Weingarten and D. N. Petcher, *Phys. Lett.* **99B**, 333 (1981).
- [57] J. Song, S. Zhao, and S. Ermon, [arXiv:1706.07561](https://arxiv.org/abs/1706.07561).
- [58] D. Levy, M. D. Hoffman, and J. Sohl-Dickstein, [arXiv:1711.09268](https://arxiv.org/abs/1711.09268).
- [59] M. Medvidovic, J. Carrasquilla, L. E. Hayward, and B. Kulchytsky, [arXiv:2012.01442](https://arxiv.org/abs/2012.01442).
- [60] S. Foreman, X.-Y. Jin, and J. C. Osborn, in *Proceedings of the 9th International Conference on Learning Representations* (2021), 5.

- [61] M. Gabri, G. M. Rotskoff, and E. Vanden-Eijnden, [arXiv:2105.12603](#).
- [62] C. Gattringer and C. B. Lang, *Quantum Chromodynamics on the Lattice* (Springer, Berlin, 2010), Vol. 788.
- [63] I. Montvay and G. Munster, *Quantum Fields on a Lattice*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, Cambridge, England, 1997).
- [64] H. J. Rothe, *Lattice Gauge Theories*, 4th ed. (World Scientific, Singapore, 2012).
- [65] S. P. Meyn and R. L. Tweedie, *Markov, Chains and Stochastic Stability* (Springer Science & Business Media, Cambridge, 2012).
- [66] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [67] W. K. Hastings, *Biometrika* **57**, 97 (1970).
- [68] A. D. Kennedy, I. Horvath, and S. Sint, *Nucl. Phys. B, Proc. Suppl.* **73**, 834 (1999).
- [69] K. Ogawa, T.-W. Chiu, and T.-H. Hsieh (TWQCD Collaboration), *Proc. Sci., LAT2009* (**2009**) 033 [[arXiv:0911.5532](#)].
- [70] Y.-C. Chen and T.-W. Chiu (TWQCD Collaboration), *Phys. Lett. B* **738**, 55 (2014).
- [71] P. de Forcrand and T. Takaishi, *Nucl. Phys. B, Proc. Suppl.* **53**, 968 (1997).
- [72] P. de Forcrand, *Nucl. Phys. B, Proc. Suppl.* **73**, 822 (1999).
- [73] M. J. Peardon, [arXiv:hep-lat/0011080](#).
- [74] M. Hasenbusch and K. Jansen, *Nucl. Phys. B, Proc. Suppl.* **106**, 1076 (2002).
- [75] M. A. Clark, *Proc. Sci., LAT2006* (**2006**) 004 [[arXiv:hep-lat/0610048](#)].
- [76] M. A. Clark and A. D. Kennedy, *Phys. Rev. Lett.* **98**, 051601 (2007).
- [77] T. A. DeGrand and P. Rossi, *Comput. Phys. Commun.* **60**, 211 (1990).
- [78] J. Kogut and L. Susskind, *Phys. Rev. D* **11**, 395 (1975).
- [79] A. S. Kronfeld, *Proc. Sci., LATTICE2007* (**2007**) 016 [[arXiv:0711.0699](#)].
- [80] C. Andrieu and G. O. Roberts, *Ann. Stat.* **37**, 697 (2009).
- [81] G. Bhanot, U. M. Heller, and I. O. Stamatescu, *Phys. Lett.* **129B**, 440 (1983).
- [82] M. Hutchinson, *Commun. Stat. Simul. Comput.* **19**, 433 (1990).
- [83] P. de Forcrand and L. Keegan, *Phys. Rev. E* **98**, 043306 (2018).
- [84] E. G. Tabak and E. Vanden-Eijnden, *Commun. Math. Sci.* **8**, 217 (2010).
- [85] E. G. Tabak and C. V. Turner, *Commun. Pure Appl. Math.* **66**, 145 (2013).
- [86] M. S. Albergo, D. Boyda, D. C. Hackett, G. Kanwar, K. Cranmer, S. Racanire, D. J. Rezende, and P. E. Shanahan, [arXiv:2101.08176](#).
- [87] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, *Phys. Rep.* **810**, 1 (2019).
- [88] S. Kullback and R. A. Leibler, *Ann. Math. Stat.* **22**, 79 (1951).
- [89] T. S. Cohen and M. Welling, [arXiv:1602.07576](#).
- [90] T. S. Cohen, M. Geiger, J. Khler, and M. Welling, in *Proceedings of the International Conference on Learning Representations* (2018).
- [91] T. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov, *Proceedings of Machine Learning Research PMLR*, Vol. 97, June (2019), pp. 1321–1330.
- [92] A. Tomiya and Y. Nagai, [arXiv:2103.11965](#).
- [93] S. Bulusu, M. Favoni, A. Ipp, D. I. Mller, and D. Schuh, *Phys. Rev. D* **104**, 074504 (2021).
- [94] M. Favoni, A. Ipp, D. I. Mller, and D. Schuh, [arXiv:2012.12901](#).
- [95] Y. LeCun, P. Haffner, L. Bottou, and Y. Bengio, in *Feature Grouping*, edited by D. Forsyth (Springer, New York, 1999).
- [96] L. Zhang, E. Weinan, and L. Wang, [arXiv:1809.10188](#).
- [97] C.-W. Huang, R. T. Chen, C. Tsirigotis, and A. Courville, [arXiv:2012.05942](#).
- [98] B. Amos, L. Xu, and J. Z. Kolter, in *Proceedings of the 34th International Conference on Machine Learning, PMLR* (2017), Vol. 70, pp. 146–155.
- [99] K. Dong, D. Eriksson, H. Nickisch, D. Bindel, and A. G. Wilson, [arXiv:1711.03481](#).
- [100] S. Ubaru, J. Chen, and Y. Saad, *SIAM J. Matrix Anal. Appl.* **38**, 1075 (2017).
- [101] D. P. Kingma and P. Dhariwal, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2018), Vol. 31.
- [102] O. Witzel, *Proc. Sci., LATTICE2018* (**2019**) 006 [[arXiv:1901.08216](#)].
- [103] J. Wambach, B.-J. Schaefer, and M. Wagner, *Acta Phys. Pol. B Proc. Suppl.* **3**, 691 (2010).
- [104] A. Frommer, V. Hannemann, B. Nockel, T. Lippert, and K. Schilling, *Int. J. Mod. Phys. C* **05**, 1073 (1994); [arXiv:hep-lat/9404013](#).
- [105] M. Luscher, *Comput. Phys. Commun.* **165**, 199 (2005).
- [106] C. Urbach, K. Jansen, A. Shindler, and U. Wenger, *Comput. Phys. Commun.* **174**, 87 (2006).
- [107] H. Hotelling, *Ann. Math. Stat.* **2**, 360 (1931).
- [108] N. Madras and A. D. Sokal, *J. Stat. Phys.* **50**, 109 (1988).
- [109] S.-H. Li and L. Wang, *Phys. Rev. Lett.* **121**, 260601 (2018).
- [110] H.-Y. Hu, S.-H. Li, L. Wang, and Y.-Z. You, *Phys. Rev. Research* **2**, 023369 (2020).
- [111] M. Lscher, *Commun. Math. Phys.* **293**, 899 (2010).
- [112] G. P. Engel and S. Schaefer, *Comput. Phys. Commun.* **182**, 2107 (2011).
- [113] H. Avron and S. Toledo, *J. ACM* **58**, 1 (2011).
- [114] J. Finkenrath, F. Knechtli, and B. Leder, *Comput. Phys. Commun.* **184**, 1522 (2013).
- [115] I. Han, D. Malioutov, and J. Shin, [arXiv:1503.06394](#).
- [116] J. Sohl-Dickstein, [arXiv:2005.06553](#).
- [117] R. A. Meyer, C. Musco, C. Musco, and D. P. Woodruff, [arXiv:2010.09649](#).
- [118] S. Passenheim and E. Hoogeboom, [arXiv:2012.13311](#).
- [119] <http://iaifi.org/>.
- [120] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell

- et al.*, *Proceedings of the 2018 IEEE High Performance Extreme Computing Conference (HPEC)* (2018), 1–6.
- [121] D. P. Kingma and J. Ba, [arXiv:1412.6980](#).
- [122] J. Zhang, T. He, S. Sra, and A. Jadbabaie, [arXiv:1905.11881](#).
- [123] A. F. Agarap, [arXiv:1803.08375](#).
- [124] B. Xu, N. Wang, T. Chen, and M. Li, [arXiv:1505.00853](#).
- [125] *Quantum Fields and Strings: A Course for Mathematicians. Vol. 1, 2*, edited by P. Deligne, P. Etingof, D. S. Freed, L. C. Jeffrey, D. Kazhdan, J. W. Morgan, D. R. Morrison, and E. Witten (1999).