

Nonperturbative trivializing flows for lattice gauge theories

Mathis Gerdes^{*}

Institute of Physics, University of Amsterdam, Amsterdam, The Netherlands

Pim de Haan[†]

CuspAI, Amsterdam, The Netherlands

Roberto Bondesan[‡]

Department of Computing, Imperial College London, Imperial College London, London, United Kingdom

Miranda C. N. Cheng[†]

Institute of Physics, University of Amsterdam, Amsterdam, The Netherlands,

Institute for Mathematics, Academia Sinica, Taipei, Taiwan

and Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Amsterdam, The Netherlands



(Received 2 June 2025; accepted 11 November 2025; published 26 November 2025)

Continuous normalizing flows are known to be highly expressive and flexible, which allows for easier incorporation of large symmetries and makes them a powerful computational tool for lattice field theories. Building on previous work, we present a general continuous normalizing flow architecture for matrix Lie groups that is equivariant under group transformations. We apply this to lattice gauge theories in two dimensions as a proof of principle and demonstrate competitive performance, showing its potential as a tool for future lattice computations.

DOI: [10.1103/31d5-hvp6](https://doi.org/10.1103/31d5-hvp6)

I. INTRODUCTION

Lattice gauge theory is a cornerstone of theoretical physics, providing nonperturbative insights into quantum field theories such as quantum chromodynamics (QCD). Traditional computational methods like Markov chain Monte Carlo [1] suffer from critical slowing down near phase transitions and in the continuum limit, leading to significant computational challenges [2].

Normalizing flows are machine learning models that produce samples from a target distribution by learning a bijective transformation between an easy-to-sample distribution and the target [3]. They have emerged as a fundamental approach to computing observables in lattice field theory, promising to accelerate computations by improving sampling efficiency and reducing autocorrelation times [4]. See [5,6] for an overview.

Explicitly preserving symmetries is crucial in designing efficient samplers based on normalizing flows. This is

because symmetries allow one to restrict the class of possible bijective transformations to be learned, thereby reducing the sample complexity of the task. The importance of symmetries is especially pronounced when simulating lattice gauge theories, where the size of the gauge symmetry group grows rapidly with the lattice size. Given their high design flexibility, continuous-time normalizing flows [7] are a promising framework. This is to be contrasted with discrete flows [4,8], where a partition of the lattice, with a partial breaking of the translation symmetry as a result, is required. For instance, in [9,10] continuous-time flows equivariant to spatial symmetries have been developed for scalar field theories with ϕ^4 coupling, reaching higher performance than their discrete-time counterpart. Other contexts in which continuous normalizing flows play a role in physics sampling include [11]. In the following, we shall simply refer to continuous-time flows as continuous flows and discrete-time flows as discrete flows.

For gauge theories, discrete normalizing flows, such as those employing residual networks [12–14] and coupling layers [8], have been developed to respect gauge symmetries. In view of the above-mentioned benefits, it is desirable to explore continuous realizations of normalizing flows for lattice gauge theories. Exactly such a method was theoretically proposed by Lüscher in the name of “trivializing maps” for gauge theories in the foundational work [15], before the advent of the modern machine learning

^{*}Contact author: m.gerdes@uva.nl

[†]On leave from CNRS, France.

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

technology. In particular, the Lüscher flow is driven by a potential that is assumed to admit a perturbative expansion in flow time. The specific Ansatz of the Lüscher flow then dictates that this perturbative expansion in flow time is also a perturbative expansion in the inverse coupling constant β [see (15)]. In particular, it has the form $\tilde{S} = \beta \sum_k (\frac{\beta t}{6})^k \tilde{S}^{(k)}$, where $\tilde{S}^{(k)}$ is independent of β . Following this idea, a machine learning-enhanced implementation of such a trivializing flow was recently reported in [16], showing promising results in sampling quality. While effective, this approach did not make full use of the general expressive power of neural networks, constraining itself instead to a specific form of expansion closely following Lüscher's original idea.

Building on these developments, we present a flexible and fully equivariant framework for continuous normalizing flows for lattice gauge theories. The goal of continuous flows is to learn the vector field driving the field transformation such that the endpoint of the flow gives a good approximation of the target density, in this case given in terms of the action. In contrast to the original Ansatz of Lüscher, the time dependence of the vector field in our flow is not constrained to be polynomial. Moreover, while retaining equivariance, we do not constrain the vector field to be given by the gradient of a scalar potential. Most importantly, our framework allows for general neural network architectures parametrizing the vector field, leading to more expressive machine-learning models. As a result, the vector field can in principle be an arbitrary function of the field configuration, limited only by the particular choice of architecture which has to be made in consideration of computational efficiency.

In particular, our framework includes the flow of Lüscher [15] and its machine learning implementation [16] as special cases. The additional flexibility of our framework gives us the freedom to choose a more or less complex neural network, as appropriate for the given application, balancing the trade-off between expressiveness and computational cost.

A feature of neural ordinary differential equation (ODE) allowing for efficient gradient computations is the adjoint sensitivity method, which we implemented in a general way. This approach allows for automatic differentiation and memory-efficient gradient computation for an arbitrary neural network architecture. This is combined with Crouch-Grossman Runge-Kutta integration schemes that can simultaneously process scalar and matrix group-valued degrees of freedom, ensuring numerical stability. Our implementation can be straightforwardly applied to any related integration problem involving real numbers and matrix groups, and any loss functions.

We used this framework to construct a novel continuous normalizing flow model for lattice gauge theories. Our method achieves state-of-the-art effective sample size—a metric quantifying the efficiency of the sampler—when

benchmarked on sampling two-dimensional SU(2) and SU(3) lattice gauge theories. The integration and Lie group infrastructure developed for this study forms part of the broader flow library `bijx` [17]. Numerical data for the figures can be found at [18].

II. GAUGE EQUIVARIANT CONTINUOUS FLOWS

The aim of normalizing flows is to approximate samples from a target distribution, which in the present context is given in terms of the actions S as

$$p(U) = \frac{e^{-S(U)}}{Z}, \quad (1)$$

where the partition function Z is a normalization constant. U is a random variable that for lattice gauge theories takes value in $SU(N)^{|E|}$ with the number of edges is given by $|E| = L^d d$ for a periodic square lattice with side length L in each of the d dimensions. To sample from p , we initially draw samples $U \sim \rho(U)$ from an easy-to-sample distribution, which we choose that associated to be the Haar measure DU . One then applies a bijective map $U \mapsto U' = \Phi(U)$ to obtain the output samples. This change of variables induces a probability distribution $q(U')$ on the proposed samples U' . To train the normalizing flow, crucially without having reliable samples from the target distribution, we minimize the reverse Kullback-Leibler (KL) divergence

$$\int DU q(U) \log \left(\frac{q(U)}{p(U)} \right) = \log Z + \mathbb{E}_{U \sim q} L(U), \quad (2)$$

where

$$L(U) = \log q(U) + S(U) \quad (3)$$

which we refer to as our loss function. Note that the partition function is a constant, and therefore drops out when computing the gradient with respect to the network parameters, and that we only need to generate samples from the model q and not the target distribution p . In practice, we will not be able to obtain a global minimum of the loss. We use the flow as the proposal distribution in a Metropolis-Hasting algorithm that guarantees asymptotic exactness of the sampling process [4], in importance sampling [19], or in conjunction with Hamiltonian Monte Carlo [20].

Continuous normalizing flows define the map Φ as the solution to an ODE, whose vector field is given by a neural network. In contrast to discrete flows, any sufficiently regular ODE is always invertible, and the change of density can be easily computed if the divergence of the vector field is numerically tractable. This approach leads to highly expressive architectures that can be made to manifestly respect the symmetries of the target theory, leading to higher parameter efficiency.

In the following, we will present architectures for gauge equivariant normalizing flows that achieve state-of-the-art performance on two-dimensional pure gauge theory for both SU(2) and SU(3). We will first introduce the necessary notations and background in the following section. For simplicity, we will initially suppress the edge index and consider a single copy of SU(N) (or a matrix Lie group G more generally), as the statements can be easily generalized to the lattice setup.

A. Ordinary differential equations on Lie groups

Let G be a matrix Lie group and $\mathfrak{g}$ the corresponding Lie algebra. In the following, we will focus on the case $G = \text{SU}(N)$, but most of the results straightforwardly apply to any matrix Lie group.

Any tangent vector $V \in T_U G$ at the point $U \in G$ can be decomposed in terms of a Lie algebra element $Z \in \mathfrak{g}$ as $V = ZU$. Consequently, a neural ODE on G is defined in terms of a neural network $Z_\theta: G \times \mathbb{R} \rightarrow \mathfrak{g}$, with parameters θ , as

$$\dot{U}(t) = Z_\theta(U, t)U(t). \quad (4)$$

The Lie algebra $\mathfrak{g} = \mathfrak{su}(N)$ is spanned by a set of basis vectors $\{T_a\}_{a=1}^{N^2-1}$, as SU(N) has dimensionality $N^2 - 1$. For instance, for $G = \text{SU}(2)$ we use $T_a = i\sigma_a$, where σ_a are the Pauli matrices, while for $G = \text{SU}(3)$ we use $T_a = i\lambda_a$, where λ_a are the Gell-Mann matrices. Requiring these vectors to form an orthonormal basis, the metric is fixed to have the normalization

$$\langle A, B \rangle = -\text{tr}(AB)/2. \quad (5)$$

Correspondingly, a basis for tangent space $T_U G$ at the point U is given by $\{T_a U\}_a$.

We can define the directional derivative of a function $f: G \rightarrow \mathbb{R}$ as

$$\partial_a f(U) = \left. \frac{d}{dt} \right|_{t=0} f(e^{tT_a} U). \quad (6)$$

The gradient of f , which yields a tangent vector at U , is then given as

$$\nabla f(U) = \sum_a \partial_a f(U) T_a U. \quad (7)$$

Consider now the situation when the initial point $U(0) \in G$ is sampled from a certain distribution q_0 on G , and evolves following the ODE of Eq. (4) up to time T . The resulting distribution q_T of $U(T)$, incorporating the change of density induced by the flow, also solves an ODE.

Denoting the log-likelihood as $\mathcal{L}(t, U) = \log q_t(U)$ and writing the right-hand side of the ODE (4) as $Z_\theta(U, t)U(t) = \sum_a Z_\theta^a(t, U)T_a U(t)$, we have [21,22]

$$\mathcal{L}(T, U) = \mathcal{L}(0, U) - \int_0^T \nabla \cdot \dot{U}(t, U) dt, \quad (8)$$

where the divergence can be computed using the Lie algebra basis as

$$\nabla \cdot \dot{U}(t, U) = \sum_a \partial_a Z_\theta^a(t, U). \quad (9)$$

In the form of an ODE, this is

$$\dot{\mathcal{L}} = -\nabla \cdot \dot{U}(t, U). \quad (10)$$

In the application in lattice gauge theories, U will be replaced by a collection of group elements $\{U_e\}_e$, one for each edge of the lattice.

B. Adjoint sensitivity method

The adjoint sensitivity method is a way of computing gradients through the solution of an ODE. Below we will give a general version of the adjoint sensitivity method. The generality of the mathematical Eqs. (11) and (12) here is reflected in our implementation of the method, which can be used for any loss and on any manifold as long as it can be embedded in $\mathbb{C}^n$. We also record the explicit form this takes in Eqs. (13) and (14) for the special case of integrating on matrix Lie groups with the KL divergence loss, as they are relevant for the lattice gauge theory application in the context of normalizing flows. Further details and explanations as well as the derivation of the specialized expressions can be found in the Appendix A, and alternative treatments can be found in [16,23,24].

For a general flow on a manifold $\mathcal{M}$ given by the ODE $\dot{z} = f_\theta(z)$, $z \in \mathcal{M}$, the gradient of the loss $L(z(T))$ given a fixed initial value $z(0)$ can be computed using

$$d_\theta L(z(T)) = - \int_T^0 a(t) \circ d_\theta f_\theta(t, z(t)) dt, \quad (11)$$

by integrating backward from T . This uses the auxiliary *adjoint state* $a(t)$, which is formally a cotangent vector of $\mathcal{M}$. It is initialized as $a(T) = d_z L|_{z(T)}$ and solved for simultaneously using

$$\dot{a}(t) = -a(t) \circ d_z f_\theta|_{z(t)}. \quad (12)$$

Here, we have used the more abstract notation from differential geometry. See for instance [25] for an introduction for physicists. In particular, for $\varphi: \mathcal{M} \rightarrow \mathcal{M}'$, $d\varphi$ denotes the corresponding differential map $d\varphi_x: T_x \mathcal{M} \rightarrow T_{\varphi(x)} \mathcal{M}'$ on the tangent spaces, “ $\circ$ ” denotes the composition of maps, and we routinely treat a cotangent vector $v \in T_x^* \mathcal{M}$ as a map $v \in T_x \mathcal{M} \rightarrow \mathbb{R}$.

For our particular problem, the adjoint state can be reduced to an element of the Lie algebra [i.e., for SU(N)]

embedded in $\mathbb{C}^n$ as a skew-symmetric traceless matrix], that follows the ODE

$$\dot{A} = [Z, A] + \nabla(\nabla \cdot \dot{U})U^{-1} - \sum_a A_a (\nabla Z^a)U^{-1}. \quad (13)$$

Again, it is initialized as $A_a(T) = \partial_a L(U(T))$, where we have specified to the loss L from (3), with $q = q_T$ the distribution at time T . After solving the ODE for A , the gradient of the loss is computed as

$$\partial_\theta L(U(T)) = - \int_T^0 \left(\langle A(t), \partial_\theta Z \rangle - \sum_a \partial_\theta \partial_a Z^a \right) dt. \quad (14)$$

A derivation of this from the general statement above can be found in the Appendix.

C. Numerical integration

Central to implementing the continuous flows for gauge theories is to define an integration scheme that works simultaneously for Lie group and scalar-valued degrees of freedom. The former is required for the field configurations themselves, while the latter arises for the probability densities as well as the gradients with respect to parameters in the adjoint sensitivity method. In principle, the adjoint state additionally introduces the tangent space, which we could identify as a Lie algebra element and treat separately. However, for our implementation we treat it simply as an element of the ambient complex space.

In principle, the Lie group elements themselves could also be treated as elements of the ambient space. However, a discrete integration scheme would inevitably result in paths leaving the target manifold. Including projection steps could address this, but the need to reliably estimate the change of density makes this naive approach unwieldy. We, therefore, implement a Crouch-Grossmann Runge-Kutta method [26,27], which can be defined for any matrix group and result in paths that remain in the desired submanifold up to numerical error. The algorithm is briefly outlined with further technical details in Sec. D of the Appendix.

D. Lattice gauge theory flow

In this section, we describe in some detail our choice for the neural vector field. In pure gauge theories on the lattice with gauge group G and square lattice $V_L \cong (\mathbb{Z}/L\mathbb{Z})^d$ with oriented edges connecting the adjacent vertices, the field configuration consists of $U_e \in G$ for each edge e . Reversing the direction of the edge gives $U_{-e} = U_e^{-1}$.

The target probability density is given by the Wilson action

$$S(U) = -\frac{\beta}{N} \sum_{x \in V_L} \sum_{\mu < \nu} \text{Re}[\text{tr} P_{\mu\nu}(x)], \quad (15)$$

where $P_{\mu\nu}(x)$ are the plaquettes, namely the product of U along edges forming the smallest closed loops. Here, x denotes the initial vertex, and μ and ν are the two directions that the loop spans. We only include one orientation by requiring $\mu < \nu$ with both μ and ν pointing in the positive direction. An illustration is given in Fig. 1 (0).

For any given $\Omega: V \rightarrow G$ we denote the corresponding gauge transformation as $\Omega \cdot U$. For instance, we have

$$U_e \mapsto (\Omega \cdot U)_e = \Omega_{x_i} U_e \Omega_{x_f}^{-1} \quad (16)$$

where $x_{i/f}$ represent the starting/end point of the edge e . Note that the traces of Wilson loops, and hence the Wilson action, are invariant under the above gauge transformation.

The guiding idea for constructing our continuous flow architecture is that the gradients of gauge-invariant quantities, such as traces of loops $\mathcal{W}$, are gauge equivariant,

$$\nabla_e \mathcal{W} \xrightarrow{U \mapsto \Omega \cdot U} \Omega_{x_i} (\nabla_e \mathcal{W}) \Omega_{x_f}^{-1}, \quad (17)$$

which is the precisely property we require of the vector field $Z_\theta(U, t)U(t)$ defining the ODE (4) for our continuous normalizing flows.

In terms of the directional derivatives, i.e. transporting the above back by right multiplication to the Lie algebra, this is equivalent to

$$\partial_a^e \mathcal{W} \mapsto \Omega_{x_i} (\partial_a^e \mathcal{W}) \Omega_{x_i}^{-1}. \quad (18)$$




FIG. 1. Loops used in the ODE architecture. The first two loops, highlighted in red, are used for the nonlinear superposition function. The anchor in the dual lattice is displayed as a dot.

If the ODE is gauge equivariant, then transformation commutes with the normalizing flow map. Since the prior we choose (the Haar measure) is by definition gauge invariant, the generated probability distribution will automatically be gauge invariant, namely $q_T(U) = q_T(\Omega \cdot U)$, just like the target Wilson action.

We will now use the Wilson loops as building blocks of our vector field. To build a sufficiently expressive ODE, we must use a larger set of loops than just the plaquettes. We denote the traces of loops of some (to be specified below) class (k) as $\mathcal{W}_{(k)}$. Figure 1 shows the specific classes of loops we consider for our experiments in two dimensions. Instead of the lattice vertices, we will use elements of the dual lattice as anchor points of the Wilson loops, denoted as $\bar{x}$ and shown in the figure as dots, as this better reflects the symmetries of the loops. In $\mathcal{W}_{(k)}$ we include all spatial alignments and shifts that keep the original anchor point enclosed in the loop and leave the orientation of the loop invariant. For instance, for each anchor point $\bar{x}$ in two dimensions, $\mathcal{W}_{(0)}$ gives rise to one trace value, namely the plaquette $\mathcal{W}_{(0)}^1 = \text{tr } P_{01}$, while $\mathcal{W}_{(1)}$ and $\mathcal{W}_{(3)}$ each gives rise to four (corresponding to up, down, left, right).

Concretely, we build the neural ODE by taking a superposition of the vectors $\partial_a^e \mathcal{W}$ with coefficients depending on gauge invariant quantities, which in particular include any element of $\mathcal{W}_{(k)}$. To allow for efficient computation of the divergence, which is needed to compute the change of density, we use a shallow architecture that first applies a neural network locally in $\bar{x}$ to a subset of $\mathcal{W}$, followed by a convolution in $\bar{x}$ that introduces nonlocal correlations. This way the divergence computation can be done in $\mathcal{O}(|E|)$ time, while for general gauge invariant coefficients, it would take $\mathcal{O}(|E|^2)$ time. See Sec. C of the Appendix. Writing, as before, $\dot{U}_e = \sum_a Z_e^a(t, U) T_a U_e$, we define

$$Z_e^a(t, U) = \sum_{k, \bar{x}} (\partial_a^e \mathcal{W}^k(\bar{x})) \Lambda_{\bar{x}}^k(t, \mathcal{W}), \quad (19)$$

with $\Lambda(t, \mathcal{W})$ a time-dependent convolution of features obtained from $\mathcal{W}$ by locally applying a neural network at each anchor point $\bar{y}$,

$$\Lambda_{\bar{x}}^k = \hat{\Lambda}^k(t) + \sum_{j, \bar{y}} C_{\bar{x}\bar{y}}^{kj}(t) g_{\bar{y}}^j(t, \mathcal{W}(\bar{y})), \quad (20)$$

where $C(t)$ is a time-dependent convolution kernel, $\hat{\Lambda}^k(t)$ are field-independent learnable parameters, and g represents a neural network.

We found that this parametrization is expressive while still allowing for a tractable divergence computation, as required for the change of density computation. Further details of the architecture used in our experiments can be found in Sec. III B, where it is applied to two-dimensional pure gauge theories.

TABLE I. Table of coefficients $c(i)$ for the conjugation equivariant target actions of Eq. (21). These are the same values as used in [8].

Target (i)	$c_1^{(i)}$	$c_2^{(i)}$	$c_3^{(i)}$
0	1	0	0
1	0.17	-0.65	1.22
2	0.98	-0.63	-0.21

III. NUMERICAL EXPERIMENTS

A. Single-variable conjugation equivariant flows

To highlight the flexibility of continuous flows and their expressive power, we first repeat the experiment in [8] and use flows to learn a conjugation-invariant density on a single copy of $\text{SU}(N)$. Specifically, we use three different families of target actions of the form

$$S^{(i)}(U) = -\frac{\beta}{N} \text{Re tr}(c_n^{(i)} U^n) \quad (21)$$

with the sets of values $c^{(i)}$ as shown in Table I.

We parametrize the ODE using a time-dependent potential denoted P_θ , which we learn as a neural network,

$$\dot{U} = \nabla P_\theta(t, U). \quad (22)$$

The gradient and the Laplacian, needed for the change of density computation following (10), are both computed using automatic differentiation, which is feasible as we only consider a single $\text{SU}(N)$ matrix here.

To ensure that the defined flow is conjugation equivariant, we define the potential exclusively in terms of conjugation invariant quantities. First, we construct 64 time features by applying the sin function with individually learnable frequencies to t , followed by a dense layer with 64 features. For U -dependent features, we use absolute, real, and imaginary parts of $\text{tr}(U^n)$ for powers $n = 1 \dots n_{\max}$. Specifically, we use

$$|\text{tr}(U^n)|, \text{Re tr}(U^n), (\text{Re tr}(U^n))^2, \\ \text{Im tr}(U^n), (\text{Im tr}(U^n))^2. \quad (23)$$

The time and U -dependent features are then concatenated and processed by a ResNet of width 512 and activation function GELU before being projected to a single real value representing the potential P_θ .

For target $c^{(1)}$ we used ResNet depth 3 and $n_{\max} = 2$, while for the other two targets we used depth 3 and $n_{\max} = 4$.

Leveraging the continuous nature of our normalizing flows, we train a single potential for each type of target and match the flow time t with the target inverse temperature β as $\beta(t) = 9 \cdot t$. For training, we use four randomly drawn end times T and corresponding betas uniformly in




FIG. 2. Comparison of the $SU(3)$ densities at $\beta = 9$ in terms of angular coordinates for the three conjugation equivariant targets listed in Table I and the respectively trained continuous normalizing flows.




FIG. 3. Effective sample size along the flow time of trained normalizing flows for each of the three families of target distributions listed in Table I in comparison to the ESS achieved by the individual discrete flows from [8].

$T \in [0.1, 1]$ and evaluate the KL loss on 1024 samples for each.

The densities evaluated in angular coordinates (see Appendix E for details) of the flow shown in Fig. 2 match the target densities visually to high accuracy. Similarly, Fig. 3 shows that the flows are able to approximate the target density for the whole range of target β values to high accuracy as measured by effective sample size (ESS).

B. Pure gauge theory in two dimensions

To demonstrate the feasibility of the model, and in particular the model architecture described in Eq. (19), we performed experiments on two-dimensional pure gauge theory (15) for $SU(2)$ and $SU(3)$, achieving state-of-the-art or better effective sample sizes. Table II shows the result of training for a 16×16 lattice, comparing to previous discrete [8] and continuous [16] architectures. Table III shows results for $SU(3)$ on a smaller 8×8 lattice but for larger values of β .

TABLE II. Effective sample sizes for normalizing flows for $SU(2)$ and $SU(3)$ on 16×16 lattice compared to previous work. Best values in bold.

ESS (%)	SU(2)		SU(3)		
	$\beta = 2.2$	$\beta = 2.7$	$\beta = 5$	$\beta = 6$	$\beta = 8$
Continuous flow	87	68	86	76	23
Bacchio <i>et al.</i> [16]	...	...	88	70	...
Boyd <i>et al.</i> [8]	80	56	75	48	...

TABLE III. Effective sample sizes for normalizing flows for $SU(3)$ on 8×8 lattice compared to previous work. Best values in bold.

ESS (%)	$\beta = 8$	$\beta = 12$
Continuous flow	64	27
Multiscale + flow [28]	35	13
Haar + flow [28]	25	3

In contrast to previous work [8,16], we have found that initially using a smaller batch size of 32 and increasing this after some steps helps to speed up training. Concretely, we train using a second-order Crouch-Grossmann integrator with 40 steps in 32-bit mode until training approximately converged, then switching to 64-bit mode and a batch size of 128. ESS evaluation was done in each case in 64-bit mode and using 40 steps of a third-order Crouch-Grossmann method.

We will now outline in more detail the architectures used for $SU(2)$ and $SU(3)$ flows, respectively.

1. Flows on $SU(2)$ lattices

When applying the general form of the architecture in (19) above to the specific experiments described here, we only used the first two classes of loops, shown in red in Fig. 1, as input to the neural network that defines the superposition coefficients Λ . Moreover, as we sum over the anchor points $\bar{x}$, to avoid redundancy we do not include all shifted copies of the Wilson loops in the gradients $\nabla_e \mathcal{W}$. We therefore denote by $\tilde{\mathcal{W}}_{(i)}$ the subset of traces $\mathcal{W}_{(i)}$ where we pick one representative of each set of shifted loops. For example, $\tilde{\mathcal{W}}_{(1)}(\bar{x})$ contains two instead of four values for each anchor point.

Let $\mathcal{W}_{(i_1:i_2)}$ denote the union of $\mathcal{W}_{(i_1)}, \dots, \mathcal{W}_{(i_2)}$. The architecture of the neural ODE used in our experiments is then given by

$$Z_e^a(t, U) = \sum_{k, \bar{x}} \left(\partial_a^e \tilde{\mathcal{W}}_{(0:7)}^k(\bar{x}) \right) \Lambda_{\bar{x}}^k(t, \mathcal{W}_{(0:1)}), \quad (24)$$

where the superposition coefficients are

$$\Lambda_{\bar{x}}^k = \hat{\Lambda}^k(t) + \sum_{j,\bar{y}} C_{\bar{x}\bar{y}}^{kj}(t) g_{\bar{y}}^j(t, \mathcal{W}_{(0:1)}(\bar{y})). \quad (25)$$

Note that for SU(2) all traces $\mathcal{W}$ are real numbers, and can be directly input into conventional dense neural networks.

We parametrize g as a ResNet of depth 3, width 128, and with GELU activation function. The time dependence is introduced by first applying a set of 16 uniformly shifted Gaussian density functions to t followed by a single dense layer with output size 16 to generate a set of features $K_g^l(t)$ with $l = 1, \dots, 16$.

The convolution kernel is chosen to be of size 5×5 . The time dependence is introduced by having 16 copies of parameters and contracting them with $K_C^l(t)$ which is constructed as above but with independently trainable parameters in the dense layer. The number of output features of g and thus input features of C , indexed by j in Eq. (25), are chosen to be 24.

2. Flows on SU(3) lattices

The architecture for SU(3) is qualitatively the same as for SU(2). However, in this case, the traces of loops are generally complex. For the input of the network g we use $\text{Re}\mathcal{W}_{(0:1)}$ and $(\text{Im}\mathcal{W}_{(0:1)})^2$. The latter is chosen to reflect the parity symmetry of the theory. Similarly, we have two output channels in the convolution corresponding to the imaginary and real parts of the gradients $\partial\mathcal{W}$. The final form of the architecture is then

$$\begin{aligned} Z_e^a(t, U) = & \sum_{k,\bar{x}} \text{Re} \left(\partial_a^e \tilde{\mathcal{W}}_{(0:7)}^k(\bar{x}) \right) \Lambda_{\bar{x}}^{(R)k}(t, \mathcal{W}_{(0:1)}) \\ & + \text{Im} \left(\partial_a^e \tilde{\mathcal{W}}_{(0:7)}^k(\bar{x}) \right) \left(D^{kj} \Im \mathcal{W}_{(0:1)}^j \right) \\ & \times \Lambda_{\bar{x}}^{(I)k}(t, \mathcal{W}_{(0:1)}), \end{aligned} \quad (26)$$

where we ensure parity symmetry by introducing a dense matrix D , and with the two superposition factors similar to the above,

$$\Lambda_{\bar{x}}^{(*)k} = \hat{\Lambda}^{(*)k}(t) + \sum_{j,\bar{y}} C_{\bar{x}\bar{y}}^{(*)kj}(t) g_{\bar{y}}^j(t, \mathcal{W}_{(0:1)}(\bar{y})). \quad (27)$$

The time dependence and network architecture are the same as those specified above for SU(2).

IV. OUTLOOK

In this work, we introduced a gauge-equivariant continuous normalizing flow framework for lattice gauge theories. Our method demonstrates competitive ESS when applied to two-dimensional SU(2) and SU(3) pure gauge theories, outperforming previous discrete and continuous approaches.

In machine learning literature, the idea of guiding the continuous flow to achieve more stability and accuracy has been discussed in various works [29,30]. In our experiment on a single Lie group flow, we have briefly explored the method matching the flow time to a predetermined evolution of target density and obtained encouraging preliminary results. It would be extended to lattice gauge theories, connecting to related work of stochastic normalizing flows [31] that have shown promising results also in higher spatial dimensions [32,33].

Scaling up normalizing flows to larger lattice sizes remains a challenge [34,35]. We believe that the broad class of continuous equivariant normalizing flows developed in this paper constitutes a promising tool for achieving large-scale flow-based lattice simulations. To facilitate scaling, the incorporation of methods such as multiscale sampling [28], transfer learning [10], and other hierarchical generation and transfer learning ideas should be further explored.

In this work, we have limited ourselves to simple dense networks when designing neural networks employed to construct the neural ODE. It would be interesting to explore other network architectures, including the recently proposed gauge equivariant convolutional networks [36], in the future. Finally, it will be useful to extend the present work to include field theories with fermions.

ACKNOWLEDGMENTS

We thank Niki Stratikopoulou for helpful discussions. The research of M. G. and M. C. N. C. is supported by the Vidi grant (Grant No. 016.Vidi.189.182) from the Dutch Research Council (NWO). M. G. was partially supported by a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 864035).

DATA AVAILABILITY

The data that support the findings of this article are openly available [18].

APPENDIX A: ADJOINT SENSITIVITY METHOD

For a smooth manifold $\mathcal{M}$, we consider the ODE $\dot{z} = f_{\theta}(t, z) \in T_z\mathcal{M}$ where θ represents some set of trainable parameters. Let Φ be the associated flow map with the properties

$$\left. \frac{d}{ds} \Phi_{t,s}(z) \right|_{s=0} = f_{\theta}(t, z), \quad (\text{A1a})$$

$$\Phi_{t+s,s'} \circ \Phi_{t,s} = \Phi_{t,s+s'}, \quad (\text{A1b})$$

and with the boundary condition $\Phi_{t,0}(z) = z$. From the above, it is clear that the map Φ depends on the trainable parameters θ through the form of the ODE.

Given a loss function $L: \mathcal{M} \rightarrow \mathbb{R}$, we are interested in computing $\partial_\theta L(z(T))$ where $z(T) = \Phi_{0,T}(z(0))$ given a fixed initial $z(0)$. At the final time T , the derivative $dL|_{z(T)}$ measures the sensitivity of the loss to small changes in the final value $z(T)$. Formally, it is a cotangent vector $dL|_{z(T)} \in T_{z(T)}^* \mathcal{M}$ which linearly maps vectors in $T_{z(T)} \mathcal{M}$ to a real number. To compute it, one needs to take into account the effect of changing the ODE at all earlier times $t < T$. To this end, we define an auxiliary *adjoint state* $a(t) \in T_{z(t)}^* \mathcal{M}$ defined as follows:

$$a(t) := d(L \circ \Phi_{t,T-t})|_{z(t)} = dL|_{z(T)} \circ d\Phi_{t,T-t}|_{z(t)}. \quad (\text{A2})$$

Like $dL|_{z(T)}$, the cotangent vector $a(t)$ is a linear map from a tangent vector in $T_{z(t)} \mathcal{M}$ to a real number. It captures the sensitivity of the loss to (infinitesimally) perturbing $z(t)$ in the direction of the tangent vector, then flowing with the ODE to time T , and evaluating the loss. More explicitly, this means

$$a(t)(X) = dL|_{z(T)}(d\Phi_{t,T-t}(X)) \quad (\text{A3})$$

for all $X \in T_{z(t)} \mathcal{M}$, where

$$d\Phi_{t,T-t}|_{z(t)}: T_{z(t)} \mathcal{M} \rightarrow T_{z(T)} \mathcal{M} \quad (\text{A4})$$

is the differential map associated to $\Phi_{t,T-t}$. Formally, $a(t)$ is the pullback of dL under the flow map Φ : $a(t) = \Phi_{t,T-t}^* dL$.

To obtain an equation for $a(t)$, we take the derivative with respect to t . First, observe that

$$a(t-s) = d(L \circ \Phi_{t,T-t} \circ \Phi_{t-s,s}) = a(t) \circ d\Phi_{t-s,s} \quad (\text{A5})$$

which follows from the composition property (A1b) and the chain rule. Then, using that the derivatives with respect to z and t commute and property (A1) we can take the time derivative to get,

$$\begin{aligned} \dot{a}(t) &= -\frac{d}{ds} a(t-s) \Big|_{s=0} \\ &= -a(t) \circ d \left(\frac{d}{ds} \Phi_{t-s,s} \Big|_{z(t-s)} \right) \Big|_{s=0} \\ &= -a(t) \circ df_\theta|_{z(t)}. \end{aligned} \quad (\text{A6})$$

Note that we also use here that the tangent space is a vector space, and as such its tangent space can be identified with itself. If we let $z(t)$ vary in $\mathcal{M}$ and consider a as a section of the cotangent bundle, $\dot{a}$ then lies in its vertical bundle.

Using this, we can solve for the adjoint state by integrating the ODE (A6) backward in time, starting from $z(T)$ and the initial value $a(T) = dL|_{z(T)}$. Given the adjoint state, the derivative with respect to parameters is computed simultaneously as

$$d_\theta L(z(T)) = - \int_T^0 a(t) \circ d_\theta f_\theta(t, z(t)) dt, \quad (\text{A7})$$

as explained in [21].

1. Lie-algebra valued adjoint state

We will show how the explicit equation of (13) for the adjoint state, represented as a Lie algebra element, follows from the general expression in Eq. (12).

We consider an ODE on $\mathcal{M} = G \times \mathbb{R}$ whose elements we denote $z = (z^{(1)}, z^{(2)}) = (U, \mathcal{L})$ where $\mathcal{L} \in \mathbb{R}$ represents the log-likelihood. In particular we may have $G = \text{SU}(N)^{|E|}$, but to improve legibility we will suppress the lattice index on $U \in G$ below.

The ODE $\dot{z} = f_\theta$ is given by $f_\theta = (\dot{U}, \dot{\mathcal{L}})$, and using the ODE for $\mathcal{L}$ of Eq. (10) we have,

$$\dot{U} = Z(t, U)U = \sum_a Z^a(t, U)T_a U, \quad (\text{A8a})$$

$$\dot{\mathcal{L}} = -\nabla_U \cdot \dot{U} = -\sum_a \partial_a Z^a(t, U). \quad (\text{A8b})$$

We integrate this forward in time from $t = 0$ to $t = T$ with the initial $U(0)$ drawn from the Haar prior $q_0(U) = \rho(U)$ and $\mathcal{L}(0) = \log \rho(U(0))$. Our loss function has the specific form

$$L(U, \mathcal{L}) = S(U) + \mathcal{L}, \quad (\text{A9})$$

which we evaluate at $z(T) = (U(T), \mathcal{L}(T))$. Note that as our manifold $\mathcal{M}$ is a product space, its cotangent space at any given point z can be split into an element in $T_{z^{(1)}}^* G$ and $T_{z^{(2)}}^* \mathbb{R} \cong \mathbb{R}$, and we can use the notation d_U and $d_{\mathcal{L}}$ to indicate the corresponding operators in the two spaces. Thus, we have

$$\begin{aligned} a(T) &= a^{(1)}(T) + a^{(2)}(T), \\ a^{(1)}(T) &= d_U S, \quad a^{(2)}(T) = d_{\mathcal{L}} \mathcal{L} = 1. \end{aligned} \quad (\text{A10})$$

We have $df_\theta = (d\dot{U}, d\dot{\mathcal{L}})$ and both vector fields only have explicit dependence on U , so we have $df_\theta = d_U f_\theta$. To obtain the specific form of $\dot{a}(t)$ from the general form of Eq. (A6) we compute

$$d_U \dot{U}_e = d_U(ZU) = (d_U Z)U + Zd_U U, \quad (\text{A11a})$$

$$d_U \dot{\mathcal{L}} = -d_U(\nabla_U \cdot \dot{U}). \quad (\text{A11b})$$

This implies that the adjoint state satisfies

$$\begin{aligned}\dot{a} &= -a \circ d_z f_\theta \\ &= -a \circ (d_U(ZU), 0) + a \circ (0, d_U(\nabla_U \cdot \dot{U})) \\ &= \dot{a}^{(1)} + \dot{a}^{(2)},\end{aligned}\quad (\text{A12})$$

where we split $\dot{a}$ into

$$\dot{a}^{(1)} = -a^{(1)} \circ d_U(ZU) + a^{(2)} \circ d_U(\nabla_U \cdot \dot{U}), \quad (\text{A13a})$$

$$\dot{a}^{(2)} = 0, \quad (\text{A13b})$$

as the component in the vertical bundle of T^*G and $T^*\mathbb{R}$, respectively. Thus, given the initial condition $a^{(2)}(T) = 1$ we have $a^{(2)}(t) = 1$ for all times, and we are left with

$$\dot{a}^{(1)} = -a^{(1)} \circ d_U(ZU) + d_U(\nabla_U \cdot \dot{U}). \quad (\text{A14})$$

Recall that $a^{(1)}(t) \in T_{U(t)}^*G$. In view of the decomposition mentioned above (4) and given $U(t)$, we can equivalently define $\tilde{a}^{(1)} \in \mathfrak{g}^*$ such that for all $v \in \mathfrak{g}$ we have $\tilde{a}^{(1)}(v) = a^{(1)}(vU)$, namely $\tilde{a}^{(1)} = a^{(1)} \circ r_U$ where r_U denotes right multiplication by $U = U(t)$. We thus have

$$\frac{d}{dt} \tilde{a}^{(1)} = \dot{a}^{(1)} \circ r_U + a^{(1)} \circ r_{ZU}, \quad (\text{A15})$$

where we have used $\dot{U} = ZU$. Using the scalar product which gives an isomorphism $\mathfrak{g}^* \rightarrow \mathfrak{g}$, we can represent $\tilde{a}^{(1)}(t)$ as a Lie algebra element $A(t)$ such that

$$\tilde{a}^{(1)}(t)(v) = a^{(1)}(t)(vU) = \langle A(t), v \rangle \quad \forall v \in \mathfrak{g}. \quad (\text{A16})$$

Similarly, defining the gradient $\nabla_U h \in T_U \mathcal{G}$ of a scalar function h as

$$d_U h(vU) = \langle \nabla_U h U^{-1}, v \rangle, \quad (\text{A17})$$

the initial condition is then given by $A(T) = \nabla_U S$. In terms of the explicit orthonormal basis $\{T_a\}$, we have $A(T) = \sum_a (\partial_a S) T_a U$.

Following a straightforward calculation using the above relations and expressing the bilinear form as trace (5), we finally obtain the adjoint state ODE

$$\dot{A} = [Z, A] + \nabla_U(\nabla_U \cdot \dot{U})U^{-1} - \sum_a A_a \nabla_U Z^a U^{-1}, \quad (\text{A18})$$

where $[Z, A]$ is the commutator and the a index is the component with respect to the Lie algebra basis. In terms of explicit basis $\{T_a\}_a$ for $\mathfrak{g}$ and the dual basis $\{T^a\}_a$ for $\mathfrak{g}^*$, we have $\tilde{a}^{(1)}(t) = \sum_a A_a T^a$ and $A(t) = \sum_a A_a T_a$. Then (A15) acting on $v \in \mathfrak{g}$ gives

$$\begin{aligned}\sum_a v^a \dot{A}_a &= \sum_{a,b} (-A_a v^b \partial_b Z^a + \partial_a \partial_b Z^b) \\ &\quad - \sum_a A_a \langle T_a, [v, Z] \rangle,\end{aligned}$$

which becomes, upon using the cyclic property of the trace,

$$\dot{A}_a = \sum_b (-A_b \partial_a Z^b + \partial_a \partial_b Z^b) + \langle T_a, [Z, A] \rangle \quad (\text{A19})$$

which is precisely (A18). To translate the general loss integral (11), repeated here,

$$d_\theta L(z(T)) = - \int_T^0 a(t) \circ d_\theta f_\theta(t, z(t)) dt, \quad (\text{A20})$$

into the specialized expression (14), we simply specialize to the case $df_\theta = (d\dot{U}, d\dot{Z})$. With $\dot{U} = ZU$ we get

$$\begin{aligned}a^{(1)} \circ d_\theta \dot{U} &= \langle A(t), \partial_\theta Z \rangle \\ a^{(2)} \circ d_\theta \dot{Z} &= - \sum_a \partial_\theta \partial_a Z^a\end{aligned}\quad (\text{A21})$$

and thus the loss is computed as

$$\begin{aligned}\partial_\theta L(z(T)) &= \int_T^0 (-\langle A(t), \partial_\theta Z \rangle + \sum_a \partial_\theta \partial_a Z^a) dt \\ &= \sum_a \int_T^0 (-A_a(t) \partial_\theta Z^a + \partial_\theta \partial_a Z^a) dt.\end{aligned}\quad (\text{A22})$$

APPENDIX B: GRADIENT IMPLEMENTATION

In practice, we implement the adjoint sensitivity method under the assumption that $\mathcal{M}$ is embedded in $\mathbb{C}^n$ for some integer n , and that we represent the values as elements of this ambient space. This is in particular the case for matrix Lie groups such as $SU(N)$, product manifolds thereof, and the real numbers that represent the log-density of samples. The implementation makes heavy use of JAX's [37] Jacobian-vector product (JVP, also known as forward mode) and vector-Jacobian products (VJP, also known as reverse mode). Given a function $f: \mathbb{C}^m \mapsto \mathbb{C}^n$, these generate new functions

$$\begin{aligned}\text{JVP}(f): T\mathbb{C}^m &\rightarrow T\mathbb{C}^n, & v &\mapsto df(v), \\ \text{VJP}(f): T^*\mathbb{C}^n &\rightarrow T^*\mathbb{C}^m, & c &\mapsto c \circ df.\end{aligned}\quad (\text{B1})$$

Numerically $T^*\mathbb{C}^m \rightarrow \cong T\mathbb{C}^m \cong \mathbb{C}^m$, and the action of a dual space element $c \in T^*\mathbb{C}^m$ on $v \in T\mathbb{C}^m$ is simply the inner product $c \cdot v \in \mathbb{C}$ on $\mathbb{C}^m$. In particular, for a real function $h: \mathbb{C}^m \mapsto \mathbb{R}$, the cotangent dh can be obtained via VJP with cotangent $1 \in \mathbb{C} \cong T^*\mathbb{C}$ as input. With these operations, the adjoint sensitivity method can thus be implemented in full generality.

1. Gradients and divergence of loops

For traces of loops, the gradient and second derivatives needed to compute the neural ODE and its divergence can be efficiently computed by inserting Lie group generators into the product of group elements U_e labeled by the edges e .

If a loop has the form

$$\mathcal{W} = \text{tr}(U_{e_1} \dots U_{e_i} \dots U_{e_n}), \quad (\text{B2})$$

the directional derivative $\partial_a^{e_i}$ acts on U_{e_i} as $\partial_a^{e_i} U_{e_i} = T_a U_{e_i}$ and $\partial_a^{e_i} U_{e_j} = 0$ for $i \neq j$. Then for the loop we have

$$\partial_a^{e_i} \mathcal{W} = \text{tr}(U_{e_1} \dots T_a U_{e_i} \dots U_{e_n}), \quad (\text{B3})$$

with a sum over all occurrences if U_{e_i} appears in more than one place. For $U_{e_i}^{-1} = U_{e_i}^\dagger$, using the fact that the generators are skew symmetric, we have $\partial_a^{e_i} U_{e_i}^\dagger = (T_a U_{e_i})^\dagger = -U_{e_i}^\dagger T_a$. Once these directional derivatives are computed, they can be propagated through the derivative of the neural network as described above. For the divergence, one has to trace over e and a . The primary challenge in the implementation is to efficiently keep track of the combinatorics of where edge elements appear and to avoid redundant computations for the second derivatives arising from the divergence calculation.

APPENDIX C: COMPUTATIONAL COST OF THE DIVERGENCE

To understand the role of architecture in relation to the computational cost of the divergence, first consider a simpler scalar vector field

$$\dot{\phi}_x = f_x(\phi, t). \quad (\text{C1})$$

Computing the divergence

$$\frac{\partial}{\partial \phi_x} f_x(\phi, t) \quad (\text{C2})$$

for a generic f would require one automatic differentiation call per lattice site value x , regardless of whether forward or backward mode is used. Let N be the number of such sites [i.e. the size of the vectors ϕ and $f(\phi)$]. As f is a function of N variables, its computational cost will scale as $\mathcal{O}(N)$ in the best case. A single backward or forward pass would thus also scale linearly in N . Since the full divergence computation requires N backward or forward passes, this leads to a total computational cost for the divergence of $\mathcal{O}(N^2)$.

If the analytic form of f is known and chosen appropriately, it may be used to reduce this computational cost. Before considering our gauge field architecture, we can illustrate this in the following toy example. Let

$f_x(\phi) = \sum_{y,i} M_{xy}^i \sin(\omega_i \phi_y)$. It is clear the divergence is in fact cheaper to evaluate than the vector field itself: $\partial f_x / \partial \phi_x = \sum_i M_{xx}^i \omega_i \cos(\omega_i \phi_x)$. This simplification arises from the linear map M and the “locally” applied function $\sin$. However, a black-box application of automatic differentiation for the divergence would not be able to exploit this fact, leading to a large amount of redundant computations.

Now, recall the general architecture of Eq. (19),

$$Z_e^a(t, U) = \sum_{k, \bar{x}} (\partial_a^e \mathcal{W}^k(\bar{x})) \Lambda_{\bar{x}}^k(t, \mathcal{W}). \quad (\text{C3})$$

Applying the chain rule, after differentiating with respect to the same link and Lie algebra component (i.e. applying ∂_a^e), the divergence becomes,

$$\begin{aligned} \partial_a^e Z_e^a(t, U) &= \sum_{k, \bar{x}} (\partial_a^e \partial_a^e \mathcal{W}^k(\bar{x})) \Lambda_{\bar{x}}^k(t, \mathcal{W}) \\ &\quad + \sum_{k, \bar{x}} \sum_{j, \bar{y}} (\partial_a^e \mathcal{W}^k(\bar{x})) (\partial_a^e \mathcal{W}^j(\bar{y})) \\ &\quad \times \frac{\partial}{\partial \mathcal{W}^j(\bar{y})} \Lambda_{\bar{x}}^k(t, \mathcal{W}). \end{aligned} \quad (\text{C4})$$

The cost of taking the second derivative of the Wilson loops for a given edge e does not increase with the lattice size, at least when the choice of included loops is held constant. We can thus focus on the second term. Since not all Wilson loops contain the link U_e , the sums over $\bar{x}$ and $\bar{y}$ do not scale with $|E|$. Thus, the number of derivatives $\partial \Lambda / \partial \mathcal{W}$ that need to be taken for any given edge remains $\mathcal{O}(1)$. However, if the network Λ is a black box, these derivatives still have to be computed with automatic-differentiation for each lattice site. Since Λ is generically a function of all field variables, this may still scale as $\mathcal{O}(|E|^2)$ —a function of $|E|$ variables evaluated $|E|$ times.

The final simplification arises from the shallow form

$$\Lambda_{\bar{x}}^k = \hat{\Lambda}^k(t) + \sum_{l, \bar{z}} C_{\bar{x}\bar{z}}^{kl}(t) g_{\bar{z}}^l(t, \mathcal{W}(\bar{z})). \quad (\text{C5})$$

Applying the derivative,

$$\frac{\partial}{\partial \mathcal{W}^j(\bar{y})} \Lambda_{\bar{x}}^k = \sum_l C_{\bar{x}\bar{y}}^{kl}(t) \frac{\partial}{\partial \mathcal{W}^j(\bar{y})} g_{\bar{y}}^l(t, \mathcal{W}(\bar{y})). \quad (\text{C6})$$

Because the nonlinearity g is applied “locally,” i.e. its output at $\bar{z}$ only depends on the Wilson loop stack at $\bar{z}$ (which does not scale with $|E|$), the sum over $\bar{z}$ above is removed and the convolution reduces to a simple pointwise multiplication. Thus, computing g and its derivatives with respect to $\mathcal{W}$ does not itself scale with $|E|$. We still have to compute this once for each link, of course, leading to a computational cost that grows linearly with $|E|$. In contrast to a black-box vector field, however, this shows that

computing the divergence for our shallow vector field architecture does not add an additional factor of $|E|$ relative to the vector field itself.

APPENDIX D: IMPLEMENTATION OF THE CROUCH-GROSSMAN RUNGE KUTTA INTEGRATION METHOD

We consider ODEs on $G \times \mathbb{R}^d$, where we simultaneously evolve both real degrees of freedom $x \in \mathbb{R}^d$ and matrix Lie group degrees of freedom $Y \in G$. To achieve this, we apply a Runge-Kutta method for the former and simultaneously a Crouch-Grossmann method [26,27] for the latter ensuring that the group structure of Y is preserved during the simultaneous time evolution of both variables.

In the Crouch-Grossman method, the Lie group element evolves according to,

$$Y_{n+1} = \exp(hb_s Z_s^{(n)}) \cdot \dots \cdot \exp(hb_1 Z_1^{(n)}) Y_n,$$

where h is the time step size, b_i are the coefficients from the Butcher tableau (Runge-Kutta coefficients) and $Z_i^{(n)} = Z(x_i^{(n)}, Y_i^{(n)}) \in \mathfrak{g}$ is the Lie algebra-valued vector field at the internal stage i that generally may depend on both $x_i^{(n)}$ and $Y_i^{(n)}$.

The internal stages $Y_i^{(n)} \in G$ are computed using,

$$Y_i^{(n)} = \exp(ha_{i,i-1} Z_{i-1}^{(n)}) \cdot \dots \cdot \exp(ha_{i,1} Z_1^{(n)}) Y_n,$$

where $a_{i,j}$ are given by the chosen Butcher tableau. While we have suppressed the explicit time-dependence in the equations above, we note that in evaluating $Z_i^{(n)}$, we set the time to be $t_i^{(n)} = t_n + c_i h$ with $c_i = \sum_j a_{i,j}$.

For the real degrees of freedom $x \in \mathbb{R}^d$, we use a standard Runge-Kutta method with the same time step h and the same Butcher-tableau coefficients $a_{i,j}, b_i$ as for the Lie group variables. The real degrees of freedom are updated as,

$$x_{n+1} = x_n + h \sum_{i=1}^s b_i f(x_i^{(n)}, Y_i^{(n)}),$$

where $f(x_i^{(n)}, Y_i^{(n)}) \in \mathbb{R}^d$ is the vector field in the ODE $\dot{x} = f(x, Y)$, governing the evolution of x and x_i are the internal stages computed similarly,

$$x_i^{(n)} = x_n + h \sum_{j=1}^{i-1} a_{i,j} f(x_j^{(n)}, Y_j^{(n)}).$$

Note that both internal stages $x_i^{(n)}$ and $Y_i^{(n)}$ are updated in parallel using their respective methods.

APPENDIX E: HAAR MEASURE AND SAMPLING

1. Haar measure for matrix Lie groups

The Haar measure defines a unique invariant measure on compact Lie groups G such as $SU(N)$ [38]. The measure, denoted dU , is invariant under left and right multiplications,

$$d(VU) = d(UV) = dU \quad \forall V, \quad U \in G.$$

This invariance ensures that the measure is uniform across the group. With respect to this measure, the Haar prior distribution is uniform and (up to volume normalization) given as $q_0(U) = 1$.

2. Sampling from the Haar distribution

Sampling from the Haar distribution on $SU(N)$ can be achieved through the construction of random unitary matrices. The method involves generating random matrices with independent complex entries, followed by an orthogonalization procedure that ensures the sampled matrix lies within the $SU(N)$ group [39].

We start with a random complex matrix $X \in \mathbb{C}^{N \times N}$, whose elements are drawn from a standard normal distribution. We then apply a QR decomposition,

$$X = QR,$$

where $Q \in \mathbb{C}^{N \times N}$ is a unitary matrix, and $R \in \mathbb{C}^{N \times N}$ is upper triangular. We then extract and normalize the diagonal part of R by dividing each element by its absolute value: $R' = \text{diag}(R)/|\text{diag}(R)|$, i.e. $R'_{ii} = R_{ii}/|R_{ii}|$, which ensures that all elements are pure phases thus making it unitary. We then need to divide by

$$\kappa = \sqrt[N]{\det(Q) \det(R')}, \quad (\text{E1})$$

to ensure that the matrix $X' = QR'/\kappa$ has a determinant of 1, so finally we have $X' \in SU(N)$. The resulting matrix X' is then uniformly distributed according to the Haar measure on $SU(N)$.

3. Haar measure for $SU(3)$

For the specific case of $SU(3)$, the Haar measure can be expressed in terms of angular coordinates that parametrize the eigenvalues of the matrix. The maximal torus of $SU(3)$ corresponds to diagonal matrices with phases θ_1, θ_2 , and θ_3 , satisfying $\sum_{i=1}^3 \theta_i = 0$ and ensuring the unit determinant condition.

The Haar measure on $SU(3)$ in terms of these angular variables is given by [40],

$$dU = d\theta_1 d\theta_2 \prod_{i < j} |e^{i\theta_i} - e^{i\theta_j}|^2. \quad (\text{E2})$$

APPENDIX F: MARKOV CHAIN MONTE CARLO AND OBSERVABLES

To verify the implementation and successful training of normalizing flows, we generated asymptotically unbiased samples using an independent Metropolis-Hastings Markov chain with the trained normalizing flows as independent proposal distributions. Denoting the learned flow distribution as $q(U)$ and the target distribution as $p(U)$, one samples proposals $U' \sim q$ which are accepted, $U^{(i)} = U'$, with probability

$$\min \left(1, \frac{q(U^{(i-1)})p(U')}{q(U')p(U^{(i-1)})} \right) \quad (\text{F1})$$

and otherwise rejected, leading to a repeat of the previous sample $U^{(i)} = U^{(i-1)}$.

The acceptance rates over a sequence of this Markov chain gives an alternative quality of fit measure for the generative models to the ESS used in the main text. Table IV gives an overview of these values, which are generally in agreement with the ESS.

Figures 4 and 5 show the ratios of observables estimated using the same Markov chains compared to analytically known values [41,42] for SU(2) and SU(3), respectively. Here, $W_{l_1 l_2}$ denote the expectation values of the real parts of the traces of rectangular Wilson loops with edge lengths l_1 and l_2 . In addition, $|\ell|^2$ is the zero-separation two-point function of Polyakov loops

$$\ell(x_2) = \text{tr} \prod_t U_1(x = (t, x_2)), \quad (\text{F2})$$

identifying axis 1 of the lattice with Euclidean time. Then $|\ell|^2 = l^*(x)l(x)$ is independent of x due to translational symmetry.

TABLE IV. Acceptance rates (in %) for the continuous normalizing flows trained for different gauge groups and lattice sizes using an independent Metropolis-Hastings Markov chain of length 20480 for all configurations of Tables II and III.

SU(2), $L = 16$		SU(3), $L = 16$			SU(3), $L = 8$	
$\beta = 2.2$	$\beta = 2.7$	$\beta = 5$	$\beta = 6$	$\beta = 8$	$\beta = 8$	$\beta = 12$
84	82	78	71	38	65	32




FIG. 4. Ratio of Monte Carlo estimates of observables to analytic values for flows on SU(2) on a 16×16 lattice. The trained flows of Table II are used as proposal distribution for an independent Metropolis-Hastings Markov chain of length 20480, and errors shown as vertical bars are estimated using a bootstrap method.




FIG. 5. Ratio of Monte Carlo estimates of observables to analytic values for flows on SU(3) on a 16×16 lattice. The trained flows of Table II are used as proposal distribution for an independent Metropolis-Hastings Markov chain of length 20480, and errors shown as vertical bars are estimated using a bootstrap method.

- [1] C. Morningstar, The Monte Carlo method in quantum field theory, [arXiv:hep-lat/0702020](#).
- [2] A. Sokal, Monte Carlo methods in statistical mechanics: Foundations and new algorithms, in *Functional Integration: Basics and Applications* (Springer, New York, 1997), pp. 131–192.
- [3] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [4] M. S. Albergo, G. Kanwar, and P. E. Shanahan, Flow-based generative models for Markov chain Monte Carlo in lattice field theory, *Phys. Rev. D* **100**, 034515 (2019).
- [5] M. S. Albergo, D. Boyda, D. C. Hackett, G. Kanwar, K. Cranmer, S. Racanière, D. J. Rezende, and P. E. Shanahan, Introduction to normalizing flows for lattice field theory, [arXiv:2101.08176](#).
- [6] M. C. N. Cheng and N. Stratikopoulou, Lecture notes on normalizing flows for lattice quantum field theories, [arXiv:2504.18126](#).
- [7] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, Neural ordinary differential equations, [arXiv:1806.07366](#).
- [8] D. Boyda, G. Kanwar, S. Racanière, D. J. Rezende, M. S. Albergo, K. Cranmer, D. C. Hackett, and P. E. Shanahan, Sampling using $\$SU(N)\$$ gauge equivariant flows, *Phys. Rev. D* **103**, 074504 (2021).
- [9] P. de Haan, C. Rainone, M. C. N. Cheng, and R. Bondesan, Scaling up machine learning for quantum field theory with equivariant continuous flows, [arXiv:2110.02673](#).
- [10] M. Gerdes, P. de Haan, C. Rainone, R. Bondesan, and M. C. N. Cheng, Learning lattice quantum field theories with equivariant continuous flows, *SciPost Phys.* **15**, 238 (2023).
- [11] M. Caselle, E. Cellini, and A. Nada, Sampling the lattice nambu-goto string using continuous normalizing flows, *J. High Energy Phys.* **02** (2024) 048.
- [12] Y. Nagai and A. Tomiya, Gauge covariant neural network for quarks and gluons, *Phys. Rev. D* **111**, 074501 (2025).
- [13] R. Abbott *et al.*, Normalizing flows for lattice gauge theory in arbitrary space-time dimension, [arXiv:2305.02402](#).
- [14] R. Abbott, A. Botev, D. Boyda, D. C. Hackett, G. Kanwar, S. Racanière, D. J. Rezende, F. Romero-López, P. E. Shanahan, and J. M. Urban, Applications of flow models to the generation of correlated lattice QCD ensembles, *Phys. Rev. D* **109**, 094514 (2024).
- [15] M. Lüscher, Trivializing maps, the Wilson flow and the HMC algorithm, *Commun. Math. Phys.* **293**, 899 (2010).
- [16] S. Bacchio, P. Kessel, S. Schaefer, and L. Vaitl, Learning trivializing gradient flows for lattice gauge theories, *Phys. Rev. D* **107**, L051504 (2023).
- [17] M. Gerdes, bijx: Bijections & normalizing flows with JAX/NNX, <https://github.com/mathisgerdes/bijx>.
- [18] M. Gerdes, P. de Haan, R. Bondesan, and M. C. N. Cheng, Non-perturbative trivializing flows for lattice gauge theories, [10.5281/zenodo.15576712](#) (2025).
- [19] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, Estimation of thermodynamic observables in lattice field theories with deep generative models, *Phys. Rev. Lett.* **126**, 032001 (2021).
- [20] D. Albandea, L. Del Debbio, P. Hernández, R. Kenway, J. Marsh Rossney, and A. Ramos, Learning trivializing flows, *Eur. Phys. J. C* **83**, 676 (2023).
- [21] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, Neural Ordinary Differential Equations, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2018), Vol. 31.
- [22] L. Falorsi, Continuous normalizing flows on manifolds, [arXiv:2104.14959](#).
- [23] A. Lou, D. Lim, I. Katsman, L. Huang, Q. Jiang, S.-N. Lim, and C. De Sa, Neural manifold ordinary differential equations, [arXiv:2006.10254](#).
- [24] L. Falorsi and P. Forré, Neural ordinary differential equations on manifolds, [arXiv:2006.06663](#).
- [25] M. Nakahara, *Geometry, Topology and Physics*, 2nd ed. (CRC Press, Bristol, 2003).
- [26] P. E. Crouch and R. Grossman, Numerical integration of ordinary differential equations on manifolds, *J. Nonlinear Sci.* **3**, 1 (1993).
- [27] M. Wandelt, Geometric integration on Lie groups and its applications in Lattice QCD, Ph.D. thesis, Universitätsbibliothek Wuppertal, 2021.
- [28] R. Abbott, M. S. Albergo, D. Boyda, D. C. Hackett, G. Kanwar, F. Romero-López, P. E. Shanahan, and J. M. Urban, Multiscale normalizing flows for gauge theories, *Proc. Sci., LATTICE2023* (2024) 035 [[arXiv:2404.10819](#)].
- [29] H. Wu, J. Köhler, and F. Noé, Stochastic Normalizing flows, [arXiv:2002.06707](#).
- [30] L. I. Midgley, V. Stimper, G. N. C. Simm, B. Schölkopf, and J. M. Hernández-Lobato, Flow annealed importance sampling bootstrap, [arXiv:2208.01893](#).
- [31] M. Caselle, E. Cellini, A. Nada, and M. Panero, Stochastic normalizing flows as non-equilibrium transformations, *J. High Energy Phys.* **07** (2022) 015.
- [32] A. Bulgarelli, E. Cellini, and A. Nada, Sampling SU(3) pure gauge theory with stochastic normalizing flows, *Proc. Sci., LATTICE2024* (2025) 040 [[arXiv:2409.18861](#)].
- [33] A. Bulgarelli, E. Cellini, and A. Nada, Scaling of stochastic normalizing flows in SU(3) lattice gauge theory, *Phys. Rev. D* **111**, 074517 (2025).
- [34] R. Abbott *et al.*, Aspects of scaling and scalability for flow-based sampling of lattice QCD, *Eur. Phys. J. A* **59**, 257 (2023).
- [35] L. Del Debbio, J. M. Rossney, and M. Wilson, Efficient modelling of trivializing maps for lattice ϕ theory using normalizing flows: A first look at scalability, *Phys. Rev. D* **104**, 094507 (2021).
- [36] M. Favoni, A. Ipp, D. I. Müller, and D. Schuh, Lattice gauge equivariant convolutional neural networks, *Phys. Rev. Lett.* **128**, 032003 (2022).

- [37] J. Bradbury *et al.*, JAX: Composable transformations of Python + NumPy programs (2018).
- [38] J. J. Duistermaat and J. A. C. Kolk, *Lie Groups*, Universitext (Springer, Berlin, Heidelberg, 2000).
- [39] F. Mezzadri, How to generate random matrices from the classical compact groups, *Not. Am. Math. Soc.* **54**, 592 (2007).
- [40] D. Bump, The Weyl Integration Formula, in *Lie Groups*, Graduate Texts in Mathematics, edited by D. Bump (Springer, New York, 2004), pp. 112–116.
- [41] A. A. Migdal, Recursion equations in gauge theories, *Sov. Phys. JETP* **42**, 413 (1975).
- [42] D. J. Gross and W. Taylor, Two-dimensional QCD is a string theory, *Nucl. Phys.* **B400**, 181 (1993).