

Flow-based sampling for lattice field theories

Gurtej Kanwar

*Albert Einstein Center, Institute for Theoretical Physics,
University of Bern, 3012 Bern, Switzerland*

E-mail: kanwar@itp.unibe.ch

Critical slowing down and topological freezing severely hinder Monte Carlo sampling of lattice field theories as the continuum limit is approached. Recently, significant progress has been made in applying a class of generative machine learning models, known as “flow-based” samplers, to combat these issues. These generative samplers also enable promising practical improvements in Monte Carlo sampling, such as fully parallelized configuration generation. These proceedings review the progress towards this goal and future prospects of the method.

*The 40th International Symposium on Lattice Field Theory (Lattice 2023)
July 31st - August 4th, 2023
Fermi National Accelerator Laboratory*

Contents

1	Introduction	2
2	Flow-based sampling	3
2.1	Normalizing flows	3
2.2	Flows using machine learning	3
2.3	Application to scalar field theory	5
2.4	Symmetries and application to gauge theory	6
2.5	Fermions and application to QCD	8
3	Current research	8
3.1	New paradigms	8
3.2	Practical gains	9
3.3	Towards QCD at scale	10
3.4	Frontiers of development	11
4	Outlook	14

1. Introduction

Lattice simulations play a vital role in the study of quantum field theories, allowing properties of a theory to be non-perturbatively computed from first principles. For example, lattice field theory calculations are key to understanding the Quantum Chromodynamics (QCD) sector of the Standard Model. Lattice field theory calculations require evaluating the lattice-regularized path integral,

$$\langle O \rangle \equiv \frac{1}{Z} \int \mathcal{D}[U] \mathcal{O}[U] e^{-S[U]}, \quad Z \equiv \int \mathcal{D}[U] e^{-S[U]}, \quad (1)$$

where \mathcal{O} is an operator of interest, S is the discretized action of the theory, $\int \mathcal{D}[U]$ indicates integration over all lattice field configurations, and Z is the partition function. Almost all calculations achieve this through Monte Carlo sampling according to the distribution

$$p[U] = e^{-S[U]} / Z \quad (2)$$

which allows unbiased estimates of $\langle O \rangle$ with systematically improvable uncertainties. Traditional approaches based on Markov Chain Monte Carlo have achieved incredible progress through the use of high-performance computing up to the exascale, but the precision of many results is limited by Monte Carlo statistics. Increasing these statistics is particularly challenging at small lattice spacings where critical slowing down causes Markov Chain autocorrelations to diverge as the continuum limit is taken [1]. Furthermore, for state-of-the-art calculations, practical challenges

associated with storage costs and efficient parallelization are encountered as the number of lattice sites increases [2, 3].

Recently, an approach to Monte Carlo sampling using *flow-based models* has been proposed and evaluated in several lattice field theory contexts, with promising early results. Such models are based on *normalizing flows* [4, 5], which in this context describe efficient transformations between distributions over lattice field configurations, often using machine learning techniques to parameterize specific components of transformations. Models can be constructed to use normalizing flows for a variety of tasks:

- Perform efficient Monte Carlo sampling
- Estimate the partition function Z
- Draw correlated samples at different bare physical parameters
- Efficiently perform physical parameter scans

These proceedings discuss the progress in this developing field and its future prospects.

2. Flow-based sampling

2.1 Normalizing flows

To get a feel for normalizing flows, we can consider a simple example in two variables. Starting from points (x, y) distributed uniformly in the unit disk, $D = \{x, y \mid x^2 + y^2 < 1\}$, consider applying the Box-Muller transformation $(x', y') = f(x, y)$ defined in terms of $r \equiv \sqrt{x^2 + y^2}$ as [6]

$$x' = \frac{x}{r} \sqrt{-2 \ln r^2} \quad \text{and} \quad y' = \frac{y}{r} \sqrt{-2 \ln r^2}. \quad (3)$$

This is an invertible map from the unit disk to the whole 2D plane. The action of the transformation stretches the points near the edge of the disk towards infinity while concentrating the points near the center, as shown in Fig. 1. The output probability density $q(x', y')$ is given in terms of the *prior density* $r(x, y) = \text{Unif}_D(x, y)$ and the Jacobian J_f of the transformation as

$$q(x', y') = r(x, y) |\det J_f|^{-1} = \frac{1}{2\pi} e^{-(x'^2 + y'^2)/2}. \quad (4)$$

This is simply the unit normal distribution over two variables. The Box-Muller transformation can thus be used to draw Gaussian-distributed variables (which are useful in many contexts) by producing samples of uniform random variables on the disk (which are easy to generate) and applying f .

2.2 Flows using machine learning

Generalizing this concept, if one can find a diffeomorphism f with a tractable Jacobian between a simple distribution and a more complicated distribution of interest, then this approach can be used to efficiently sample the latter distribution. In most contexts, however, either such a transformation is not known explicitly or it is expensive to apply the function or its Jacobian. In the lattice field theory context, the trivializing map [7] is one example of a flow from a gauge theory at infinite




Figure 1: Example of the Box-Muller transformation, which can be considered a simple normalizing flow.

lattice spacing to one at finite lattice spacing; however, the map and its Jacobian are defined by integrating ordinary differential equations (ODEs), which was demonstrated to be too costly to provide a significant advantage in practice [8].

An alternate approach is to *approximate* the target distribution p by optimizing the choice of flow function f within a family of functions to produce an output distribution q that is close to the target distribution. Working with a variational family that is explicitly defined to ensure the function and its Jacobian determinant can be efficiently evaluated restricts to computationally tractable choices [9]. Families of functions can be defined by a variety of methods, but approaches based on machine learning are appealing: neural networks can efficiently parameterize large classes of functions [10], existing frameworks can be used to efficiently evaluate such functions, and stochastic gradient descent can be straightforwardly applied to “train” (optimize) the choice of function. For a comprehensive review of progress in parameterizing normalizing flows using machine learning methods, see Ref. [11].

To simplify the task of ensuring the Jacobian determinant remains tractable while defining an expressive family of flows, one of two broad approaches is commonly used:

Discrete learnable flows: The function f is defined as the composition of a discrete sequence of functions g_1, \dots, g_n , where the Jacobian determinant of f can be efficiently computed by

$$\det J_f = \det J_{g_1} \cdots \det J_{g_n}. \quad (5)$$

One such approach is to define the g_i as “coupling layers”, which transform one subset of the variables conditioned on the complementary subset, resulting in a simple triangular Jacobian [12, 13]. Each g_i is parametrized by free parameters which can collectively be optimized to reproduce a target distribution.

Continuous learnable flows: The function f is defined by integrating an ODE [14–16]. To ensure f is a diffeomorphism, the ODE can be defined in terms of the gradient of a scalar function φ as

$$\frac{d}{dt} U(t) = \nabla \varphi(U(t); t), \quad (6)$$

giving $U(t) = V + \int_0^t dt' \nabla \varphi(U(t'); t')|_{U(0)=V}$ and $f[V] = U(T)$ in terms of the total flow time T . For gauge fields, the differentiation, integration, and addition must be appropriately defined on

the manifold of gauge configurations [7]. The Jacobian determinant can also be computed with an analogous ODE,

$$\ln \det J = - \int_0^T dt \nabla^2 \varphi(U(t); t). \quad (7)$$

Here, the scalar function φ is parameterized by free parameters which can be optimized to reproduce a target distribution. The trivializing map of Ref. [7] is a particular parameter-free choice of continuous flow.

Optimizing the parameters of a flow defined by one of the methods above requires a target “loss function” to be minimized, typically via stochastic gradient descent. A useful loss function for the lattice field theory context should provide a measure of the distance between the model distribution $q[U]$ and the known target distribution $p[U] = \exp(-S[U])/Z$ without requiring many (or any) existing samples of the distribution. These features set apart the challenge of optimizing flow-based models for lattice field theory from many other machine learning contexts. Taking advantage of the explicitly known action $S[U]$, the reverse Kullback-Leibler divergence [17] can be used to *self-train* a model [15], i.e., to optimize the flow function f by performing stochastic gradient descent using only samples of field configurations drawn from the model and evaluations of the action on those configurations. This concept of self-training is similar in spirit to methods applied in non-flow approaches, such as for Self-Learning Monte Carlo methods [18–29].

2.3 Application to scalar field theory

In Ref. [30], these concepts were applied to demonstrate for the first time that critical slowing down could be essentially eliminated using flow-based models. The focus of this early study was an interacting scalar field theory on a 1 + 1D lattice, described by the action

$$S[\phi] = \sum_x \partial_\mu \phi(x) \partial^\mu \phi(x) + \frac{M^2}{2} \phi(x)^2 + \lambda \phi(x)^4. \quad (8)$$

In this work, flow-based models used discrete learnable flows with a coupling layer architecture that alternately transformed even and odd sites conditioned on the complementary subset [12, 13]. Each such coupling layer locally applies the affine transformation $\phi'(x) = e^{s(x)} \phi(x) + t(x)$ with parameters $s(x)$ and $t(x)$ determined from neural network outputs, as shown on the left side of Fig. 2. This transformation has a triangular Jacobian with simple determinant,

$$J_{ij} \equiv \partial \phi'_i / \partial \phi_j = \begin{bmatrix} I & \\ \blacksquare & \delta_{ij} e^{s_i} \end{bmatrix} \implies \ln \det J = \sum_i s_i. \quad (9)$$

The right side of Fig. 2 schematically depicts the composed sequence of coupling layers which, when optimized, transform configurations sampled from a simple Gaussian prior distribution to configurations distributed approximately according to $p[\phi] = \exp(-S[\phi])/Z$, in terms of the scalar field theory action $S[\phi]$ in Eq. (8).

In the lattice field theory context, maintaining quantifiable uncertainties is also an important requirement. While a flow that only approximates a given target distribution produces biased samples, the use of a flow-based model with computable model density $q[\phi]$ allows any biases to be systematically corrected. For example, samples from the model can be used as proposals in a




Figure 2: Discrete learnable flow architecture for scalar field theory demonstration in Ref. [30].

Markov Chain, with usual Metropolis acceptance steps applied to avoid bias [30], or observables can be directly reweighted using the weights $p[\phi]/q[\phi]$ [31].

By now, many works have demonstrated flow-based sampling for scalar field theory with larger lattice sizes, smaller lattice spacings, and in the broken symmetry phase [32–40]. Since this early work, there have also been many developments in the direction of incorporating gauge symmetries, incorporating dynamical fermions, and scaling to address challenging problems in lattice QCD, as we now discuss. A further overview of these developments can be found in Ref. [41].

2.4 Symmetries and application to gauge theory

Symmetries of the lattice action play an important role in recovering the correct continuum limit, simplifying renormalization, and simplifying analysis. By applying Metropolis acceptance steps or reweighting, one is guaranteed to recover the correct physical distribution, including all symmetries. Nevertheless, incorporating symmetries directly in the definition of variational families of flow functions may improve both the training time and the final quality of the resulting flow-based model. Intuitively, all parameters of such a family describe the distribution in the meaningful physical directions, while symmetric directions remain “flat” by construction; this ensures that gradients provide good information about all parameters and may reduce the total number of model parameters required to construct models of similar quality.

The distribution of a flow-based model is defined by both the prior density $r[U]$ and the flow f . A sufficient condition to produce an invariant model is that, for all generators t of the symmetry group, [42]

1. The prior density is invariant, $r[t \cdot U] = r[U]$, and
2. The flow is equivariant, $f[t \cdot U] = t \cdot f[U]$.

The requirement on the prior distribution is often straightforward to satisfy for lattice field theory symmetries. Building flows that are equivariant under such symmetries requires more care.

In practice, it is important to exactly incorporate gauge symmetry into flows to make progress in applying flow-based models to lattice gauge theories, such as QCD. A procedure to build gauge-equivariant coupling layers was introduced in Refs. [44, 45], enabling computationally efficient discrete flows to be constructed for gauge theories. As shown schematically in Fig. 3, such coupling layers involve transforming gauge links $U_\mu(x)$ in terms of equivariant transformations of untraced




Figure 3: Gauge equivariant transformation defined in terms of inner transformations $W_\ell(x) \rightarrow W'_\ell(x)$ of untraced Wilson loops. Using the generalized staple $V_\ell(x)$, the transformation of each Wilson loop can be assigned to particular gauge links $U_\mu(x)$. Figure reprinted from Ref. [43].




Figure 4: Measured autocorrelation time of topological charge (τ_Q^{int} , left) and estimates of observables (Wilson loops $\langle W_{\ell \times \ell} \rangle$ and topological susceptibility χ_Q , right) in a 1+1D $U(1)$ gauge theory [44]. Sampling with flow-based models (Flow) results in significantly reduced topological freezing in the continuum limit ($\beta \rightarrow \infty$) and more precise estimates versus Hybrid Monte Carlo (HMC) and Heatbath (HB) using identical ensemble sizes. Figures adapted from Refs. [43, 44].

Wilson loops $W_\ell(x)$, which satisfy the simpler same-site transformation $W_\ell(x) \rightarrow \Omega(x)W_\ell(x)\Omega^\dagger(x)$ under gauge symmetry. In this case, custom neural networks suitable for acting on the compact $U(1)$ and $SU(N)$ manifolds also needed to be defined [45, 46], highlighting the broader consideration that in many cases specialized machine learning developments need to be made to fit the lattice field theory context.

Figure 4 depicts an example of the significant improvements possible once such flows are constructed and optimized for lattice gauge theories, in this case for a 1+1D $U(1)$ gauge theory [44]. Here, topological freezing in the continuum limit ($\beta \rightarrow \infty$) is significantly mitigated when sampling using flows. The result is far more precise observable estimates, especially for topological quantities such as the topological susceptibility, given the same ensemble sizes. In this example, the cost per configuration in the ensemble was roughly equivalent in all three methods.

The gauge-equivariant transformations described above are by no means the only approach to defining equivariant transformations. Recent works have developed and applied several new architectures for gauge-equivariant transformations that are suitable whenever invertibility is not required [28, 47–53]. It is also possible to incorporate gauge equivariance into continuous flows by




Figure 5: Free energy estimates in a scalar ϕ^4 theory using a flow-based estimator (flow) vs a naive estimator using Hybrid Monte Carlo data (hmc). Figure reprinted from Ref. [34].

working with *gauge-invariant* scalar potentials φ [54]. In the view of the author, developing architectures for gauge-equivariant flow transformations that are both efficient and expressive remains one of the areas of this research requiring the most further investigation.

2.5 Fermions and application to QCD

It is noteworthy that the methods described so far are already sufficient to apply normalizing flows to theories with fermions and theories in arbitrary spacetime dimensions, including the particularly challenging task of lattice QCD, as demonstrated by exploratory studies performed in Refs. [54–59]. These studies have shown that there are no theoretical obstacles to flow-based sampling for almost all lattice field theories of interest. However, how to best design and optimize flow-based models in these more general contexts remains an area of active study, as we now discuss.

3. Current research

3.1 New paradigms

Though motivated by the problem of critical slowing down, using flows to transform between lattice field theory distributions opens the door to several potential applications beyond directly sampling field configurations. These new paradigms are outlined below.

Partition functions. The reverse Kullback-Leibler divergence used to optimize flows provides a lower bound on the partition function, Z , or conversely an upper bound on the free energy, $-\ln Z$ [15]. Building on this observation, an unbiased, low-variance estimator for the partition function was recently introduced [31, 34]. Figure 5 depicts the improvements in precision that was for example achieved in Ref. [34] using this method to estimate the free energy of a scalar field theory. No modifications to the flow-based model architecture or training are required to apply this method, meaning that much more precise thermodynamic observables automatically become available if flow-based Monte Carlo sampling is used.

Parameter dependence. Flows can be defined with additional dependence on the bare physical parameters of the lattice action. In this case, if the parameterization of the flow is sufficiently flexible,

one could hope to simultaneously define flows across an entire range of action parameters [37, 40, 60]. Just as ordinary machine learning models have demonstrated the ability to effectively interpolate and extrapolate from a given data set, a similar training strategy could be applied for flows. Optimizing a parameterized family of flows on several representative choices of bare couplings, one aims to use the resulting model to efficiently interpolate and extrapolate to sampling a nearby values of the couplings. This for example could save significant computational effort when tuning bare parameters before ensemble generation.

Correlated samples. In addition to considering a family of flows from a common prior distribution to theories defined by many choices of physical parameters, as above, one can also construct flows that map *between* theories at nearby points in parameter space [61]. This approach has the potential to produce low-noise estimates of observables that require correlated samples from a pair of actions, for example in the Feynman-Hellmann method for evaluating matrix elements. Whereas flow-based direct sampling requires optimizing a flow between a simple prior distribution and more challenging target distribution, which typically are very different, this approach to correlated samples only requires constructing flows between very similar distributions.

Parallel tempering. Parallel tempering involves the execution of several Markov chains in parallel, each sampling an action interpolating between a distribution that is easier to sample and the target distribution [62]. The efficiency of the method is controlled by the exchange rate between neighboring chains. By using flows to perform an exchange, a high acceptance rate could be achieved, with a perfect flow corresponding to a 100% acceptance rate [63, 64]. This approach is related to out-of-equilibrium simulations, which can similarly benefit from using normalizing flows as intermediate steps [35, 65], and to the general framework of stochastic normalizing flows [66].

Sign problems. By defining flows to map into the complexified domain of field configurations, similar techniques may also be used to define and optimize deformations of the lattice path integral, which continue to be explored as potential solutions to sign problems [67–71].

3.2 Practical gains

Using flow-based sampling in state-of-the-art lattice calculations would not only address difficult autocorrelations arising from critical slowing down, but would also practically simplify calculations in several ways, as outlined below.

Parallel sampling. In a Monte Carlo scheme where samples are drawn directly, many such samples can be produced in parallel. This structure maps very naturally onto the large multi-node machines presently used for state-of-the-art lattice QCD calculations, and results in almost perfect weak scaling. This benefit carries over even to the context of flow-based Markov Chain Monte Carlo, where normalizing flows are used to propose steps that are accepted or rejected using a Metropolis step. Despite the apparent serialization, the proposals are independent of the state of the chain, meaning they can be generated in parallel and arbitrarily in advance.

Storage-free ensembles. A surprising benefit of direct flow-based sampling is the possibility of eliminating long-term storage (and similarly communication) costs of ensembles of field configurations. With flow-based sampling, each configuration is produced independently starting from an

easily-constructed sample of a prior distribution. In general, the sample from the prior distribution can be replaced with just the *state of the random number generator* that was used to produce it, which is a much smaller object. For example, if each sample from the prior distribution is produced using a known seed, it would be sufficient to store or communicate the seed alone to be able to deterministically reconstruct the final field configuration on demand. This is a classic tradeoff between data size versus additional computation which is exploited in traditional compression algorithms. In this sense, a good normalizing flow can be considered an excellent compression scheme for lattice field configurations!

3.3 Towards QCD at scale

Among many interesting applications of lattice simulations, state-of-the-art lattice QCD represents by far the biggest computational challenge. At the same time, the field stands to make significant progress if a breakthrough in Monte Carlo sampling for lattice QCD could be achieved. As such, many developments in flow-based sampling have been oriented towards the goal of lattice QCD. By now, significant progress has been made in this direction.

This progress can be viewed in three broad stages:

- Proof-of-principle studies of solving critical slowing down, incorporating gauge symmetries, and incorporating fermions (see Secs. 2.3, 2.4, and 2.5);
- Works addressing architectural and training challenges presented by dynamical fermions; and
- Applications to higher spacetime dimensions and more degrees of freedom.

The progress in the latter two stages is discussed briefly below. In the view of the author, we are now seeing progress in the final stage that will result in impact to some state-of-the-art calculations in the next couple of years, likely first through one or more of the new paradigms discussed in Sec. 3.1 rather than direct flow-based ensemble generation.

Dynamical fermions. In lattice calculations, dynamical fermions are normally integrated out to produce an effective action over the remaining bosonic degrees of freedom that is written in terms of a determinant of the Dirac operator. In principle, flow-based models can be optimized to reproduce this target distribution just as for an ordinary bosonic theory. Early studies demonstrated that it was possible to train models to closely reproduce this effective action using exact evaluations of the determinant [55, 56]. However, just as for Hybrid Monte Carlo (HMC), this approach does not scale favorably with the number of lattice sites. This challenge can be avoided as in HMC by using the pseudofermion representation of the determinant [55, 58, 72], for example by using a joint model as shown in Fig. 6 to sample bosonic fields and pseudofermions simultaneously. The “marginal” normalizing flow used to sample the effective action is combined with a “conditional” flow that produces pseudofermion samples along with a *normalized* estimate of their probability density, effectively providing an estimator of the determinant. In addition, Ref. [72] showed that the method can be combined with even-odd or Hasenbusch preconditioning as usual, and noise from the pseudofermion estimator of the determinant can be reduced by averaging over multiple samples.




Figure 6: Structure of a normalizing flow model that jointly produces a bosonic field U (for example a gauge field) and a pseudofermion field ϕ . In this case, a marginal model first produces the bosonic field, then a conditional model produces the pseudofermion conditioned on the bosonic field. Figure reprinted from Ref. [72].

Higher dimensions and more degrees of freedom. While the technology described above is sufficient to apply flow-based sampling to almost any lattice field theory of interest, there are practical challenges associated with training models with many degrees of freedom. There is also still much ground that can be gained by experimenting with varying the architectures of flow-based models, as has been demonstrated by recent advances [54, 59, 73]. With such refinements and more dedicated studies, flow-based sampling has now been applied to the $N_f = 2$ Schwinger model with up to 128×128 lattice sites [73], to 3 + 1D Yang-Mills [54, 59], and to $N_f = 2$ QCD in a small volume [58].

3.4 Frontiers of development

To take the remaining steps towards flow-based sampling for lattice field theories at state-of-the-art scale, including towards improved lattice QCD calculations, continued development is needed. The many frontiers of these developments can unfortunately not be covered here in detail. Instead, only the most recent developments in several promising directions are discussed below.

New ways to construct flows. The flexibility of flows composed from a discrete sequence of coupling layers depends on each transformation being sufficiently expressive and their composition enabling further expressivity. As discussed above, the requirement of enforcing gauge equivariance in coupling layers restricts their form and likely their expressivity. Recent work has sought new ways to define gauge equivariant layers that are both individually expressive and compose well [59], enabling applications to theories in higher spacetime dimensions.

Recently, several works have also investigated learnable continuous flows defined by the integration of a continuous transformation [54, 57]. Using a physics-inspired basis for the learnable potential, parameterized by collections of Wilson loops of varying sizes, promising results have been demonstrated for sampling gauge theories using very few free parameters. The challenge, as for discrete flows, is to parameterize flows that can be expressive enough to capture the difficult distributions associated with realistic lattice gauge theories in higher spacetime dimensions and at fine lattice spacings. Additionally, for continuous flows, accurate integration of the ODEs defining the flow and its Jacobian are required. Nevertheless, the simplicity of incorporating symmetries in this approach is attractive—defining families of *invariant* potentials tends to be a less complex task than defining *equivariant* coupling layers.

It is noteworthy that the benefits of both approaches may be combined: Ref. [59] introduced *residual coupling layers*, which define discrete invertible transformations by the gradient of a scalar potential. This architecture is currently being further explored.

Meanwhile, recent works have explored improving normalizing flows by incorporating stages that work in a Fourier momentum basis, rather than the lattice position basis [36, 74]. These efforts have so far demonstrated promising results for quantum mechanics problems and scalar field theory.

Finally, normalizing flows constructed from a hierarchical sequence of transformations is another natural approach to model the hierarchy of scales in typical lattice field theory problems. This can be considered an application of “in-painting” which has been studied thoroughly in the image generation community, including through the use of normalizing flows [12]. This approach has been explored in only a few lattice field theory examples so far [75, 76].

New ways to incorporate flows. The focus of this review has largely been on applying normalizing flows to convert easily produced samples of a simple prior distribution to independent samples of a more complicated distribution. However, flows as flexible invertible transformations of field configurations can potentially be incorporated in lattice field theory calculations in many ways. Some of these ways have already been discussed in the new paradigms covered in Sec. 3.1:

- Normalizing flows may be used to produce a correlated field configuration distributed according to one action from a field configuration drawn from the other.
- Similarly, flows may be used to map a configuration to a more typical configuration under a target action when proposing exchanges between neighboring chains in parallel tempering.
- Flows may be used to define contour deformations to solve sign problems.

Flows can also be incorporated into more standard Markov Chain sampling approaches. On one hand, both the trivializing map and more general flows have been applied to transform to a more easily sampled distribution on which standard HMC methods can be applied for faster mixing [7, 8, 38, 39, 57, 75]. On the other hand, several recent works have explored *flow-based HMC*, in which HMC trajectories are themselves defined by learned flow transformations [77, 78]. In contrast to directly producing samples from a hierarchical normalizing flow, as discussed above, recent works have also considered using domain decomposition to combined flows in inner blocks with traditional Markov Chain steps in an outer procedure [36, 73, 79]. This approach has the benefit of exploiting the expected decorrelation of subvolumes of field configurations that are separated by sufficient physical distance, circumventing the known challenges surrounding directly sampling field configurations at large physical volume [72].

Application to new theories. In recent years, several other lattice theories have been approached using normalizing flows. In particular, flows have been applied to the effective Nambu-Goto action of a string worldsheet [80, 81], the XY model [60, 79], and CP^{N-1} models [65, 82]. In the view of the author, it is particularly interesting to further study whether critical slowing down can be eliminated for CP^{N-1} models, as these models reproduce many of the features of QCD [83, 84] while remaining computationally inexpensive. Parallel tempering methods, potentially combined with normalizing flows, have made progress against topological freezing in CP^{N-1} models [65, 85, 86],




Figure 7: Left: Model quality (effective sample size) as a function of optimization time using the standard gradient estimator versus an improved path gradient estimator for a 1 + 1D scalar field theory. Right: Model quality (effective sample size, ESS) as a function of the number of gradient descent steps using the standard gradient estimator versus an improved gradient estimator using control variates for a 3 + 1D $SU(3)$ gauge theory. Figures adapted from Refs. [59, 91].

and it remains to be seen whether all significant critical slowing down effects can be addressed by such means or other flow-based sampling approaches.

New software. Exploring normalizing flow techniques requires implementing models with components and architectures not currently used in the machine learning community, precluding the use of off-the-shelf implementations. Some early steps have been taken in the direction of making code specific to lattice field theory available in the form of a Jupyter notebook [87] and a Julia package [88], each with implementations of specific flow models, as well as a python implementation of flow-based HMC [77]. Recently, Ref. [89] introduced a fully general package aimed at collecting components and tools across the various applications of flows for lattice field theory. In addition, one of the main practical challenges of studying normalizing flows at scale is the absence of highly optimized codes such as those available for traditional HMC simulations. While machine learning libraries provide the ability to target hardware accelerators and perform relatively optimized operations, reaching state-of-the-art scales will require significant further development, especially towards efficient performance in a distributed computing context.

New training methods. The choice of training (optimization) scheme used to determine a good flow function from a variational family can have significant effects both on the data efficiency of training and the final quality of the resulting flow-based model. Early works largely used simple stochastic estimators of gradients of the reverse Kullback-Leibler divergence for stochastic gradient descent optimization. Recently, however, a number of works have explored lower-variance estimators of the same gradients as a means of improving the convergence, demonstrating in several cases significantly better training times and results [54, 59, 90–93]. For example, Fig. 7 demonstrates the stark improvement in model quality when working with improved estimators to optimize flow-based models for examples of 1 + 1D scalar field theory [91] and 3 + 1D $SU(3)$ gauge theory [59].

For theories with a spontaneously broken symmetry, the target distribution features either multiple isolated modes, corresponding to distinct classical minima of the action, or extended

modes, corresponding to a continuum of minima. These structures can also present difficulties to optimization using the reverse Kullback-Leibler divergence. Recently, several works have studied these challenges, proposing several practical methods to avoid training difficulties [32, 94].

Finally, further exploration has demonstrated that the time for optimization and the final model quality are both sensitive to many hyperparameter choices of the models and optimization procedure [95]. It is clear that in many cases training costs can be improved, sometimes by orders of magnitude. One systematic study was performed of the scaling of training costs when working with a particular model architecture, demonstrating a steep increase in the required training time as the continuum limit was approached [96]. This work highlights the need to continue to investigate the training procedure and model architecture. Whether similar scaling holds for more modern flow-based models at large scales remains an interesting open question, meriting careful future investigation.

4. Outlook

In the past few years, rapid progress has been made towards the goal of implementing effective normalizing flows for lattice field theory configurations. Flow-based models have been shown to essentially eliminate critical slowing down and topological freezing for certain theories. New components and architectures have been developed to extend flow-based models to apply to a wide variety of theories: theories with gauge and fermionic degrees of freedom, theories in arbitrary spacetime dimensions, and theories with non-trivial phase structure. Significant progress has specifically been made in the direction of applying flow-based sampling for the particularly challenging task of large-scale lattice QCD calculations, including several applications to theories with light dynamical fermions, pure-gauge theories in $3 + 1D$, and $N_f = 2$ QCD in a small volume.

There are of course many open questions remaining and exciting new directions of research currently under way. For example, recent works on continuous normalizing flows have demonstrated that, for several gauge theory applications, a particularly parameter-efficient representation of normalizing flows is possible, inspired by a generalization of trivializing maps. Whether and how these results can be combined with existing work on discrete normalizing flows remains to be seen. Meanwhile significant strides have been made in our understanding of how to optimize normalizing flows for lattice field theory. Yet how to best architect and optimize normalizing flows to minimize training costs, and the precise scaling of those costs, is still unclear. Developing a mature and optimized software library on which the community can build particular models and ensembles is also an important step towards applying flow-based models at state-of-the-art scales. We have seen recent, ongoing efforts in this direction.

It is likely that we will see the first demonstration of practical benefits beyond a proof of principle in applications that incorporate normalizing flows as components within existing calculations, for example by reducing variance in contexts where correlated samples are required or perhaps in improvements to HMC through the inclusion of normalizing flow components. The long term future of such a young research area cannot be reliably predicted. However, given our present understanding of the challenges and features of normalizing flows, we can consider one potential avenue: It is possible that the maximum impact on Monte Carlo sampling for lattice field theory will be achieved through a hybrid method, taking advantage of the natural hierarchy of scales of

most lattice field theory problems to organize Monte Carlo sampling in a nested fashion, with independent sampling of inner subvolumes via normalizing flows and traditional Markov Chain updates of boundary degrees of freedom. Several works are beginning to investigate this approach already. This scheme might naturally be combined with practical improvements, such as highly parallelized evaluation across subvolumes, elimination of storage costs of inner degrees of freedom, and exponential variance reduction techniques based on hierarchical sampling [97, 98]. If such a sampling scheme can be realized, it has the potential to drastically increasing the precision of lattice field theory calculations across a variety of fields, enabling a deeper understanding of physics ranging from condensed matter theory to QCD and particle physics beyond the Standard Model.

Acknowledgments

I am grateful to the organizers for the invitation to present this review and outlook. I would also like to thank Phiala Shanahan for helpful comments on this manuscript and all of my collaborators for their significant contributions towards the portion of the work reviewed here that is our own. This work is supported by the Swiss National Science Foundation (SNSF) under grant 200020_200424.

References

- [1] U. Wolff, *Critical slowing down*, *Nucl. Phys. B Proc. Suppl.* **17** (1990) 93.
- [2] USQCD collaboration, *Status and Future Perspectives for Lattice Gauge Theory Calculations to the Exascale and Beyond*, *Eur. Phys. J. A* **55** (2019) 199 [1904.09725].
- [3] P. Boyle et al., *Lattice QCD and the Computational Frontier*, in *Snowmass 2021*, 3, 2022 [2204.00039].
- [4] E.G. Tabak and E. Vanden-Eijnden, *Density estimation by dual ascent of the log-likelihood*, *Communications in Mathematical Sciences* **8** (2010) 217.
- [5] E.G. Tabak and C.V. Turner, *A family of nonparametric density estimation algorithms*, *Communications on Pure and Applied Mathematics* **66** (2013) 145.
- [6] G.E. Box and M.E. Muller, *A note on the generation of random normal deviates*, *The annals of mathematical statistics* **29** (1958) 610.
- [7] M. Luscher, *Trivializing maps, the Wilson flow and the HMC algorithm*, *Commun. Math. Phys.* **293** (2010) 899 [0907.5491].
- [8] G.P. Engel and S. Schaefer, *Testing trivializing maps in the Hybrid Monte Carlo algorithm*, *Comput. Phys. Commun.* **182** (2011) 2107 [1102.1852].
- [9] D. Rezende and S. Mohamed, *Variational inference with normalizing flows*, in *International Conference on Machine Learning*, pp. 1530–1538, PMLR, 2015 [1505.05770].
- [10] K. Hornik, M. Stinchcombe and H. White, *Multilayer feedforward networks are universal approximators*, *Neural networks* **2** (1989) 359.

- [11] G. Papamakarios, E. Nalisnick, D.J. Rezende, S. Mohamed and B. Lakshminarayanan, *Normalizing flows for probabilistic modeling and inference*, *The Journal of Machine Learning Research* **22** (2021) 2617 [1912.02762].
- [12] L. Dinh, D. Krueger and Y. Bengio, *NICE: Non-linear independent components estimation*, 1410.8516.
- [13] L. Dinh, J. Sohl-Dickstein and S. Bengio, *Density estimation using Real NVP*, 1605.08803.
- [14] R.T.Q. Chen, Y. Rubanova, J. Bettencourt and D. Duvenaud, *Neural Ordinary Differential Equations*, 1806.07366.
- [15] L. Zhang, L. Wang et al., *Monge-ampère flow for generative modeling*, 1809.10188.
- [16] Q. Zhang and Y. Chen, *Diffusion normalizing flow*, *Advances in Neural Information Processing Systems* **34** (2021) 16280 [2110.07579].
- [17] S. Kullback and R.A. Leibler, *On information and sufficiency*, *The annals of mathematical statistics* **22** (1951) 79.
- [18] L. Huang and L. Wang, *Accelerated Monte Carlo simulations with restricted Boltzmann machines*, *Phys. Rev. B* **95** (2017) 035105 [1610.02746].
- [19] J. Liu, Y. Qi, Z.Y. Meng and L. Fu, *Self-learning Monte Carlo method*, *Phys. Rev. B* **95** (2017) 041101 [1611.09364].
- [20] J. Liu, H. Shen, Y. Qi, Z.Y. Meng and L. Fu, *Self-learning monte carlo method and cumulative update in fermion systems*, *Phys. Rev. B* **95** (2017) 241104 [1611.09364].
- [21] X.Y. Xu, Y. Qi, J. Liu, L. Fu and Z.Y. Meng, *Self-learning quantum Monte Carlo method in interacting fermion systems*, *Phys. Rev. B* **96** (2017) 041119 [1612.03804].
- [22] Y. Nagai, H. Shen, Y. Qi, J. Liu and L. Fu, *Self-learning Monte Carlo method: Continuous-time algorithm*, *Phys. Rev. B* **96** (2017) 161102 [1705.06724].
- [23] H. Shen, J. Liu and L. Fu, *Self-learning monte carlo with deep neural networks*, *Phys. Rev. B* **97** (2018) 205140 [1801.01127].
- [24] C. Chen, X.Y. Xu, J. Liu, G. Batrouni, R. Scalettar and Z.Y. Meng, *Symmetry-enforced self-learning Monte Carlo method applied to the Holstein model*, *Phys. Rev. B* **98** (2018) 041102 [1802.06177].
- [25] Y. Liu, W. Wang, K. Sun and Z.Y. Meng, *Designer Monte Carlo simulation for the Gross-Neveu-Yukawa transition*, *Phys. Rev. B* **101** (2020) 064308 [1910.07430].
- [26] Y. Nagai, M. Okumura and A. Tanaka, *Self-learning Monte Carlo method with Behler-Parrinello neural networks*, *Phys. Rev. B* **101** (2020) 115111 [1807.04955].

- [27] M. Medvidovic, J. Carrasquilla, L.E. Hayward and B. Kulchytskyy, *Generative models for sampling of lattice field theories*, in *34th Conference on Neural Information Processing Systems*, 12, 2020 [2012.01442].
- [28] Y. Nagai and A. Tomiya, *Gauge covariant neural network for 4 dimensional non-abelian gauge theory*, 2103.11965.
- [29] Y. Nagai and A. Tomiya, *Self-learning Monte Carlo with equivariant Transformer*, 2306.11527.
- [30] M.S. Albergo, G. Kanwar and P.E. Shanahan, *Flow-based generative models for Markov chain Monte Carlo in lattice field theory*, *Phys. Rev. D* **100** (2019) 034515 [1904.12072].
- [31] K.A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller and P. Kessel, *Asymptotically unbiased estimation of physical observables with neural samplers*, *Phys. Rev. E* **101** (2020) 023304 [1910.13496].
- [32] D.C. Hackett, C.-C. Hsieh, M.S. Albergo, D. Boyda, J.-W. Chen, K.-F. Chen et al., *Flow-based sampling for multimodal distributions in lattice field theory*, 2107.00734.
- [33] P. de Haan, C. Rainone, M.C.N. Cheng and R. Bondesan, *Scaling Up Machine Learning For Quantum Field Theory with Equivariant Continuous Flows*, 2110.02673.
- [34] K.A. Nicoli, C.J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel et al., *Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models*, *Phys. Rev. Lett.* **126** (2021) 032001 [2007.07115].
- [35] M. Caselle, E. Cellini, A. Nada and M. Panero, *Stochastic normalizing flows as non-equilibrium transformations*, *JHEP* **07** (2022) 015 [2201.08862].
- [36] J. Komijani and M.K. Marinkovic, *Generative models for scalar field theories: how to deal with poor scaling?*, *PoS LATTICE2022* (2023) 019 [2301.01504].
- [37] M. Gerdes, P. de Haan, C. Rainone, R. Bondesan and M.C.N. Cheng, *Learning Lattice Quantum Field Theories with Equivariant Continuous Flows*, *SciPost Phys.* **15** (2023) 238 [2207.00283].
- [38] D. Albandea, L. Del Debbio, P. Hernández, R. Kenway, J. Marsh Rossney and A. Ramos, *Learning trivializing flows*, *Eur. Phys. J. C* **83** (2023) 676 [2302.08408].
- [39] D. Albandea, L. Del Debbio, P. Hernández, R. Kenway, J. Marsh Rossney and A. Ramos, *Learning Trivializing Flows in a ϕ^4 theory from coarser lattices*, *PoS LATTICE2023* (2023) 013.
- [40] A. Singha, D. Chakrabarti and V. Arora, *Conditional normalizing flow for Markov chain Monte Carlo sampling in the critical region of lattice field theory*, *Phys. Rev. D* **107** (2023) 014512.

- [41] K. Cranmer, G. Kanwar, S. Racanière, D.J. Rezende and P.E. Shanahan, *Advances in machine-learning-based sampling motivated by lattice quantum chromodynamics*, *Nature Rev. Phys.* **5** (2023) 526 [2309.01156].
- [42] J. Köhler, L. Klein and F. Noé, *Equivariant flows: exact likelihood generative learning for symmetric densities*, in *International Conference on Machine Learning*, pp. 5361–5370, PMLR, 2020 [2006.02425].
- [43] G. Kanwar, *Machine Learning and Variational Algorithms for Lattice Field Theory*, Ph.D. thesis, MIT, 2021. 2106.01975.
- [44] G. Kanwar, M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, S. Racanière et al., *Equivariant flow-based sampling for lattice gauge theory*, *Phys. Rev. Lett.* **125** (2020) 121601 [2003.06413].
- [45] D. Boyda, G. Kanwar, S. Racanière, D.J. Rezende, M.S. Albergo, K. Cranmer et al., *Sampling using $SU(N)$ gauge equivariant flows*, *Phys. Rev. D* **103** (2021) 074504 [2008.05456].
- [46] D.J. Rezende, G. Papamakarios, S. Racanière, M.S. Albergo, G. Kanwar, P.E. Shanahan et al., *Normalizing flows on tori and spheres*, 2002.02428.
- [47] M. Favoni, A. Ipp, D.I. Müller and D. Schuh, *Lattice Gauge Equivariant Convolutional Neural Networks*, *Phys. Rev. Lett.* **128** (2022) 032003 [2012.12901].
- [48] Y. Namekawa, K. Kashiwa, H. Matsuda, A. Ohnishi and H. Takase, *Improving efficiency of the path optimization method for a gauge theory*, *Phys. Rev. D* **107** (2023) 034509 [2210.05402].
- [49] C. Lehner and T. Wettig, *Gauge-equivariant neural networks as preconditioners in lattice QCD*, *Phys. Rev. D* **108** (2023) 034503 [2302.05419].
- [50] D. Knüttel, C. Lehner and T. Wettig, *Gauge-equivariant multigrid neural networks*, *PoS LATTICE2023* (2023) 037.
- [51] J. Aronsson, D.I. Müller and D. Schuh, *Geometrical aspects of lattice gauge equivariant convolutional neural networks*, 2303.11448.
- [52] C. Lehner and T. Wettig, *Gauge-equivariant pooling layers for preconditioners in lattice QCD*, 2304.10438.
- [53] K. Holland, A. Ipp, D.I. Müller and U. Wenger, *Fixed point actions from convolutional neural networks*, *PoS LATTICE2023* (2023) 038 [2311.17816].
- [54] S. Bacchio, P. Kessel, S. Schaefer and L. Vaitl, *Learning trivializing gradient flows for lattice gauge theories*, *Phys. Rev. D* **107** (2023) L051504 [2212.08469].
- [55] M.S. Albergo, G. Kanwar, S. Racanière, D.J. Rezende, J.M. Urban, D. Boyda et al., *Flow-based sampling for fermionic lattice field theories*, *Phys. Rev. D* **104** (2021) 114507 [2106.05934].

- [56] M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, G. Kanwar, S. Racanière et al., *Flow-based sampling in the lattice Schwinger model at criticality*, *Phys. Rev. D* **106** (2022) 014514 [2202.11712].
- [57] P. Boyle, T. Izubuchi, L. Jin, C. Jung, C. Lehner, N. Matsumoto et al., *Use of Schwinger-Dyson equation in constructing an approximate trivializing map*, *PoS LATTICE2022* (2023) 229 [2212.11387].
- [58] R. Abbott et al., *Sampling QCD field configurations with gauge-equivariant flow models*, *PoS LATTICE2022* (2023) 036 [2208.03832].
- [59] R. Abbott et al., *Normalizing flows for lattice gauge theory in arbitrary space-time dimension*, 2305.02402.
- [60] M. Dibak, L. Klein, A. Krämer and F. Noé, *Temperature steerable flows and boltzmann generators*, *Physical Review Research* **4** (2022) L042005.
- [61] R. Abbott, A. Botev, D. Boyda, D.C. Hackett, G. Kanwar, S. Racanière et al., *Applications of flow models to the generation of correlated lattice QCD ensembles*, to appear (2024) .
- [62] R.H. Swendsen and J.-S. Wang, *Replica Monte Carlo Simulation of Spin-Glasses*, *Phys. Rev. Lett.* **57** (1986) 2607.
- [63] M. Invernizzi, A. Krämer, C. Clementi and F. Noé, *Skipping the replica exchange ladder with normalizing flows*, *The Journal of Physical Chemistry Letters* **13** (2022) 11643–11649.
- [64] R. Abbott et al., *Practical applications of machine-learned flows on gauge fields*, to appear (2024) .
- [65] C. Bonanno, A. Nada and D. VDACCHINO, *Out-of-equilibrium simulations to fight topological freezing*, *PoS LATTICE2023* (2023) 005 [2310.11979].
- [66] H. Wu, J. Köhler and F. Noé, *Stochastic normalizing flows*, *Advances in Neural Information Processing Systems* **33** (2020) 5933 [2002.06707].
- [67] S. Lawrence and Y. Yamauchi, *Normalizing Flows and the Real-Time Sign Problem*, *Phys. Rev. D* **103** (2021) 114509 [2101.05755].
- [68] J.M. Pawłowski and J.M. Urban, *Flow-based density of states for complex actions*, *Phys. Rev. D* **108** (2023) 054511 [2203.01243].
- [69] M. Rodekamp, E. Berkowitz, C. Gäntgen, S. Krieg, T. Luu and J. Ostmeier, *Mitigating the Hubbard sign problem with complex-valued neural networks*, *Phys. Rev. B* **106** (2022) 125139 [2203.00390].
- [70] M. Rodekamp, E. Berkowitz, M. Dincă, C. Gäntgen, S. Krieg and T. Luu, *From Theory to Practice: Applying Neural Networks to Simulate Real Systems with Sign Problems*, 2311.18312.

- [71] W. Detmold, G. Kanwar, Y. Lin, P.E. Shanahan and M.L. Wagman, *Signal-to-noise improvement through neural network contour deformations for 3D SU(2) lattice gauge theory*, *PoS LATTICE2023* (2023) 043 [2309.00600].
- [72] R. Abbott et al., *Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions*, *Phys. Rev. D* **106** (2022) 074506 [2207.08945].
- [73] J. Finkenrath, *Tackling critical slowing down using global correction steps with equivariant flows: the case of the Schwinger model*, 2201.02216.
- [74] S. Chen, O. Savchuk, S. Zheng, B. Chen, H. Stoecker, L. Wang et al., *Fourier-flow model generating Feynman paths*, *Phys. Rev. D* **107** (2023) 056001 [2211.03470].
- [75] N. Matsumoto, R.C. Brower and T. Izubuchi, *Decimation map in 2D for accelerating HMC*, *PoS LATTICE2023* (2023) 033 [2312.04800].
- [76] R. Abbott et al., *Multiscale normalizing flows for gauge theories*, to appear (2024) .
- [77] S. Foreman, T. Izubuchi, L. Jin, X.-Y. Jin, J.C. Osborn and A. Tomiya, *HMC with Normalizing Flows*, *PoS LATTICE2021* (2022) 073 [2112.01586].
- [78] S. Foreman, X.-Y. Jin and J.C. Osborn, *MLMC: Machine Learning Monte Carlo for Lattice Gauge Theory*, *PoS LATTICE2023* (2023) 036 [2312.08936].
- [79] A. Faraz, A. Singha, D. Chakrabarti and V. Arora, *Scalable Lattice Sampling using Factorized Generative Models*, 2308.08615.
- [80] M. Caselle, E. Cellini and A. Nada, *Sampling the lattice Nambu-Goto string using Continuous Normalizing Flows*, 2307.01107.
- [81] M. Caselle, E. Cellini and A. Nada, *Sampling Nambu-Goto theory using Normalizing Flows*, *PoS LATTICE2023* (2023) 015.
- [82] C. Chamness, K. Orginos and D. Kovner, *Trivializing Flow in 2D O(3) sigma model*, *PoS LATTICE2023* (2023) 008.
- [83] A. D’Adda, M. Luscher and P. Di Vecchia, *A $1/n$ Expandable Series of Nonlinear Sigma Models with Instantons*, *Nucl. Phys. B* **146** (1978) 63.
- [84] E. Witten, *Instantons, the Quark Model, and the $1/n$ Expansion*, *Nucl. Phys. B* **149** (1979) 285.
- [85] M. Hasenbusch, *Fighting topological freezing in the two-dimensional CPN-1 model*, *Phys. Rev. D* **96** (2017) 054504 [1706.04443].
- [86] M. Berni, C. Bonanno and M. D’Elia, *Large- N expansion and θ -dependence of 2d CP $^{N-1}$ models beyond the leading order*, *Phys. Rev. D* **100** (2019) 114509 [1911.03384].
- [87] M.S. Albergo, D. Boyda, D.C. Hackett, G. Kanwar, K. Cranmer, S. Racanière et al., *Introduction to Normalizing Flows for Lattice Field Theory*, 2101.08176.

- [88] A. Tomiya and S. Terasaki, *GomalizingFlow.jl: A Julia package for Flow-based sampling algorithm for lattice field theory*, [2208.08903](#).
- [89] K.A. Nicoli, C. Anders, L. Funcke, K. Jansen, S. Nakajima and P. Kessel, *NeuLat: a toolbox for neural samplers in lattice field theories*, *PoS LATTICE2023* (2023) 286.
- [90] P. Bialas, P. Korcyl and T. Stebel, *Gradient estimators for normalising flows*, [2202.01314](#).
- [91] L. Vaitl, K.A. Nicoli, S. Nakajima and P. Kessel, *Path-gradient estimators for continuous normalizing flows*, in *International Conference on Machine Learning*, pp. 21945–21959, PMLR, 2022 [[2206.09016](#)].
- [92] L. Vaitl, K.A. Nicoli, S. Nakajima and P. Kessel, *Gradients should stay on path: better estimators of the reverse-and-forward kl divergence for normalizing flows*, *Machine Learning: Science and Technology* **3** (2022) 045006 [[2207.08219](#)].
- [93] P. Białas, P. Korcyl and T. Stebel, *Training normalizing flows with computationally intensive target probability distributions*, [2308.13294](#).
- [94] K.A. Nicoli, C.J. Anders, T. Hartung, K. Jansen, P. Kessel and S. Nakajima, *Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories*, *Phys. Rev. D* **108** (2023) 114501 [[2302.14082](#)].
- [95] R. Abbott et al., *Aspects of scaling and scalability for flow-based sampling of lattice QCD*, *Eur. Phys. J. A* **59** (2023) 257 [[2211.07541](#)].
- [96] L. Del Debbio, J.M. Rossney and M. Wilson, *Efficient modeling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability*, *Phys. Rev. D* **104** (2021) 094507 [[2105.12481](#)].
- [97] M. Luscher and P. Weisz, *Locality and exponential error reduction in numerical lattice gauge theory*, *JHEP* **09** (2001) 010 [[hep-lat/0108014](#)].
- [98] M. Cè, L. Giusti and S. Schaefer, *Domain decomposition, multi-level integration and exponential noise reduction in lattice QCD*, *Phys. Rev. D* **93** (2016) 094507 [[1601.04587](#)].