

Sampling $U(1)$ gauge theory using a retrainable conditional flow-based model

Ankur Singha¹,[✉] Dipankar Chakrabarti¹,[✉] and Vipul Arora²

¹*Department of Physics, Indian Institute of Technology Kanpur, Kanpur-208016, India*

²*Department of Electrical Engineering, Indian Institute of Technology Kanpur, Kanpur-208016, India*



(Received 24 July 2023; accepted 10 October 2023; published 30 October 2023)

Sampling topological quantities in the Monte Carlo simulation of lattice gauge theory becomes challenging as we approach the continuum limit ($a \rightarrow 0$) of the theory. In this work, we introduce a conditional normalizing flow model to sample $U(1)$ gauge theory in two dimensions, aiming to mitigate the impact of topological freezing when dealing with smaller values of the $U(1)$ bare coupling. To train the conditional flow model, we utilize samples generated by hybrid Monte Carlo method, ensuring that the autocorrelation in topological quantities remains low. Subsequently, we employ the trained model to extrapolate the coupling parameter to values where training was not performed. We thoroughly examine the quality of the model in this region and generate uncorrelated samples, significantly reducing the occurrence of topological freezing. Furthermore, we propose a retrainable approach that utilizes the model's own samples to enhance the generalization capability of the conditional model. This method enables sampling for coupling values far beyond the initial training region, expanding the model's applicability.

DOI: [10.1103/PhysRevD.108.074518](https://doi.org/10.1103/PhysRevD.108.074518)

I. INTRODUCTION

In the field of lattice theory, Monte Carlo simulation methods are utilized to sample configurations of the lattice. These samples are generated based on a distribution determined by the action of the lattice theory. The choice of the action parameter value used in generating the lattice samples determines the computational cost of the simulation. As we approach the critical region of a lattice system or move towards the continuum limit of a lattice field theory, the samples become highly correlated. Within the critical region, the integrated autocorrelation time, which measures the level of correlation, increases rapidly and diverges at the critical point. In a finite-size lattice, the critical point corresponds to the peak of the autocorrelation curve, leading to a phenomenon known as critical slowing down [1,2]. This has been a major problem in the simulation of lattice systems near the critical region. If proper care is not taken, this may lead to biased estimation of observable.

Monte Carlo simulation encounters similar difficulties when exploring the topological sector in lattice gauge theory, such as $U(1)$ gauge theory in 2D. When we move towards finer lattices, the topological charge freezes, i.e., the hybrid Monte Carlo (HMC) sampling is restricted only

to one or two topological sectors for a longer simulation run. This implies a larger integrated autocorrelation time for topological charge which is well known as topological freezing. Numerous efforts have been made to mitigate the impact of critical slowing down in statistical systems and lattice quantum field theory [3–5]. However, in lattice gauge theory it still remains a challenging task.

These days, machine learning-based solutions to the critical slowing down problem has become popular. Various machine learning algorithms have been applied for statistical physics, condensed matter problems and lattice field theory [6–29]. Several generative learning algorithms have recently been developed to avoid the difficulty in lattice field theory [30–41]. In the flow-based approach [30,33] normalizing flow (NF) models are trained for a action parameter with reverse Kullback-Leibler (KL) divergence. But this kind of self-learning method has a major issue of mode collapse which may lead to inefficient modeling of complicated multimodal distributions [10,36]. In contrast a generative model trained with forward KL are mode covering and applying the Metropolis-Hastings (MH) algorithm can produce the correct statistics. In the $U(1)$ gauge theory, this issue of mode collapse is discussed in Sec. IV A. In our earlier works [42,43] we presented a method for sampling lattice configurations near the critical regions using conditional normalizing flow (C-NF) (conditional normalizing flow has also recently been used in [44]) and conditional generative adversarial network to reduce the problem of critical slowing down for scalar field theory and lattice Gross-Neveu model in two dimensions. We have

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


shown that the C-NF model [42] trained in the noncritical region can produce samples for parameter values in the critical region. In this work, we propose an application of equivariant flow to construct a conditional flow model for sampling $U(1)$ gauge theory in two dimensions. To the best of our knowledge the conditional flow model, based on the action parameter, has not been employed in the sampling of lattice gauge theory.

In $U(1)$ gauge theory, the target distribution for a given action parameter β can be defined as


$$p(U_\mu(n)|\beta) = \frac{1}{Z} e^{-S(U,\beta)}, \quad (1)$$

where $U_\mu(n)$ denotes the $U(1)$ link variable field, β denotes the lattice action parameter and Z is the partition function defined as $Z = \sum_U e^{-S(U,\beta)}$. The action parameter β is the conditional parameter, and we train a conditional flow model on lattice ensembles corresponding to different β values. We partition the action parameter into two sets β_L and β_S based on the integrated autocorrelation time (τ_{int}) of topological charge. β_L corresponds to large β values where topological freezing is dominant in HMC simulations as shown in Fig. 1(a) and β_S corresponds to smaller values of β where the τ_{int} is small, and hence the samples fluctuate among different topological sectors as shown in Fig. 1(b). Due of the lower autocorrelation time, the HMC simulation cost is minimal in the β_S set. Hence, we train a C-NF $\tilde{p}(U_\beta)$ with HMC samples from $p(U, \beta_S)$.

We train the C-NF model to be a generalised model over β parameters. The model is then extrapolated to



(a) $\beta = 7$



(b) $\beta = 3.5$

FIG. 1. The fluctuation of topological charge for 50k Markov steps simulated by HMC at (a) $\beta = 7$ and (b) $\beta = 3.5$.

larger values of β , i.e., in the β_L set to generate different ensembles of lattice configurations. However, the extrapolated model may not provide samples from the true target distribution. But the exactness can be guaranteed by using the MH algorithm. Therefore, we use the extrapolated model $\tilde{p}(U(\beta_L))$ at large β values as proposal for constructing a Markov Chain via an independent MH algorithm [30]. The quality of the extrapolated model may deteriorate as we move further away from the training region, particularly for distant β values. In such cases, we can employ a retrainable method to sample at those β values. This retrainable method utilizes the model's own samples at intermediate β value to enhance the generalization of the C-NF model, enabling better sampling capabilities at points that are far away from the training set.

II. $U(1)$ GAUGE THEORY

The lattice action for $U(1)$ gauge theory in $1+1$ dimensions can be written as

$$S(U) = -\beta \sum_{n,\mu < \nu} \text{Re}[U_{\mu\nu}(n)], \quad (2)$$

where, the plaquette $U_{\mu\nu}$ is defined as

$$U_{\mu\nu} = U_\mu(n)U_\nu(n+\hat{\mu})U_\mu(n+\hat{\nu})U_\nu(n). \quad (3)$$

$U_\mu(n)$ can be written in terms of angular variable as $U_\mu(n) = e^{i\theta_\mu(n)}$. The action is symmetric under the transformation

$$U_\mu(n) \rightarrow e^{i\alpha(n)}U_\mu(n)e^{-i\alpha(n+\hat{\mu})}. \quad (4)$$

In terms of angular variable the plaquette becomes

$$U_{\mu\nu}(n) = e^{i[\theta_\mu(n)+\theta_\nu(n+\hat{\mu})-\theta_\nu(n+\hat{\nu})-\theta_\mu(n)]} \quad (5)$$

$$= e^{i\theta_{\mu\nu}(n)}. \quad (6)$$

So, the action can be written as

$$S(U) = -\beta \sum_n \cos \theta_{\mu\nu}(n), \quad (7)$$

which is symmetric under

$$\theta_\mu(n) \rightarrow \alpha(n) + \theta_\mu(n) - \alpha(n + \hat{\mu}), \quad (8)$$

where $\alpha(n) \in \mathbb{R}$.

The action in Eq. (7) is used in Eq. (1) for sampling in both HMC and C-NF model. The observable we calculate here is the topological charge defined as

$$Q = \frac{1}{2\pi} \sum_n \arg[U_{\mu\nu}(n)], \quad (9)$$

$$\text{where } \arg[(U_{\mu\nu})] \in [-\pi, \pi]. \quad (10)$$

III. C-NF MODEL EMPLOYING EQUIVARIANT FLOW

In a flow-based model, starting with a gauge symmetric prior distribution does not guarantee that the resulting output distribution will also be gauge symmetric. For the output distribution to maintain gauge symmetry, the flow transformation must be an equivariant transformation, meaning it should commute with the gauge transformation. Implementing such equivariance in a flow-based model is nontrivial. However, if the flow transformation only affects the gauge invariant quantity, then the condition mentioned above is easily satisfied. In such cases, the flow-based model can maintain gauge symmetry by exclusively transforming the gauge invariant component. In $U(1)$ gauge theory, the link variable $U_\mu(n)$ are not itself gauge invariant. For such construction one has to find one to one correspondence between the gauge invariant quantity and the link variable. For example, the plaquette variable $P_{\mu\nu}(n)$ is one such gauge invariant quantity which can be used for construction of equivariant flow. We use the equivariant flow construction following Ref. [33]. In 2D, the link variables are tensor of shape $(2, L, L)$ and the plaquette variables are tensor of shape (L, L) . The update in $P_{\mu\nu}$ can be translated into update of a single U_μ using suitable mask pattern. Note that a single update of a link variable in 2D will update two adjacent plaquettes. Therefore, in a coupling layer we need three subclass of input plaquettes as described below, obtained by suitable mask namely active, passive and frozen [33]. The active variables ($P_{\mu\nu}^a$) are transformed by the coupling layer, the frozen variable ($P_{\mu\nu}^f$) are not transformed by this coupling layer and the passive variables ($P_{\mu\nu}^{ps}$) are those which gets transformed indirectly due to a link update. We choose the mask such that the update in the link variable U_μ is sparse and hence scalable to larger lattice sizes. We use a total of 32 coupling layers in the C-NF model.

The link update in a coupling layer happens in two steps, in the first step, we take care of translating $P_{\mu\nu}$ to U_μ and the second step corresponds to the actual flow in the $P_{\mu\nu}(n)$ variable. For the second step, we use the noncompact projection (NCP) flow [45]. In NCP flow, we use fully convolutional layers to transform the frozen variables of the coupling layers. We have used 48 such coupling layers, and each coupling layer has three hidden convolutional neural networks with a number of filters 16, 32, 16. We use LeakyReLU activation in the hidden layers. Since we use coupling layers the conditional parameters can be concatenated to frozen variables and passed through the neural network, which is evaluated in the forward direction only.

IV. NUMERICAL EXPERIMENTS AND RESULTS

In this section, we present the details of our numerical experiment and the resulting outcomes. We outline the

training and sampling processes and the dataset preparation used to train the C-NF model.

A. Mode collapse in $U(1)$ gauge theory

In the literature, there are several flow-based works which try to accelerate the sampling of lattice field theory using reverse KL (RKL) divergence. However, one major disadvantage RKL comes with is the mode-collapsing behavior. On the other hand, flow-based models trained using forward KL (FKL) are mode covering. In this section, we investigate the problem of mode collapse in a flow-based model for $U(1)$ gauge theory. We train two flow-based models $q_R(U, \Theta)$ and $q_F(U, \Psi)$ using RKL and FKL divergence, respectively. After training both models to 50% effective sample size (ESS), we estimate the negative log-likelihood (NLL). Two kinds of NLL for the flow models are used for the investigation, namely \mathcal{L}_1 and \mathcal{L}_2 ,

$$\mathcal{L}_1 = E_{U \sim p(U|\beta)}[\log q(U, \Theta)], \quad (11)$$

where $q(U, \Theta)$ is either trained with the FKL or RKL and the expectation E is taken over samples generated from the true distribution (e.g., HMC simulation). This type of NLL is best suited for detecting mode collapse in a generative model.

$$\mathcal{L}_2 = E_{U \sim q_U, \Theta}[\log p(U|\beta)], \quad (12)$$

where in this case, the expectation is taken over the samples generated from the model itself and the expectation of $p(U)$ is estimated. This metric alone is not sufficient to investigate the mode collapse as explained in the Appendix. However, by the combination of \mathcal{L}_1 and \mathcal{L}_2 , one can detect the mode collapse.

For HMC $\mathcal{L}_1 = \mathcal{L}_2$ and is given by

$$\mathcal{L}_{hmc} = E_{x \sim p(U|\beta)}[\log p(U|\beta)] \quad (13)$$

We have estimated all three NLL types, \mathcal{L}_1 , \mathcal{L}_2 , and \mathcal{L}_{hmc} shown in the Table I. We see that the \mathcal{L}_1 is quite high for the RKL and highly deviates from the \mathcal{L}_{hmc} . For RKL, \mathcal{L}_1 is 270.61, which is quite higher than the \mathcal{L}_2 (196.88). This indicates that the RKL model has not explored the distribution's parameter space as HMC does. Hence, RKL generates samples in a specific region of the distribution,

TABLE I. Comparison of the different NLL types. For HMC \mathcal{L}_1 and \mathcal{L}_2 are equal.

NLL comparison		
Model	\mathcal{L}_1	\mathcal{L}_2
RKL [34]	270.611	196.882
FKL	199.715	196.982
HMC	196.666	

leading to the disagreement in \mathcal{L}_2 and \mathcal{L}_1 . On the other hand, for the FKL, \mathcal{L}_1 is low and very close to the \mathcal{L}_{hmc} . Moreover, \mathcal{L}_1 and \mathcal{L}_2 are also quite close (196.98 and 199.71) for FKL. More discussions on this issue are given in the Appendix.

Note that an RKL model alone is sufficient to observe the mode collapse, and one can use the method [34] for $U(1)$ gauge theory to check the large disagreement of \mathcal{L}_1 and \mathcal{L}_2 . The observed difference remains constant while training, even if we increase the ESS. This is what one expects in an online training process.

B. Training dataset

To train the C-NF model, we generate 10 different ensembles of lattice configurations, each corresponding to a different value of $\beta = \beta_S$: $\{1.0, 1.5, 1.8, 2, 2.2, 2.5, 2.8, 3, 3.2, 3.5\}$. We choose the training dataset such that there is no topological freezing. The largest value of β_S is 3.5, where integrated autocorrelation time is ≈ 23.45 . The behavior of topological charge is shown in the Fig. 1(b) for $\beta_S = 3.5$. For larger β the topological freezing sets in as shown in Fig. 1(a). To prepare training ensembles, we employ HMC simulation and adjust the HMC parameter for each β_S value to achieve an acceptance rate of approximately $\approx 85\%$. We have used a nonuniform ensemble size for each β_S . For the largest value of β_S , we generate 15,000 samples and as we move to the next lower β_S we reduce ensemble size by 500. We perform all the numerical experiments on a 16×16 lattice.

C. Training and sampling

In the C-NF model, the input lattice configurations are tensors with a shape of $(2, 16, 16)$. These tensors are concatenated with their corresponding ensemble labels or conditional parameters, denoted as $\mathbb{1}\beta$. The tensor $\mathbb{1}$ has the same shape as the lattice configurations, $(2, 16, 16)$, with all elements set to identity. It implies that all masking patterns will be applied to the condition tensors along with the lattice configurations.

During training, we randomly select a batch from any value of β_S . We use a batch size of 512 to calculate the gradient at each iteration, and the model weights are updated after every 10 iterations. The C-NF model is trained using forward KL divergence and an Adam optimizer with an initial learning rate of 0.001. Increasing the ESS of the C-NF model can be challenging and may reach a plateau during the training. We incorporate a decay of learning rate of 0.5 at intervals of 25,000 training iterations. We have found this approach effective in increasing the ESS for the C-NF model. It's important to note that the C-NF model may exhibit overfitting and perform well only for the training β values. To overcome this, we condition only every fourth coupling layer allowing for better generalization and performance on unseen β values. We stop the training when the ESS reach above 30% and the

increment is less than 3% for next 5000 consecutive iterations. Note that, with further training or a more optimized architecture, one can potentially achieve a higher ESS. However, our objective is to assess the sample quality as we extrapolate for large β values. While training, we also monitor the acceptance rate periodically after every 10,000 iterations in the MH algorithm. At the end of training, we achieve an acceptance rate of approximately 65% in the training region.

After training, we extrapolate the model for large β values, $\beta_L = \{5.5, 6, 6.5, 7, 7.5\}$. Using the extrapolated model, we generate proposals for the MH step in order to construct a Markov chain. For each value of β , we generate an ensemble consisting of 10^6 lattice configurations. Once training is over, the flow-based model enables us to efficiently generate such large ensembles without any significant challenges.

D. Results

We calculate different observables on the ensemble generated from the C-NF model. In Fig. 2 we have shown the integrated autocorrelation time for HMC and C-NF model calculated for a topological charge. During the estimation of observables on the ensemble, we consider every 20th configuration to reduce any effect of autocorrelation. We see that for HMC simulation, the effect of topological freezing increases rapidly as we move towards larger β . Since the generation in the C-NF model is inherently uncorrelated, the autocorrelation depends on the acceptance rate in MH step. There is a massive gain as we move towards larger β where we have not trained the model. In the extrapolated region, the acceptance rate in the MH step is almost constant. This indicates that our model has learned the conditional distribution very well. In Fig. 3, we show that the acceptance rate is between 60% to 50% for the extrapolated parameter values in β_L .




FIG. 2. Integrated autocorrelation time calculated from the C-NF model and HMC simulation. The solid (green) line represents τ_{int} in the HMC simulation and dashed (orange) line represents τ_{int} in C-NF.




FIG. 3. Acceptance rate calculated from the extrapolated C-NF model for different β values. The acceptance rate shows a very small decline over a wide range of nontraining values of β .

This allows one to store a single model and generate ensembles at multiple β values.

Figure 4 shows the fluctuation of topological charge with the Monte Carlo Markov chain (MCMC) time. We see clear freezing in the HMC Markov chain. In contrast, the C-NF model has reduced it significantly. This is a significant gain one can obtain from a flow-based model.

E. Retraining method for distant β values

If we want to extrapolate the model far away from the training region, the acceptance rate may decrease. This is due to the lack of good generalization for the conditional model over a wide range of β values. For example, if we want to generate samples for $\beta = 9$, then the MH acceptance rate drops to $\sim 30\%$. We address this issue by utilizing samples from nontraining regions of β values. For extrapolation at $\beta = 9$, first, we generate samples from the C-NF model for an intermediate β , say $\beta = 6$, where the acceptance rate is $\sim 55\%$. We add these newly generated samples and retrain the C-NF model using forward KL for 10,000 iterations along with the previously used samples. We reduce 50% samples for each $\beta \in \beta_S$. This will improve the generalization of the C-NF model for larger β values.




FIG. 4. Topological freezing is shown for 50,000 Markov chains in both the HMC simulation and the C-NF model. A significant gain from the C-NF model in reducing the effect of topological freezing is clearly visible.

We observed an increase in acceptance rate with the retraining up to $\sim 40\%$ for $\beta = 9$. Note that if β is not too far from the training region and the ESS does not fall off drastically, retraining is not required.

V. SUMMARY AND CONCLUSION

Based on our previous work, it has been observed that the conditional flow model has proven to be successful when applied to scalar and Gross-Neveu models. Building and testing the conditional flow-based model on gauge theory is crucial to advance further in the development of lattice QCD.

In the simulation of $U(1)$ gauge theory using the MCMC method, the correlation between the samples becomes stronger as lattice spacing decrease i.e., β increases. As β reaches higher values, a phenomenon called topological freezing is observed in the MCMC simulation. To tackle this issue, we have developed a conditional flow-based model (C-NF) specifically designed for sampling the $U(1)$ gauge theory in 2D when β is large. This model takes β as a conditioning parameter and is trained using ensembles of β values where there are low autocorrelations and minimal effects of topological freezing. The training process continues until the ESS for the training ensembles reaches approximately 30%.

Once the training is completed, we employ the model to generate samples for larger β values using the Metropolis-Hastings algorithm. The quality of the extrapolated model depends on how well it has learned a generalized distribution over the β parameter. We achieve an acceptance rate of approximately 50–60% in the MH algorithm across a wide range of β values that were not included in the training data. For $\beta = 7.5$, the acceptance rate in the MH algorithm is approximately 52%. However, as β increases further, the acceptance rate gradually decreases. In such cases, if necessary, one can use a retrainable method for sampling points far beyond the training region. This method utilizes the model's samples at intermediate β values to improve the conditional generation of samples.

While the conditional flow model has shown effectiveness in certain cases, it is generally not scalable for sampling high-dimensional distributions. One possible future direction is to tackle the challenging task of constructing a conditional flow model for sampling high-dimensional lattice gauge theory.

ACKNOWLEDGMENTS

The work is supported by SERB under the Grant No. CRG/2021/003466.

APPENDIX: MODE COLLAPSE IN A FLOW MODEL

Let us examine a multimodal target distribution, denoted as $p(x)$, which encompasses two modes as illustrated in Fig. 5. To represent this distribution, we employ an NF




FIG. 5. Mode collapse in a flow model.

model denoted as $q(x, \Theta)$. This NF model can be trained using either forward KL or reverse KL divergence.

We examine two types of NLL metrics to investigate mode collapse in the NF model, aiming to determine which one is more effective in identifying a mode collapse.

We use HMC as our baseline for studying the mode collapse phenomenon.

In \mathcal{L}_2 , samples are drawn from the distribution $p(x)$, where $x \sim p(x)$, and the density is subsequently estimated using the NF model. In other words, we are estimating $E_{x \sim p(x)}[\log q(x, \Theta)]$.

When there is no mode collapse, the model generates samples in both modes and as a result, $E_{x \sim p(x)}[\log q(x, \Theta)]$ aligns with \mathcal{L}_{hmc} .

Now, consider the scenario in which mode collapse occurs. Since the model learns only a single mode, $q(x, \Theta)$ becomes substantially elevated at that specific mode, as illustrated in Fig. 5. Consequently, \mathcal{L}_1 will be notably higher compared to the situation where all modes are learned. Therefore, this serves as a suitable metric for identifying mode collapse.

In \mathcal{L}_2 , samples are generated from the trained NF model and the expectation value of $\log p(x)$ is estimated. If we assume there is no mode collapse, then the model will generate samples $x \sim q(x)$ in both modes and $E_x[\log p(x)]$ will match with \mathcal{L}_{hmc} . If there is a mode collapse, the model will generate samples only in a single mode. But the estimate $E_x[\log p(x)]$ will still match with \mathcal{L}_{hmc} as the estimated quantity is $\log p(x)$. Thus it does not serve as a good metric for mode collapse detection.

-
- [1] Ulli Wolff, Critical slowing down, *Nucl. Phys. B, Proc. Suppl.* **17**, 93 (1990).
 - [2] Stefan Schaefer, Rainer Sommer, and Francesco Virotta, Critical slowing down and error analysis in lattice QCD simulations, *Nucl. Phys.* **B845**, 93 (2011).
 - [3] Alberto Ramos, Playing with the kinetic term in the HMC, *Proc. Sci. LATTICE2012* (2012) 193.
 - [4] Arjun Singh Gambhir and Kostas Orginos, Improved sampling algorithms in lattice QCD, *Proc. Sci. LATTICE2014* (2015) 043 [arXiv:1506.06118].
 - [5] Michael G. Endres, Richard C. Brower, William Detmold, Kostas Orginos, and Andrew V. Pochinsky, Multiscale Monte Carlo equilibration: Pure Yang-Mills theory, *Phys. Rev. D* **92**, 114516 (2015).
 - [6] David Albandea, Pilar Hernández, Alberto Ramos, and Fernando Romero-López, Topological sampling through windings, *Eur. Phys. J. C* **81**, 873 (2021).
 - [7] Kai Zhou, Lingxiao Wang, Long-Gang Pang, and Shuzhe Shi, Exploring QCD matter in extreme conditions with machine learning, arXiv:2303.15136.
 - [8] Christoph Lehner and Tilo Wettig, Gauge-equivariant pooling layers for preconditioners in lattice QCD, arXiv:2304.10438.
 - [9] Jimmy Aronsson, David I. Müller, and Daniel Schuh, Geometrical aspects of lattice gauge equivariant convolutional neural networks, arXiv:2303.11448.
 - [10] Kim A. Nicoli, Christopher J. Anders, Tobias Hartung, Karl Jansen, Pan Kessel, and Shinichi Nakajima, Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories, arXiv:2302.14082.
 - [11] Simone Bacchio, Pan Kessel, Stefan Schaefer, and Lorenz Vaitl, Learning trivializing gradient flows for lattice gauge theories, *Phys. Rev. D* **107**, L051504 (2023).
 - [12] Kai Zhou, Gergely Endrődi, Long-Gang Pang, and Horst Stöcker, Regressive and generative neural networks for scalar field theory, *Phys. Rev. D* **100**, 011501 (2019).
 - [13] Dian Wu, Lei Wang, and Pan Zhang, Solving statistical mechanics using variational autoregressive networks, *Phys. Rev. Lett.* **122**, 080602 (2019).
 - [14] Jan M. Pawłowski and Julian M. Urban, Reducing auto-correlation times in lattice simulations with generative adversarial networks, *Mach. Learn.* **1**, 045011 (2020).
 - [15] Kim A. Nicoli, Shinichi Nakajima, Nils Strodthoff, Wojciech Samek, Klaus-Robert Müller, and Pan Kessel, Asymptotically unbiased estimation of physical observables with neural samplers, *Phys. Rev. E* **101**, 023304 (2020).
 - [16] Lei Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
 - [17] Japneet Singh, Mathias Scheurer, and Vipul Arora, Conditional generative models for sampling and phase transition indication in spin systems, *SciPost Phys.* **11**, 043 (2021).

- [18] Juan Carrasquilla, Machine learning for quantum matter, *Adv. Phys.* **5**, 1797528 (2020).
- [19] Junwei Liu, Huitao Shen, Yang Qi, Zi Yang Meng, and Liang Fu, Self-learning Monte Carlo method and cumulative update in fermion systems, *Phys. Rev. B* **95**, 241104 (R) (2017).
- [20] Johanna Vielhaben and Nils Strodthoff, Generative neural samplers for the quantum Heisenberg chain, *Phys. Rev. E* **103**, 063304 (2021).
- [21] Chuang Chen, Xiao Yan Xu, Junwei Liu, George Batrouni, Richard Scalettar, and Zi Yang Meng, Symmetry-enforced self-learning Monte Carlo method applied to the Holstein model, *Phys. Rev. B* **98**, 041102(R) (2018).
- [22] Giacomo Torlai and Roger G. Melko, Learning thermodynamics with Boltzmann machines, *Phys. Rev. B* **94**, 165134 (2016).
- [23] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [24] Pengfei Zhang, Huitao Shen, and Hui Zhai, Machine learning topological invariants with neural networks, *Phys. Rev. Lett.* **120**, 066401 (2018).
- [25] Ryan Abbott *et al.*, Normalizing flows for lattice gauge theory in arbitrary space-time dimension, [arXiv:2305.02402](https://arxiv.org/abs/2305.02402).
- [26] Andrea Cocco, Marco Letizia, Humberto Reyes-Gonzalez, and Riccardo Torre, On the curse of dimensionality for normalizing flows, [arXiv:2302.12024](https://arxiv.org/abs/2302.12024).
- [27] Ryan Abbott *et al.*, Sampling QCD field configurations with gauge-equivariant flow models, *Proc. Sci. LATTICE2022* (2023) 036.
- [28] Ryan Abbott *et al.*, Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions, *Phys. Rev. D* **106**, 074506 (2022).
- [29] Matteo Favoni, Andreas Ipp, and David I. Müller, Applications of lattice gauge equivariant neural networks, *EPJ Web Conf.* **274**, 09001 (2022).
- [30] M. S. Alberg, 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).
- [31] Michael S. Alberg, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Julian M. Urban, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan, Flow-based sampling for fermionic lattice field theories, *Phys. Rev. D* **104**, 114507 (2021).
- [32] Phiala E. Shanahan, Daniel Trewartha, and William Detmold, Machine learning action parameters in lattice quantum chromodynamics, *Phys. Rev. D* **97**, 094506 (2018).
- [33] Gurtej Kanwar, Michael S. Alberg, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan, Equivariant flow-based sampling for lattice gauge theory, *Phys. Rev. Lett.* **125**, 121601 (2020).
- [34] Michael S. Alberg, Denis Boyda, Daniel C. Hackett, Gurtej Kanwar, Kyle Cranmer, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan, Introduction to normalizing flows for lattice field theory, [arXiv:2101.08176](https://arxiv.org/abs/2101.08176).
- [35] Michael S. Alberg, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Fernando Romero-López, Phiala E. Shanahan, and Julian M. Urban, Flow-based sampling in the lattice Schwinger model at criticality, *Phys. Rev. D* **106**, 014514 (2022).
- [36] Daniel C. Hackett, Chung-Chun Hsieh, Michael S. Alberg, Denis Boyda, Jiunn-Wei Chen, Kai-Feng Chen, Kyle Cranmer, Gurtej Kanwar, and Phiala E. Shanahan, Flow-based sampling for multimodal distributions in lattice field theory, [arXiv:2107.00734](https://arxiv.org/abs/2107.00734).
- [37] Pim de Haan, Corrado Rainone, Miranda C. N. Cheng, and Roberto Bondesan, Scaling up machine learning for quantum field theory with equivariant continuous flows, [arXiv:2110.02673](https://arxiv.org/abs/2110.02673).
- [38] Michele Caselle, Elia Cellini, Alessandro Nada, and Marco Panero, Stochastic normalizing flows for lattice field theory, *Proc. Sci. LATTICE2022* (2023) 005.
- [39] Michele Caselle, Elia Cellini, Alessandro Nada, and Marco Panero, Stochastic normalizing flows as non-equilibrium transformations, *J. High Energy Phys.* **07** (2022) 015.
- [40] Michele Caselle, Elia Cellini, and Alessandro Nada, Sampling the lattice Nambu-Goto string using continuous normalizing flows, [arXiv:2307.01107](https://arxiv.org/abs/2307.01107).
- [41] P. R. Dinesh, Locality-constrained autoregressive cum conditional normalizing flow for lattice field theory simulations, [arXiv:2304.01798](https://arxiv.org/abs/2304.01798).
- [42] Ankur Singha, Dipankar Chakrabarti, and Vipul Arora, Conditional normalizing flow for Markov Chain Monte Carlo sampling in the critical region of lattice field theory, *Phys. Rev. D* **107**, 014512 (2023).
- [43] Ankur Singha, Dipankar Chakrabarti, and Vipul Arora, Generative learning for the problem of critical slowing down in lattice Gross-Neveu model, *SciPost Phys. Core* **5**, 052 (2022).
- [44] Jan M. Pawłowski and Julian M. Urban, Flow-based density of states for complex actions, *Phys. Rev. D* **108**, 054511 (2023).
- [45] Danilo Jimenez Rezende, George Papamakarios, Sébastien Racanière, Michael S. Alberg, Gurtej Kanwar, Phiala E. Shanahan, and Kyle Cranmer, Normalizing flows on tori and spheres, [arXiv:2002.02428](https://arxiv.org/abs/2002.02428).