

RECEIVED: October 23, 2023

REVISED: March 29, 2024

ACCEPTED: April 2, 2024

PUBLISHED: May 7, 2024

# Diffusion models as stochastic quantization in lattice field theory

L. Wang,<sup>a,b,c</sup> G. Aarts<sup>d,e</sup> and K. Zhou<sup>id f,c,\*</sup>

<sup>a</sup>*Interdisciplinary Theoretical and Mathematical Sciences Program (iTHEMS), RIKEN, Saitama 351-0198, Wako, Japan*

<sup>b</sup>*Shanghai Research Center for Theoretical Nuclear Physics, National Natural Science Foundation of China and Fudan University, Shanghai 200438, China*

<sup>c</sup>*Frankfurt Institute for Advanced Studies, Ruth Moufang Strasse 1, D-60438, Frankfurt am Main, Germany*

<sup>d</sup>*Department of Physics, Swansea University, SA2 8PP, Swansea, U.K.*

<sup>e</sup>*European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT\*) & Fondazione Bruno Kessler, Strada delle Tabarelle 286, 38123 Villazzano (TN), Italy*

<sup>f</sup>*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen),*

*2001 Longxiang Boulevard, Longgang District, Shenzhen 518172, China*

*E-mail: [lwang@fias.uni-frankfurt.de](mailto:lwang@fias.uni-frankfurt.de), [g.aarts@swansea.ac.uk](mailto:g.aarts@swansea.ac.uk), [zhoukai@cuhk.edu.cn](mailto:zhoukai@cuhk.edu.cn)*

**ABSTRACT:** In this work, we establish a direct connection between generative diffusion models (DMs) and stochastic quantization (SQ). The DM is realized by approximating the reversal of a stochastic process dictated by the Langevin equation, generating samples from a prior distribution to effectively mimic the target distribution. Using numerical simulations, we demonstrate that the DM can serve as a global sampler for generating quantum lattice field configurations in two-dimensional  $\phi^4$  theory. We demonstrate that DMs can notably reduce autocorrelation times in the Markov chain, especially in the critical region where standard Markov Chain Monte-Carlo (MCMC) algorithms experience critical slowing down. The findings can potentially inspire further advancements in lattice field theory simulations, in particular in cases where it is expensive to generate large ensembles.

**KEYWORDS:** Algorithms and Theoretical Developments, Lattice Quantum Field Theory, Non-Perturbative Renormalization, Stochastic Processes

ARXIV EPRINT: [2309.17082](https://arxiv.org/abs/2309.17082)

---

\*Corresponding author.

---

**Contents**

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Stochastic quantization</b>	<b>3</b>
2.1	Fokker-Planck equation and observables	3
2.2	Langevin simulations	5
<b>3</b>	<b>Diffusion models</b>	<b>6</b>
3.1	Denoising model	6
3.2	Diffusion model as stochastic quantization	8
3.3	Building effective actions	9
3.4	Effective action in a toy model	10
<b>4</b>	<b>Diffusion models for lattice <math>\phi^4</math> scalar fields</b>	<b>12</b>
4.1	Configuration generation and model setups	13
4.2	Broken phase	14
4.3	Symmetric phase	14
4.4	Autocorrelation time	17
<b>5</b>	<b>Conclusion and outlook</b>	<b>18</b>
<b>A</b>	<b>U-Net architecture</b>	<b>20</b>
<b>B</b>	<b>Skilling-Hutchinson estimator</b>	<b>21</b>
<b>C</b>	<b>Statistics examination</b>	<b>21</b>
<b>D</b>	<b>Acceptance rate</b>	<b>23</b>

---

**1 Introduction**

To obtain physical observables in lattice field theory, it is essential to approximate the path integral as a sum over field configurations. Monte Carlo methods rely on random sampling from the probability distribution, determined by the action of the system [1]. These methods can have high computational costs, particularly in the vicinity of a critical point in parameter space [2]. Generative models, as a class of machine learning (ML) algorithms, can be trained to generate new data that follow the same distribution as a given data set, representing the underlying target distribution [3]. An important potential application of generative models in lattice QCD is therefore to improve the efficiency of Monte Carlo simulations [4].

Depending on the way the likelihood is estimated, roughly two main categories of generative models are being developed. In implicit maximum likelihood estimation (MLE), Generative Adversarial Networks (GANs) can learn to generate new configurations via a min-max game using training on a dataset prepared using, e.g., a standard lattice simulation. A well-trained GAN can subsequently generate new samples, which should follow the same

statistical distribution as the original dataset, but with reduced computational cost. Additionally, GANs can be used to improve the interpretability of lattice QCD simulations by generating visualizations of the system’s behavior [5, 6]. In explicit MLE, flow-based models have recently been proposed to improve the efficiency of lattice simulations [7–20]. Flow-based models are able to directly approach the underlying distribution of the quantum field without preparing training data. Although the method may encounter “mode-collapse” and scalability problems [21, 22], it allows for the generation of independent new configurations that can be used in MC simulations, again reducing the computational cost. The current progress in applying generative models also includes designing novel neural network architectures for specific field theories, e.g., autoregressive networks [23, 24] and gauge-equivariant neural networks [8, 18, 24–27]. These approaches suggest that generative models have the potential to advance the capabilities of lattice QCD simulations. A comprehensive review can be found in refs. [28, 29].

Recently, the ML community has developed a new class of deep generative models known as Diffusion Models (DMs) [30]. Impressive success has been obtained in generating high-quality images via stochastic denoising processes, as showcased in prominent applications [31] such as DALL-E 2 [32] and Stable Diffusion [33]. As a prevalent class of generative models encoded in a stochastic process, DMs are grounded in Markov chains, and their implementation involves the utilization of variational inference methods. For an adoption in high-energy physics experiments, see, e.g., refs. [34, 35].

In numerical lattice field theory, there is an approach alternative to Monte Carlo methods to generate field configurations which relies on solving a stochastic differential equation, i.e., a Langevin equation. Its origin is stochastic quantization (SQ) as proposed by Parisi and Wu [36]. The basic idea of SQ is to view the quantum system as the thermal equilibrium limit of a hypothetical stochastic process with respect to a fictitious time. For gauge fields, it does not require gauge fixing [36]. Comprehensive reviews include refs. [37, 38]. For theories with a sign or complex weight problem, such as Quantum Chromodynamics at nonzero baryon density or theories in real (rather than Euclidean) time, the Langevin process can be extended to complex Langevin dynamics [39], which evades the sign problem in certain cases [40–44], see also the reviews [45–48]. Although the method has been put on firmer theoretical grounds [49–53], convergence and boundary condition issues continue to hinder obtaining a fully satisfying approach. Recent work to address this includes refs. [54–57].

In this work, we first point out the correspondence between generative DMs from the ML community and SQ in quantum field theory. Then, we explore using DMs to generate lattice field configurations. Specifically, we demonstrate for a two-dimensional  $\phi^4$  lattice field theory that, combined with a hybrid Monte Carlo (HMC) algorithm, DMs can serve as an efficient global sampler along a Markov chain, leading to shorter autocorrelation times. We provide self-consistency checks and numerical evidence that this approach is able to capture the SQ perspective-induced stochastic dynamics of quantum field theories.

The paper is organized as follows. In the subsequent section, we provide a brief overview of the concept of SQ. In section 3, we present details of DMs. It is argued that the denoising process of DMs can explicitly serve as a SQ procedure. In section 4, we explore the application of DMs to a two-dimensional scalar  $\phi^4$  field theory. Additionally, we demonstrate that DMs

can accurately estimate the physical action in an unsupervised manner using the probabilistic flow approach. In section 5, we summarize the correspondence between DMs and SQ, and outline how this novel method can motivate further research, in particular for theories where it is expensive to sample from the prior distribution using standard approaches. Appendix A contains information on the U-Net architecture employed, while appendix B contains a brief discussion on the Skilling-Hutchinson estimator. A study of the dependence of observables and the acceptance rate on various system and DM parameters can be found in appendix C and D respectively.

## 2 Stochastic quantization

In Euclidean quantum field theory, as an alternative quantization scheme, one may invoke stochastic quantization (SQ) [36]. We only consider real actions here, for which SQ is well understood and convergence is guaranteed [37]. Starting from a generic Euclidean path integral with Euclidean action  $S_E$ ,

$$Z = \int D\phi e^{-S_E}. \tag{2.1}$$

SQ introduces a *fictitious* time,  $\tau$ , for the field,  $\phi$ , whose evolution is described by Langevin dynamics,

$$\frac{\partial\phi(x, \tau)}{\partial\tau} = -\frac{\delta S_E[\phi]}{\delta\phi(x, \tau)} + \eta(x, \tau), \tag{2.2}$$

where the noise term routinely satisfies,

$$\langle\eta(x, \tau)\rangle = 0, \quad \langle\eta(x, \tau)\eta(x', \tau')\rangle = 2\alpha\delta(x - x')\delta(\tau - \tau'), \tag{2.3}$$

with  $\alpha$  being the diffusion constant. Denoting the distribution as

$$p(\phi) = \frac{e^{-S_E(\phi)}}{Z}, \tag{2.4}$$

the drift term in the Langevin equation can be written as

$$K(\phi) \equiv -\frac{\delta S_E}{\delta\phi(x)} = \frac{\delta \log p(\phi)}{\delta\phi(x)}. \tag{2.5}$$

In the long-time limit ( $\tau \rightarrow \infty$ ) the system reaches an equilibrium state of the dynamics [38] if the drift term has a damping effect, which is best analysed using the Fokker-Planck equation (see below). In addition, introducing the stochasticity is only one strategy for the quantization. The *fictitious* time itself can serve as a clue to develop microcanonical quantization, which describes a deterministic evolution with an ordinary differential equation [58]. We will revisit this idea and its connection with DMs in section 3.3.

### 2.1 Fokker-Planck equation and observables

The Fokker-Planck equation [59] describes the time evolution of the probability distribution of field configurations, subject to a stochastic process, such as the Langevin equation. Also known

as the Kolmogorov forward equation, in the context of SQ it describes the time evolution of the probability density function  $P[\phi, \tau]$  for field configurations  $\phi$  at time  $\tau$  [37, 60], as follows,

$$\frac{\partial P[\phi, \tau]}{\partial \tau} = \int d^n x \left\{ \frac{\delta}{\delta \phi} \left( \alpha \frac{\delta}{\delta \phi} + \frac{\delta S_E}{\delta \phi} \right) \right\} P[\phi, \tau]. \quad (2.6)$$

Here  $n$  is the dimension of the Euclidean spacetime on which the field theory is formulated. Given this probability distribution function, the expectation value of an observable  $\mathcal{O}[\phi]$  is then given by [37]

$$\langle \mathcal{O}[\phi] \rangle_\tau = \int D\phi \mathcal{O}[\phi] P[\phi, \tau], \quad (2.7)$$

where the left-hand side is understood as an average over many stochastic processes and on the right-hand side the integral is taken over all field configurations  $\phi$ .

In the long-time limit, the probability distribution  $P[\phi, t]$  approaches an equilibrium distribution, which is the stationary solution of the Fokker-Planck equation (2.6),

$$P_{\text{eq}}[\phi] \propto e^{-\frac{1}{\alpha} S_E[\phi]}. \quad (2.8)$$

If  $\alpha$  is set (formally) as  $\hbar$  and the equilibrium distribution is normalised, expectation values of observables are indeed observables in a quantum theory, with

$$\langle \mathcal{O}[\phi] \rangle_{\tau \rightarrow \infty} = \frac{\int D\phi \mathcal{O}(\phi) e^{-\frac{1}{\hbar} S_E(\phi)}}{\int D\phi e^{-\frac{1}{\hbar} S_E(\phi)}} = \langle \mathcal{O}[\phi] \rangle_{\text{quantum}}, \quad (2.9)$$

demonstrating that a stochastic process according to eq. (2.2) can be used to quantize a system.

Note that convergence to the stationary solution (2.8) follows from properties of the Fokker-Planck Hamiltonian associated to the Fokker-Planck equation (2.6) [37]. In short, writing the time-dependent solution  $P[\phi, t]$  as

$$P[\phi, t] = e^{-\frac{1}{2\alpha} S_E[\phi]} Q[\phi, t], \quad (2.10)$$

it is easy to show that  $Q[\phi, t]$  satisfies the equation

$$\frac{\partial Q[\phi, \tau]}{\partial \tau} = -H_{\text{FP}} Q[\phi, \tau], \quad (2.11)$$

where the Fokker-Planck Hamiltonian reads

$$H_{\text{FP}} = \int d^n x L^\dagger L, \quad (2.12)$$


with

$$L = \sqrt{\alpha} \frac{\delta}{\delta \phi} + \frac{1}{2\sqrt{\alpha}} \frac{\delta S_E}{\delta \phi}, \quad L^\dagger = -\sqrt{\alpha} \frac{\delta}{\delta \phi} + \frac{1}{2\sqrt{\alpha}} \frac{\delta S_E}{\delta \phi}. \quad (2.13)$$

The Fokker-Planck Hamiltonian is hence bounded from below by zero. Provided its spectrum is gapped, time-dependent components in eq. (2.11) will decay exponentially and only the solution  $H_{\text{FP}} Q[\phi] = 0$  survives. This solution is given by

$$LQ[\phi] = 0 \quad \Rightarrow \quad Q[\phi] \propto e^{-\frac{1}{2\alpha} S_E[\phi]}, \quad (2.14)$$

which, when combined with eq. (2.10), results again in the equilibrium solution (2.8), completing the standard proof of convergence.



**Figure 1.** Sketch of a Langevin simulation. Each line represents a different trajectory, evolving from an initial state at  $\tau = 0$  to one at a later time,  $T$ . The initial configurations are sampled from a prior but naive distribution,  $\phi \sim P[\phi, \tau = 0]$ . After the simulation, the final configurations follow the equilibrated distribution,  $\phi \sim P[\phi, \tau = T] \sim P_{\text{target}}[\phi]$ .

## 2.2 Langevin simulations

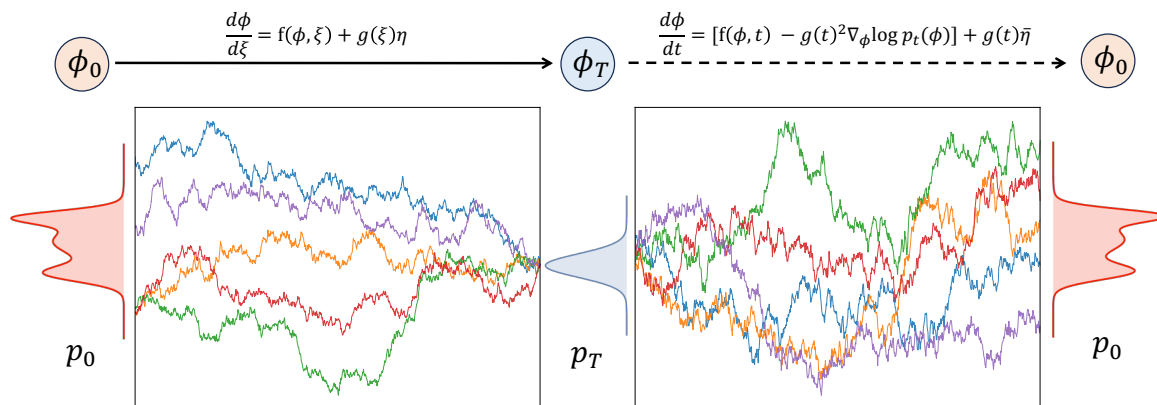
In most cases, the Fokker-Planck equation cannot be solved analytically. Instead, one solves the stochastic process numerically by discretizing the Langevin equation eq. (2.2). Using Itô calculus, the simplest discretized Langevin equation, using time step  $\Delta\tau$ , is

$$\phi(x, \tau_{n+1}) = \phi(x, \tau_n) + K[\phi(x, \tau_n)]\Delta\tau + \sqrt{\Delta\tau}\eta(x, \tau_n), \tag{2.15}$$

with  $\langle \eta(x, \tau_n)\eta(x', \tau_{n'}) \rangle = 2\alpha\delta(x - x')\delta_{nn'}$ .

A typical calculation includes, firstly, an initial condition to prepare field configurations  $\{\phi(x, \tau = 0)\}$ ; secondly, the thermalization procedure, i.e. evolve the system to reach equilibrium by iterating through a sufficient number of time steps with an appropriate time step,  $\Delta\tau$ ; thirdly, after thermalization, record the values of the physical observable  $\mathcal{O}[\phi]$  for each configuration in an ensemble, e.g. after a time  $T$ . This is repeated for numerous configurations to improve the statistical accuracy of the ensemble average. Figure 1 demonstrates one such typical Langevin process for a one-dimensional case. The target distribution is encoded via the drift, cf. eq. (2.5).

This procedure provides an estimate of the quantum expectation value for the observables  $\mathcal{O}[\phi]$ . Since the simple discretization scheme above results in finite-stepsize errors linear in  $\Delta\tau$ , one has to repeat the analysis to extrapolate to vanishing stepsize. The Langevin simulation provides an alternative to Monte Carlo approaches and has been used in simulations of lattice QCD [61]. A characteristic is that no accept/reject step is used, which makes the algorithm not exact (as demonstrated by e.g., the finite-stepsize corrections). This can be remedied by introducing such an accept/reject step, which results in more sophisticated algorithms that outperform the basic Langevin algorithm and are hence favoured (such as Hybrid Monte Carlo). The absence of an accept/reject step is turned into a feature for theories with a complex-valued Boltzmann weight, for which importance sampling is not



**Figure 2.** A sketch of the forward diffusion process (left panel) and the reverse denoising process (right panel). The two stochastic processes are described by two stochastic differential equations. The target distribution is typically unknown but learnt from the training data.

possible. The application of SQ to these theories is commonly referred to as complex Langevin dynamics [39], which is still an active area of research [40–57], but not further explored here.

### 3 Diffusion models

Turning now to deep learning, DMs denote a class of deep generative algorithms that simulate a stochastic process to generate samples following a desired target distribution [62]. The fundamental concepts underlying DMs involve modeling the data distribution through a continuous-time diffusion process or a discrete-time Markov chain, wherein noise is incrementally added to or removed from the data. This generative framework has gained prominence in recent works, such as score-based models (SBMs) [63, 64] and denoising diffusion probabilistic models (DDPMs) [65]; for a comprehensive review, see ref. [30]. In this section, we will first introduce the denoising perspective of DMs and then point out the connections of DMs to SQ.

#### 3.1 Denoising model

In the context of DMs, the denoising model can be perceived as a mapping that reconstructs the original data from its noisy variant, which is attained through the application of a diffusion process. Similar to the Langevin dynamics introduced above, the first process (also called forward process) gradually introduces noise to the data, effectively “smoothing” the underlying (but typically unknown) probability distribution of data. Then the denoising model aims to learn its *inverse process*, i.e., denoising the data by eliminating the added noise. Upon training, the denoising model can be employed to generate samples following the data distribution by executing the reverse diffusion process only. This sampling process commences with random noise and iteratively applies the trained denoising model to “clean” it until convergence, resulting in a sample representative of the target data distribution.

In the limit of continuous time, the forward process follows a stochastic differential equation (SDE),

$$\frac{d\phi}{d\xi} = f(\phi, \xi) + g(\xi)\eta(\xi), \tag{3.1}$$

where  $\xi \in [0, T]$  is the Langevin time,  $\eta$  is the noise term, satisfying  $\langle \eta(\xi)\eta(\xi') \rangle = 2\alpha\delta(\xi - \xi')$  with  $\alpha \equiv 1/2$  as a convention in the machine learning community, and  $f(\phi, \xi)$  is the drift term. Note that the noise and drift term have the same size as the field  $\phi$ . Finally, the square of  $g(\xi)$  is the time-dependent scalar *diffusion coefficient*. The forward diffusion process can be modeled as the solution of such a generic SDE, provided the drift and diffusion coefficient are chosen appropriately.

As figure 2 demonstrates, the aim is to obtain a sample from the data distribution  $p_0$  by starting from a sample of the prior distribution  $p_T$  and reversing the above diffusion process. This is achieved by solving a reverse SDE, representing a diffusion process evolving backward in time, which reads [66]

$$\frac{d\phi}{dt} = \left[ f(\phi, t) - g^2(t)\nabla_{\phi} \log p_t(\phi) \right] + g(t)\bar{\eta}(t), \tag{3.2}$$

where the reverse time  $t \equiv T - \xi$  and  $\bar{\eta}$  is a noise term in the reverse time direction. Importantly, the drift term now contains two terms, including the gradient of the logarithm of  $p_t(\phi)$ , the probabilistic distribution at time  $t$  in the forward diffusion process. This reverse SDE can be computed once we have specified the drift term and diffusion coefficient of the forward SDE, and determined the gradient of  $\log p_t(\phi)$  for each  $t \in [0, T]$ .

The key to implementing a denoising model lies in accurately estimating the probability distribution,  $p_t(\phi)$ , under a well-designed forward diffusion process of eq. (3.1). In score-based models [63, 64],  $\nabla_{\phi} \log p_t(\phi)$  is named as the *score* of each sample. Routinely, one can introduce a perturbation kernel for each forward diffusion step<sup>1</sup> as  $p_{\xi}(\phi_{\xi}|\phi_0) \equiv \mathcal{N}(\phi_{\xi}; \phi_0, \sigma_{\xi}^2\mathbf{I})$ , where  $\phi_0 \sim p_0$ . Thus, the *score* can be computed from

$$p_{\xi}(\phi_{\xi}) = \int p_{\xi}(\phi_{\xi}|\phi_0)p_0(\phi_0)d\phi_0. \tag{3.3}$$

A neural network  $s_{\theta}(\phi, t)$  can be trained to approximate the *score* conveniently, which is named as *score matching* [67, 68]. The parameters of the neural network are optimized with the objective,

$$\mathcal{L}_{\theta} = \sum_{i=1}^N \sigma_i^2 \mathbb{E}_{p_0(\phi_0)} \mathbb{E}_{p_i(\phi_i|\phi_0)} \left[ \|s_{\theta}(\phi_i, \xi) - \nabla_{\phi_i} \log p_i(\phi_i|\phi_0)\|_2^2 \right], \tag{3.4}$$

where the interval  $T$  has been divided into  $N$  segments with index  $i$  (further details are given below). The optimal score-based model  $s_{\hat{\theta}}(\phi, t)$  with  $\hat{\theta} = \arg \min_{\theta} \mathcal{L}_{\theta}$  should match the score at every time-step. Sampling turns to solving the reverse SDE shown in eq. (3.2), with the **score** replaced with  $s_{\hat{\theta}}(\phi, t)$ ,

$$\frac{d\phi}{dt} = \left[ f(\phi, t) - g^2(t)s_{\hat{\theta}}(\phi, t) \right] + g(t)\bar{\eta}(t). \tag{3.5}$$

For a concrete and concise example, by putting the forward drift term  $f(\phi, \xi)$  to zero and simplifying the diffusion coefficient, see, e.g., ref. [64], as  $g(\xi) \equiv \sigma^{\xi}$ , one gets a *variation*

---

<sup>1</sup>By adding to samples a sequence of Gaussian noise with scales,  $\sigma_{\min} = \sigma_1 < \sigma_2 < \dots < \sigma_N = \sigma_{\max}$ , so the starting point is  $p_1(\phi_1) \simeq p_0(\phi_0)$  when  $\sigma_{\min}$  is small enough, and the end point is  $p_N(\phi) \simeq \mathcal{N}(\phi; \mathbf{0}, \sigma_{\max}^2\mathbf{I})$  when  $\sigma_{\max}$  is large enough. Hereafter, we abbreviate  $\phi(\xi)$  as  $\phi_{\xi}$ , or  $\phi_i$  for the discretised case.



*expanding* forward process whose transition kernel remains Gaussian,

$$p_\xi(\phi_\xi|\phi_0) = \mathcal{N}\left(\phi_\xi; \phi_0, \frac{1}{2\log\sigma}(\sigma^{2\xi} - 1)\mathbf{I}\right). \quad (3.6)$$

We note that both eqs. (3.1) and (3.2) are Langevin equations.<sup>2</sup> In the context of the reverse SDE, given the gradients  $\nabla_\phi \log p(\phi, t)$ , it is convenient to generate samples from a prior distribution in a stochastic evolution process, as introduced in section 2.2.

### 3.2 Diffusion model as stochastic quantization

The reverse SDEs of DMs are mathematically related to the Langevin dynamics. To explore the connection of DMs to stochastic quantization in more detail, we choose the *variance expanding* picture of the DMs. Its denoising process transfers  $\phi_T$  back to  $\phi_0$  with the same marginal densities as the forward process, but evolving according to

$$\frac{d\phi}{dt} = -g(t)^2 \nabla_\phi \log p_t(\phi) + g(t)\bar{\eta}(t), \quad (3.7)$$

where time  $t$  flows backward from  $T$  to 0. With the redefinition of time,  $\tau \equiv T - t$  ( $d\tau \equiv -dt$ ), and denoting  $g_\tau = g(T - \tau)$ ,  $q_\tau(\phi) = p_{T-\tau}(\phi)$  of the forward process, the SDE now reads

$$\frac{d\phi}{d\tau} = g_\tau^2 \nabla_\phi \log q_\tau(\phi) + g_\tau \bar{\eta}(\tau) \quad (3.8)$$

and in discretised form

$$\phi(\tau_{n+1}) = \phi(\tau_n) + g_{\tau_n}^2 \nabla_\phi \log q_{\tau_n}[\phi(\tau_n)]\Delta\tau + g_{\tau_n} \sqrt{\Delta\tau} \bar{\eta}(\tau_n), \quad (3.9)$$

where the noise term  $\langle \bar{\eta}(\tau_n) \rangle = 0$ ,  $\langle \bar{\eta}(\tau_n) \bar{\eta}(\tau_{n'}) \rangle = 2\bar{\alpha} \delta_{nn'}$ , with  $\bar{\alpha} \equiv \hbar = 1$  as a convention in quantum field theories. Because the mean value of the noise term remains zero, the sign in front of the noise term is unrelated to the stochastic process. We note that  $g^2(\tau)$  is referred to as a kernel, rescaling both the drift term and the variance of the noise term [37]. Its effect can be absorbed by rescaling time with  $g^2(\tau)$ , or equivalently absorb it in the time step,  $g_\tau^2 \Delta\tau$ .

The derivation of the corresponding Fokker-Planck is straightforward and one finds

$$\frac{\partial p_\tau(\phi)}{\partial \tau} = \int d^n x \left\{ g_\tau^2 \frac{\delta}{\delta \phi} \left( \bar{\alpha} \frac{\delta}{\delta \phi} + \nabla_\phi S_{\text{DM}} \right) \right\} p_\tau(\phi), \quad (3.10)$$

where

$$\nabla_\phi S_{\text{DM}} \equiv -\nabla_\phi \log q_\tau(\phi) \quad (3.11)$$

is the score trained on the forward diffusion process, and  $p_\tau(\phi)$  is the probability density of  $\phi_\tau$  in the reverse diffusion process. One can immediately derive the following equilibrium condition, locally in time,

$$p_{\text{eq}}(\phi) \propto e^{-S_{\text{DM}}/\bar{\alpha}}, \quad (3.12)$$

---

<sup>2</sup>This approach has been utilized in Bayesian learning as stochastic gradient Langevin dynamics [69]. In contrast to standard stochastic gradient-descent (SGD) algorithms, stochastic gradient Langevin dynamics introduces stochasticity into the parameter updates, thereby avoiding collapses into local minima.

where we recall that  $\bar{\alpha} = 1$  and that  $S_{\text{DM}}$  is an effective action that will be learned in DMs. When the noise scales are on par with the quantum scale, that is,  $O(\bar{\alpha}) \sim O(\hbar)$ , eq. (3.12) encompasses the quantum fluctuations exhibited in eq. (2.8). In the  $\tau \rightarrow T$  limit,  $p_{\tau=T}(\phi) \rightarrow P[\phi, T]$ , where  $P[\phi, T]$  is nothing but the time evolution of the PDF in SQ. This implies that the “equilibrium state” of SQ can be achieved by denoising from a naive distribution. Concurrently, sampling from a DM is equivalent to optimizing a stochastic trajectory to approach the “equilibrium state” in eq. (2.8) given a naive distribution from the SQ perspective.

However, reverting the noise configuration back to the physics configuration is a highly intricate process. In principle, one could solve the reverse SDE of eq. (3.8) to depict this denoising process, while computing the “time-dependent” drift term remains intractable. A particular type of deep neural network, known as the U-Net, is used here. The architecture of this network is discussed in more detail in the appendix A. The U-Net is tailored to parameterize the *score function*,  $\hat{\mathbf{s}}_{\theta}(\phi, \tau)$ , for estimating this time-dependent drift term,  $-\nabla_{\phi} \log q_{\tau}(\phi)$ . It is designed to accept two inputs: time and a time-dependent configuration within a trajectory, while its output is the same size as the input (field configuration).

### 3.3 Building effective actions

To derive effective actions for quantum fields in DMs requires the *probability flow* ODE formulation [64, 70], which provides a quantitative way to compute the negative log-likelihood of the field configuration. Assuming a differentiable one-to-one mapping  $\mathbf{h}$  that transforms a data sample  $\phi_0 \sim p_0$  to one from a prior distribution  $\mathbf{h}(\phi_0) = \phi_T \sim p_T$ , the likelihood of  $p_0(\phi_0)$  can be computed using the change-of-variable formula as

$$p_0(\phi_0) = p_T(\mathbf{h}(\phi_0)) |\det(J_{\mathbf{h}}(\phi_0))|, \tag{3.13}$$

where  $J_{\mathbf{h}}$  denotes the Jacobian of the mapping  $\mathbf{h}$ , and it is assumed that the prior distribution  $p_T$  is easy to evaluate. ODE trajectories also define a one-to-one mapping from  $\phi_0$  to  $\phi_T$ . In our case, for quantum field theories,<sup>3</sup>

$$\frac{d\phi}{dt} = \left[ f(\phi, t) - g(t)^2 \nabla_{\phi} \log p_t(\phi) \right]. \tag{3.14}$$

It aims to find the trajectory of the state variable  $\phi$  that has the same marginal probability density  $p_t(\phi)$  as the stochastic process described by the SDE, i.e., eq. (3.1). An instantaneous change-of-variable formula can be derived to connect the probability of  $p_0(\phi_0)$  and  $p_T(\phi_T)$ , as,

$$p_0(\phi_0) = \exp \left[ \int_0^T dt \nabla \cdot \tilde{f}(\phi_t, t) \right] p_T(\phi_T), \tag{3.15}$$

where the divergence is the trace of the Jacobian, with redefinition of  $\tilde{f}(\phi_t, t) \equiv f(\phi, t) - g(t)^2 \nabla_{\phi} \log p_t(\phi)$ . The practical computation of the trace can be found in appendix B by using the Skilling-Hutchinson estimator [71].

---

<sup>3</sup>The derivations can be found in ref. [59] and appendix D.1 of ref. [64]. The difference for the factors in front of  $g(t)^2$  is due to the convention of  $\bar{\alpha} = 1$  in quantum field theories.

From the perspective of probability current conservation, the stochastic process we derived in the previous section intrinsically connects with flow-based models, which have recently been widely explored in current lattice calculations [11, 28]. In a flow-based model, one can construct a flow mapping with a neural network whose parameters are  $\{\theta\}$ , which represents a bijective transformation between the physical distribution  $p(\phi)$  and a prior distribution, e.g.,  $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Although it is common to build the flow mapping in discrete steps, the transformation  $\tilde{f}$  which follows,

$$\log p(\phi) = \log p(\mathbf{z}) - \log \left| \det \frac{\partial \tilde{f}}{\partial \mathbf{z}} \right|, \tag{3.16}$$

can also be designed in continuous steps.<sup>4</sup> It refers to the continuous normalizing flow [18], in which the parameterized mapping,  $\tilde{f}_{\hat{\theta}}$ , is the solution to a neural ODE for time  $t \in [0, T]$ ,

$$\frac{d\phi(t)}{dt} = \tilde{f}_{\theta}(\phi(t), t), \tag{3.17}$$

with boundary conditions:  $\phi(0) \equiv \mathbf{z}$ ,  $\phi(T) \equiv \phi$ . In which, a neural network with parameters,  $\tilde{f}_{\theta}(\phi, t)$ , can be introduced to represent the dynamical kernel. The probabilistic flow follows,

$$\frac{d \log p(\phi(t))}{dt} = - \left( \nabla_{\phi} \cdot \tilde{f}_{\theta} \right) (\phi(t), t), \tag{3.18}$$

with boundary  $p(\phi(0)) \equiv p(\mathbf{z})$ . Comparing the mapping and ODEs in eqs. (3.13) and (3.14) with eqs. (3.16) and (3.17), one can immediately find that the flow mapping,  $\tilde{f}_{\theta}(\phi(t), t)$ , is equivalent to  $\tilde{f}(\phi_t, t)$  in eq. (3.15) for DMs.

### 3.4 Effective action in a toy model

From eqs. (3.8) and (3.10), one can find a “time-dependent” stochastic process. In the variance-expanding DM, the forward process of eq. (3.1) features a predetermined time-dependent diffusion noise denoted as  $\sigma^{T-\tau}$ , which ensures field configurations will eventually be reduced to a noise configuration. The effective action and the corresponding denoising process can be understood from a field deformation perspective: *the field is continuously deformed following a stochastic process from the noise distribution to the physical distribution.*

To demonstrate this in action, we introduce an over-simplified 0+0-dimensional field theory, i.e., a toy model which only has one degree of freedom,

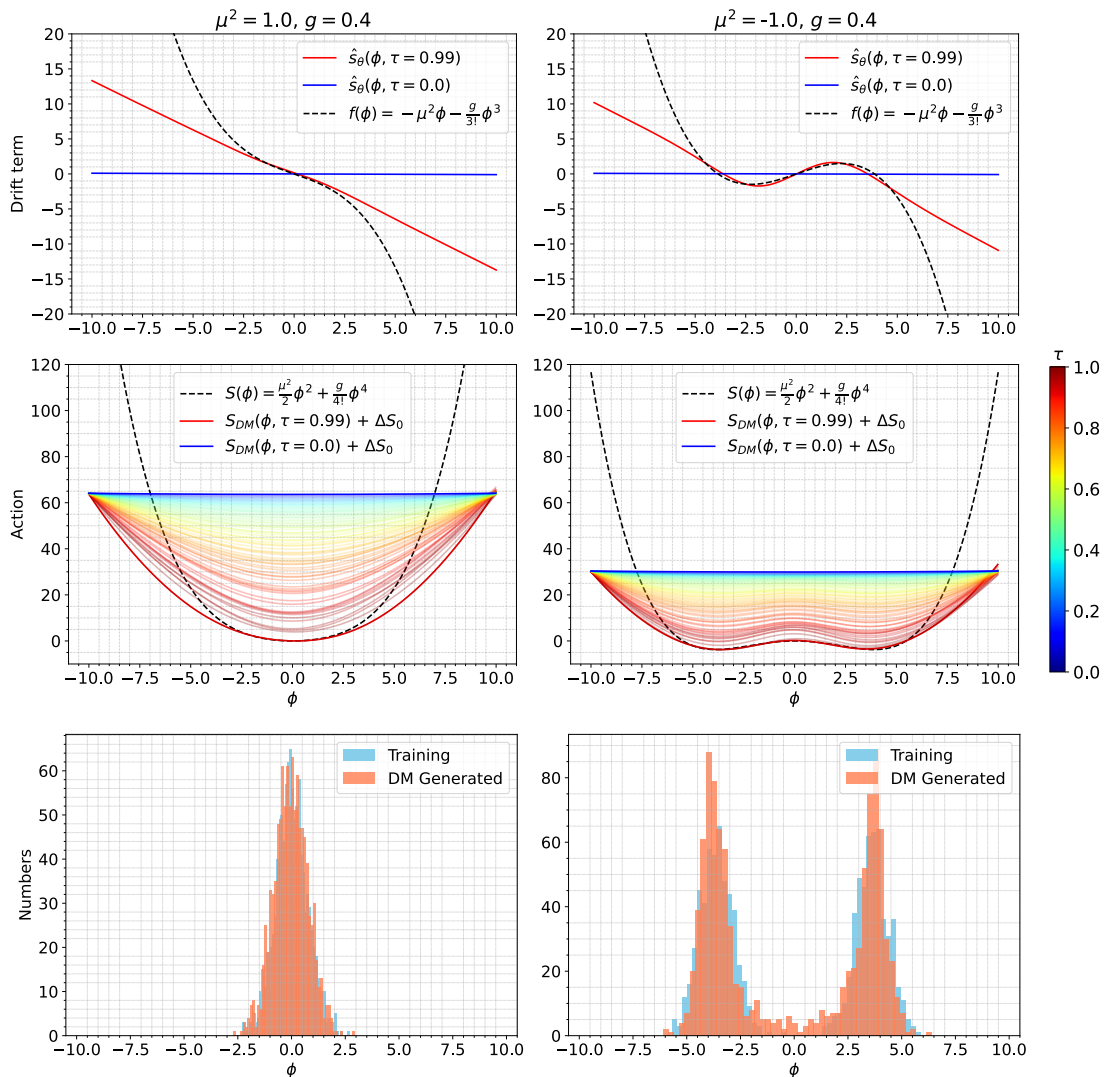
$$S(\phi) = \frac{\mu^2}{2} \phi^2 + \frac{g}{4!} \phi^4, \quad f(\phi) = - \frac{\partial S(\phi)}{\partial \phi} = -\mu^2 \phi - \frac{g}{3!} \phi^3, \tag{3.19}$$

where the physical action and drift term are determined by the parameters,  $\mu^2$  and  $g$ . We prepared 5120 configurations as training data-sets in two set-ups:  $\mu_1^2 = 1.0, g_1 = 0.4$  (single-well action), and  $\mu_2^2 = -1.0, g_2 = 0.4$  (double-well action). A one-to-one neural network with time-embedding is implemented to represent the score function  $\mathbf{s}_{\theta}(\phi, \tau)$ .

After 500 epochs of training, the learned score function  $\hat{\mathbf{s}}_{\theta}(\phi, \tau)$  is seen to approximate the drift term  $f_{\tau}(\phi)$  in the upper row of figure 3. Here the solid blue line represents the

---

<sup>4</sup>In fact, if one enforces  $\log p(\mathbf{z}) = \text{const.}$ , the transformation becomes the *trivializing* flow, see more details in refs. [12, 15, 19, 72].



**Figure 3.** Drift terms (upper row) and effective actions (middle row) learned by the DM as a function of  $\phi$  in both single-well (left column) and double-well (right column) actions, for various values of the time  $\tau$  during the stochastic process. The action is shifted by a constant  $\Delta S_0$ . The dashed lines indicate the exact values. 1024 samples generated using the exact and the DM are shown in the bottom row.

score function at the starting time ( $\tau = 0$ ), the rainbow-colored lines indicate different  $\tau > 0$  with a color-bar at rightest, while the solid red line indicates the score function at the end time ( $\tau = 0.99$ ) in a backward process of the well-trained DM. The learned effective action  $S_{DM}(\phi, \tau) = \int^\phi \hat{s}_\theta(\tilde{\phi}, \tau) d\tilde{\phi}$  is shown in the middle row of figure 3; it approximates the physical action as  $\tau$  increases. We have shifted  $S_{DM}$  with a constant  $\Delta S_0$ , which is the difference between  $\min[S(\phi)]$  and  $\min[S_{DM}(\phi, \tau)]$ . Generally, the learned drift terms and effective actions are accurate approximations in both the single and double-well cases, around  $\phi \sim 0$ . Samples generated with the trained DM are compared with samples from the underlying theory in the bottom row of figure 3.

The evolution trajectory of DM resembles an exact renormalizing group (RG) flow and its inversion. One intuitive understanding is that the gradually increasing noise scale (in

the forward process of the DM) in coordinate space implicitly corresponds to a gradually decreasing momentum scale  $\Lambda$  in an RG.<sup>5</sup> With  $g_{\tau \rightarrow T} = \sigma^{T-\tau} \rightarrow 1$ ,  $\Lambda \rightarrow \infty$ , the full quantum theory is recovered. Although some works have discussed the running noise scales from a colored noise perspective in a Langevin process [73, 74], the well-trained DM provides an exact RG flow in the *fictitious* time space,

$$\frac{\partial S_{\text{DM}}(\phi, \tau)}{\partial \tau} = \frac{\partial}{\partial \tau} \int^{\phi} \hat{s}_{\theta}(\tilde{\phi}, \tau) d\tilde{\phi}. \quad (3.20)$$

Because the probability density is positive-definite, the equation has a fixed point when  $\partial p_{\tau}(\phi)/\partial \tau = 0$ . The RG flow is the evolution trajectory of the system from the trivial distribution at  $\tau = 0$  to the equilibrium state at  $\tau = T$ , which has been shown in the middle panel of figure 3.

#### 4 Diffusion models for lattice $\phi^4$ scalar fields

Lattice scalar field theories have been studied extensively, with applications spanning a wide range of topics. Scalar fields are frequently employed to showcase the efficacy of numerical algorithms, including with machine learning. Here we will implement a DM to generate configurations for a two-dimensional scalar field theory with quartic interaction, both in the broken and the symmetric phase.

We consider a real scalar field in  $d$  *Euclidean* dimensions with the action,

$$S_E = \int d^d x \left( \frac{1}{2} \sum_{\mu=1}^d (\partial_{\mu} \phi_0)^2 + \frac{1}{2} m_0^2 \phi_0^2 + \frac{\lambda_0}{4!} \phi_0^4 \right), \quad (4.1)$$

where the subscript specifies bare quantities, including mass  $m_0$ , coupling  $\lambda_0$ , and field  $\phi_0(x)$ . Discretising the derivative on a lattice with lattice spacing  $a$  in all directions,  $\partial_{\mu} \phi_0(x) = [\phi_0(x + a\hat{\mu}) - \phi_0(x)]/a$ , with  $\hat{\mu}$  the unit vector along the  $\mu$ -direction, and redefining the field and parameters to yield dimensionless combinations, gives the lattice action (see e.g. ref. [75])

$$S_E = \sum_x \left[ -2\kappa \sum_{\mu=1}^d \phi(x)\phi(x + \hat{\mu}) + (1 - 2\lambda)\phi^2(x) + \lambda\phi^4(x) \right], \quad (4.2)$$

where  $\kappa$  is the hopping parameter, related to the bare mass parameter via

$$(am_0)^2 = \frac{1 - 2\lambda}{\kappa} - 2d, \quad (4.3)$$

---

<sup>5</sup>From eq. (3.6), one can sample to get the evolving field (for a given initial field  $\phi_0$ ) at any time during the forward process as  $\phi_{\tau}(\mathbf{x}) = \phi_0(\mathbf{x}) + \sqrt{\frac{\sigma^{2\tau}-1}{2 \log \sigma}} \epsilon(\mathbf{x})$  with  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  a Gaussian distributed random noise, which in momentum space is  $\phi_{\tau}(p) = \phi_0(p) + \sqrt{\frac{\sigma^{2\tau}-1}{2 \log \sigma}} \epsilon(p)$ . Considering the natural distribution of momentum modes for the physical fields, the above evolution will perturb (smear) higher momentum modes faster because of the gradually increasing noise level. This resembles somehow the functional renormalization group (FRG) procedure to progressively integrate out the high momentum modes. Conversely, in the reverse denoising process, the low momentum IR physics will be first generated, and only at a later time the UV details are filled by the denoising model, as seen in figure 4.

and  $\lambda = a^{4-d}\kappa^2\lambda_0/6$  denotes the dimensionless coupling constant describing field interactions. Both parameters are positive. The field has been rescaled according to  $a^{d/2-1}\phi_0 = (2\kappa)^{1/2}\phi$ .

In the case of  $d \geq 2$ , for each coupling  $\lambda$  the hopping parameter exhibits a critical value  $\kappa_c(\lambda)$ , at which a second-order phase transition occurs. This quantum phase transition corresponds to a spontaneous  $\mathbb{Z}_2$  symmetry broken above the critical point in continuum limit [76]. When the mass term becomes negative, the broken phase emerges classically. At the classical level, the critical value for the hopping parameter is given by

$$\kappa_c^{\text{cl}}(\lambda) = \frac{1 - 2\lambda}{2d}, \tag{4.4}$$

which is determined by the vanishing of the mass term, but quantum fluctuations will change the value [75]. As  $\kappa$  decreases across the critical value, the system transitions from the broken phase to the symmetric phase.


#### 4.1 Configuration generation and model setups

To implement the DM, we first prepare field configurations of the scalar  $\phi^4$ -theory in  $d = 2$  dimensions on a  $32 \times 32$  lattice, at fixed  $\lambda = 0.022$ , as defined above. To evaluate the performance of the model, we generated 5120 samples for training and 1024 samples for testing in the broken and symmetric phases at  $\kappa = 0.5$  and  $\kappa = 0.21$  respectively. The reason to choose these specific parameter values is that at  $\lambda = 0.022$  the phase transition occurs at approximately  $\kappa_c \simeq 0.239$ , as derived from eq. (4.4).


**Hybrid Monte Carlo set-up.** In order to prepare the dataset, we employ the Monte Carlo Markov Chain (MCMC) approach. The classical Metropolis-Hastings algorithm for local updating often suffers from long autocorrelation times, so we instead opt for the hybrid Monte Carlo (HMC) method. By incorporating classical Hamiltonian equations of motion [77], the HMC method improves the performance of the MCMC and allows larger steps to be taken while maintaining acceptable acceptance rates. In our setup, with the selected set of action parameters, we prepare configurations on a 2-dimensional lattice of size  $32 \times 32$ . For each Markov Chain in HMC, we have a “100-step burn-in loop” for thermalization, and after 64 additional steps, we select the final step of the 64-step update as a configuration. Here one step indicates one MC trajectory with a well-established HMC integrator. Each configuration in the training data-set is from one different HMC trajectory.

**Diffusion Model set-up.** We adopt the variance expanding DM as shown in eq. (3.6) with  $\sigma = 25$ . The diffusion time is set as  $T = 1$ . In the training procedure the stepsize is not fixed, but we randomly sample the time points in the interval  $[0, T]$ , with the number of time points equal to the batch size in each batch optimization. The batch size is 64, and the maximum number of epochs for training is 250 to ensure convergence. In the sampling procedure, we use the Euler-Maruyama approach to solve the reverse-time SDE with time stepsize  $T/N = 0.002$ . The U-Net architecture [78] is utilized to represent  $\mathbf{s}_\theta$  in eq. (3.4); more details can be found in appendix A. Unless otherwise indicated, we generated 1024 samples from the well-trained DM to demonstrate its performance.

We deployed the HMC and DM algorithms on different devices. For the chosen parameter values ( $\kappa = 0.21, \lambda = 0.022, 32 \times 32$  lattice), the HMC generates samples on CPU (Intel Xeon



**Figure 4.** The denoising process for generating four independent configurations from a well-trained DM.



**Figure 5.** Comparison of the distribution of the magnetization with 1024 samples generated from the well-trained DM and using HMC.

Gold 6226R @ 2.90GHz) with 1.57 configuration/s. The DM generates samples on GPU (Nvidia GeForce RTX 3090 @ 1.40Ghz) with 97.52 configuration/s.

## 4.2 Broken phase

We start in the broken phase. Here one expects large clusters, in which field configurations fluctuate around a positive or negative average value, usually referred to as the vacuum expectation value, order parameter or magnetisation. We demonstrate that the clustering behavior of field configurations can be successfully captured by the well-trained DM. Figure 4 visualizes the denoising process. Each row in the figure represents a different sample, and each column represents a different time point in the denoising process. The first column represents noise samples randomly drawn from the prior distribution, i.e.,  $p \approx \mathcal{N}\left(\phi; \mathbf{0}, \frac{1}{2\log\sigma}(\sigma^2 - 1)\mathbf{I}\right)$ , while the fifth column represents the generated samples obtained by denoising. The arrow at the bottom of figure 4 indicates the direction of time evolution in the stochastic differential equation, as per eq. (3.2).


The order parameter or magnetization is defined as

$$\langle M \rangle = \frac{1}{V} \left\langle \sum_x \phi(x) \right\rangle, \quad (4.5)$$

where  $V = N^d$  denotes the number of lattice sites. A comparison between HMC generated testing samples and DM generated samples is shown in figure 5, where we see consistency between the distributions containing 1024 samples each.

## 4.3 Symmetric phase

We now turn to the symmetric phase and take  $\kappa = 0.21$ , below but near the critical value. In the thermodynamic limit, the phase transition can be characterized by a divergence in the connected two-point susceptibility. This susceptibility is a measure of how a system



**Figure 6.** Comparison of the distribution of the magnetization in the symmetric phase from the test data-set (HMC) and from the well-trained DM. The number of independent configurations is 1024 in both cases.

data-set	$\langle M \rangle$	$\chi_2$	$U_L$
Training (HMC)	$0.0012 \pm 0.0007$	$2.5160 \pm 0.0457$	$0.1042 \pm 0.0367$
Testing (HMC)	$0.0018 \pm 0.0015$	$2.4463 \pm 0.1099$	$-0.0198 \pm 0.1035$
Generated (DM)	$0.0017 \pm 0.0015$	$2.4227 \pm 0.1035$	$0.0484 \pm 0.0959$

**Table 1.** Comparison of observables calculated on different datasets with lower error bounds determined using the statistical jackknife method. The sample size is comparable for the bottom two rows and a factor of 5 larger for the training set (top row).

responds to small perturbations and is defined as

$$\chi_2 = V \left( \langle M^2 \rangle - \langle M \rangle^2 \right). \tag{4.6}$$


In addition, the Binder cumulant [79] is a dimensionless quantity often used to identify phase transitions, defined as a ratio of moments of a distribution,

$$U_L = 1 - \frac{1}{3} \frac{\langle M^4 \rangle}{\langle M^2 \rangle^2}. \tag{4.7}$$

It is sensitive to the shape of the distribution, i.e., the kurtosis of the order parameter, and is particularly useful for detecting and analyzing phase transitions in finite-size systems.

In figure 6 we compare the distributions of the magnetization obtained independently from the trained DM and using HMC. Good consistency between the distributions is observed. To make this more quantitative we present in table 1 the magnetization  $\langle M \rangle$ , two-point susceptibility  $\chi_2$ , and the Binder cumulant  $U_L$  obtained using both the trained DM and with HMC. For the latter, expectation values are shown both for the training set and the testing set, with the former having a factor of five more configurations, which is reflected in the uncertainty. We conclude that the DM is capable of accurately capturing the statistical





**Figure 7.** Correlation between the effective action estimated from the trained DM using the probabilistic ODE and from the physical action on 1024 samples.

properties of the system, as reflected by these observables, and is in good agreement with HMC, a well-established method for sampling from the target distribution. A study of the dependence on the size of the training set and the generated data set can be found in the appendix C. We note that no accept-reject step has been applied in the DM.

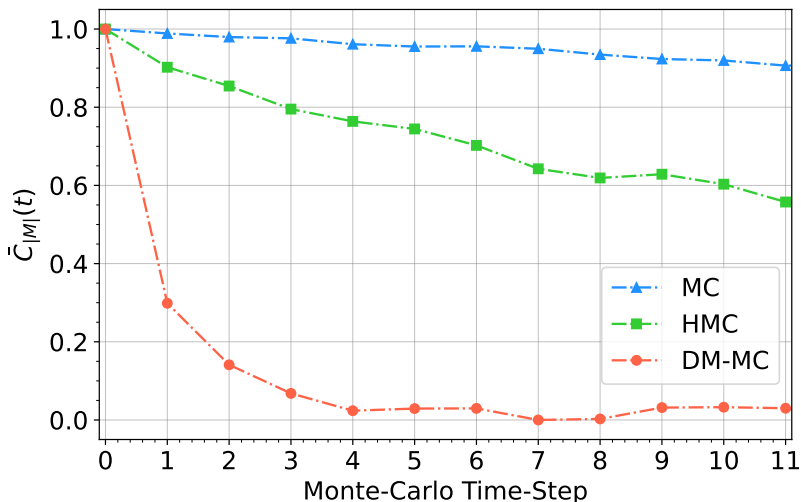
The likelihood computation method described in section 3.3 enables us to estimate the action for the field configurations generated by the trained DM. By comparing the action estimation from the trained DMs with the true action, see eq. (4.2), we can confirm whether the DM has captured the underlying physics correctly. For this purpose, we first generate an ensemble of field configurations using the trained DM. We then compute the likelihoods,  $q(\phi)$ , for these configurations using the probability flow ODE approach. By applying the change-of-variable formula, we are able to compute the action values for these field configurations, as

$$S_{\text{eff}} = -\log q(\phi), \tag{4.8}$$

where the constant term  $Z$  has been omitted. In figure 7 we compare the evaluated effective action with the one obtained directly from the physical action, i.e., eq. (4.2). The black dashed line means the  $x$ -axis is proportional to the  $y$ -axis, and the green sites are action values estimated from the probabilistic ODE (on the  $x$ -axis) and calculated from the physical action (on the  $y$ -axis).

The coefficient of determination,  $R^2$ , is a useful metric for quantitatively measuring how well the estimated action values, denoted as  $S_{\text{eff}}(\phi_i)$ , match the actual ones, denoted as  $S(\phi_i)$ . Here  $i$  indicates the index in sampled configurations. We use

$$R^2 = 1 - \frac{\sum_i [S(\phi_i) - S_{\text{eff}}(\phi_i)]^2}{\sum_i [S(\phi_i) - \bar{S}(\phi)]^2}, \tag{4.9}$$



**Figure 8.** Comparison of the normalized autocorrelation function of  $|M|$  for the Metropolis-Hastings MC algorithm, HMC, and DM-based MC, using 1024 samples at  $\kappa = 0.27$ ,  $\lambda = 0.22$  and  $L = 32$ .

where  $\bar{S}(\phi)$  is the average value for the actual action,

$$\bar{S}(\phi) = \frac{1}{N} \sum_{i=1}^N S(\phi_i). \tag{4.10}$$

It measures the proportion of variation in the dependent variable (in this case, the action) that is predictable from the independent variable (in this case, the DM-generated samples).  $R^2$  values close to 1 indicate that the DM provides an accurate estimate of the action, suggesting that the trained DM has effectively learned the underlying physics of the scalar field theory correctly. In figure 7, with the default dataset size, i.e., 5120 samples for training, the  $R^2$  value is 0.960 calculated on the 1024 generated configurations.

#### 4.4 Autocorrelation time

Critical slowing down [2] emerges when simulations approach a critical point and the correlation length of the system diverges. This should be explored at  $\kappa$  values close to the phase transition, which is located at  $\kappa_c \simeq 0.27$  (for  $L = 32, \lambda = 0.022$ ); see e.g. refs. [6, 74, 80].

Critical slowing down can be studied using the autocorrelation function of an observable  $O$ , defined as

$$C_O(t) = \langle (O_{t_0} - \langle O_{t_0} \rangle)(O_{t_0+t} - \langle O_{t_0+t} \rangle) \rangle = \langle O_{t_0} O_{t_0+t} \rangle - \langle O_{t_0} \rangle \langle O_{t_0+t} \rangle, \tag{4.11}$$

where  $t_0$  is the starting time and  $t$  denotes the time along the Markov chain. In an exponential fit,  $C_O(t) \sim \exp\{-t/t_c\}$ , the characteristic time  $t_c$  should scale as a power of the correlation length  $\xi$  around the critical point, i.e.,  $t_c \sim \xi^z$  (the correlation length  $\xi$  should not be confused with the time variable in the forward SDE). The normalized autocorrelation functions,  $\bar{C}_O(t) \equiv C_O(t)/C_O(0)$  can be used for a clearer comparison.

We investigate this for the “absolute value” of the magnetization, defined as

$$|M| = \frac{1}{V} \left\langle \sum_x |\phi(x)| \right\rangle, \tag{4.12}$$

and compare various updates. Because the likelihood of samples can be computed, as shown in the previous section, we can use the Metropolis-Hastings (MH) algorithm to construct the asymptotically exact Markov chain sampler with the well-trained DM. We refer to this as the DM-based Monte-Carlo method, following the convention of flow-based models. The acceptance criterion is corrected as

$$p_{\text{accept}}(\phi_{\text{proposal}} | \phi^{i-1}) = \min \left( 1, \frac{q(\phi^{i-1}) p(\phi_{\text{proposal}})}{p(\phi^{i-1}) q(\phi_{\text{proposal}})} \right), \quad (4.13)$$

where  $\phi^{i-1}$  is the previous configuration and  $\phi_{\text{proposal}}$  is sampled from the well-trained DM. Further, we can estimate the efficiency gain of the DM by comparing the behavior of the autocorrelation function  $C_{|M|}(t)$  with the one obtained using a Metropolis-Hastings Monte-Carlo (MC) and the HMC algorithm on 1024 configurations, as shown in figure 8. A more comprehensive discussion on the dependence of the acceptance rate on the parameters and the number of training epochs can be found in appendix D.

Beside, one can compute the integrated autocorrelation time [6, 12, 81] for a more quantitative comparison. It is defined as

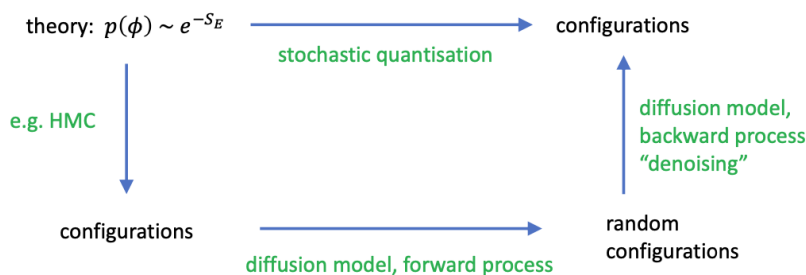
$$\tau_{O,\text{int}} = \frac{1}{2} + \frac{1}{C_O(0)} \sum_{t=1}^{t_{\text{max}}} C_O(t). \quad (4.14)$$

For the algorithms discussed, integrated autocorrelation times are  $\tau_{|M|,\text{int}}(\text{MC}) = 79.984$ ,  $\tau_{|M|,\text{int}}(\text{HMC}) = 41.354$ , and  $\tau_{|M|,\text{int}}(\text{DM-MC}) = 2.360$ , where we calculated the autocorrelation time with  $t_{\text{max}} = 100$  MC trajectories on 1024 configurations. To get a more precise estimation with uncertainties, one can also choose the automatic windowing procedure in the computation [82, 83]. These comparisons demonstrate that our proposed method has the potential to significantly suppress autocorrelations along a Markov chain.

## 5 Conclusion and outlook

In this work, we proposed a new method to generate quantum field configurations using diffusion models, popular in the ML community. We highlighted the connection with stochastic quantization (SQ), an approach to quantize field theories based on a stochastic process in a fictitious time direction. In SQ the drift term in the stochastic process is known, as it is derived from the probability distribution one wants to sample from, which is determined by the quantum theory under consideration. In DMs the drift term is not known but is learned from data in a forward stochastic process. After estimating the drift term, the trained DM is run in the opposite direction and independent configurations are created from noise, the “denoising” process. This latter part is akin to the dynamics in SQ, albeit with a learned and time-dependent drift term. In figure 9, we summarize the relation between stochastic quantisation and diffusion models schematically.

We have demonstrated the approach in a simple toy model and in a two-dimensional scalar field theory with a  $\phi^4$  interaction, both in the symmetric and in the broken phase. We have shown that the DM can serve as a global sampler to assist methods based on the Monte Carlo principle. We note that one can include an accept-reject step at the end of each




**Figure 9.** Schematic overview of the relation between stochastic quantisation and diffusion models in the case of lattice field theory.

trajectory. We have found in our numerical results that the well-trained DM can successfully reproduce configurations in both phases and that autocorrelation times are reduced along a Markov chain. This improvement in sampling efficiency can be beneficial for studying lattice field theories and should be analyzed further.

There are various directions to explore in the future. The forward (“smoothing”) and backward (“denoising”) processes are reminiscent of (inverse) renormalization group transformations and it would be interesting to make this precise. Because the action can be estimated relatively accurately, it is feasible to train a DM without preparing a training data-set, which needs to be investigated. The connection between DMs and SQ offers new perspectives. It is well known how to implement SQ for non-abelian gauge theories; hence it should be possible to formulate DMs for the non-abelian case. An intriguing direction is the following: imagine one has generated a set of configurations for a quantum field theory with fermions, such as QCD or the Schwinger model. Since the fermions are “integrated out”, the configurations are given in terms of the gauge fields only. By using these configurations as the starting point of a DM, one can learn an effective stochastic process which incorporates the effect of the fermions implicitly. This may be used to create additional configurations in a numerically “cheap” manner. An accept-reject step based on the underlying theory can be included, which would require the evaluation of the fermion determinant. Finally SQ can also be used for theories with a sign or complex action problem, and hence a non-positive definite distribution. This is commonly known as complex Langevin (CL) dynamics. It is known that CL may become less reliable along long trajectories in the stochastic process. Combining CL with DMs, one may consider fairly short trajectories, but subsequently use the resulting configurations to train the DM to generate additional configurations to increase statistics.

### Acknowledgments

We thank Profs. Tetsuo Hatsuda, Shuzhe Shi and Xu-Guang Huang for helpful discussions. We thank ECT\* and the ExtreMe Matter Institute EMMI at GSI, Darmstadt, for support during the ECT\*/EMMI workshop *Machine learning for lattice field theory and beyond* in June 2023 during the preparation of this paper. The work is supported by (i) the CUHK-Shenzhen university development fund under grant No. UDF01003041 and the BMBF funded KISS consortium (05D23RI1) in the ErUM-Data action plan (K.Z.), (ii) the AI grant of SAMSON AG, Frankfurt (K.Z. and L.W.), (iii) Xidian-FIAS International Joint Research



**Figure 10.** Sketch of the U-Net Architecture, where the gray lines with arrows indicate skip connections between encoding paths and decoding paths. In this study, the default shape of the input is set as  $(\cdot, 1, 32, 32)$ , where the first dimension is the batch size, the second indicates the channel, and the final two represent dimensionalities of the field configuration.

Center (L.W), (iv) STFC Consolidated Grant ST/T000813/1 (G.A.). K.Z. also thanks the donation of NVIDIA GPUs from NVIDIA Corporation. L.W. also thanks the Natural Science Foundation of Shanghai (Grant No.23JC1400200) for supporting the computation resources when visiting in Fudan University.

**Note added.** Recent related work introduces the stochastic process into the hybrid Monte-Carlo algorithm [84] and explores the correspondence between the exact renormalizing group (ERG) and DMs based upon the heat equation [85].

## A U-Net architecture

U-Net architectures are utilized to present  $s_\theta(\phi, t)$  with a time-embedding input. This architecture is widely used in tasks that require semantic segmentation due to its ability to capture both local and global information in an image. The model is initialized with a function that calculates the standard deviation of the perturbation kernel, a list of channels for feature maps at each resolution, and an embedding dimension for Gaussian random feature embeddings. The model consists of an **encoding path** (contracting path) and a **decoding path** (expansive path), which are characteristic of the U-Net architecture, as figure 10 shows.

The **encoding path** consists of four convolutional layers (conv1 to conv4), each followed by a dense layer (dense1 to dense4), and a group normalization layer (gnorm1 to gnorm4). The convolutional layers progressively increase the number of channels from 1 to 32, 64, 128, 256. The dense layers incorporate information from the time embedding into the feature maps, and the group normalization layers normalize the feature maps across the channels.

The **decoding path** consists of four transposed convolutional layers (tconv4 to tconv1), each followed by a dense layer (dense5 to dense7), and a group normalization layer (tgnorm4 to tgnorm2). The transposed convolutional layers progressively decrease the number of channels back to 1. The dense layers incorporate information from the time embedding into the feature maps, and the group normalization layers normalize the feature maps across the channels. The decoding path also includes skip connections from the encoding path, which help the model to better localize and learn representations with less loss of spatial information.

The model uses the Swish activation function, which is a smooth, non-monotonic function that has been found to work better than ReLU in deeper models. In the forward pass, the model first obtains the Gaussian random feature embedding for the input time  $t$ . It then passes the input  $\phi$  through the encoding path, where it applies the convolutional, dense, group normalization, and activation layers sequentially. It then passes the output of the encoding path through the decoding path, where it applies the transposed convolutional, dense, group normalization, and activation layers sequentially, while also adding the skip connections from the encoding path. Finally, it normalizes the output by the standard deviation of the perturbation kernel at time  $t$  and returns it.

## B Skilling-Hutchinson estimator

The divergence function in eq. (3.15) can be hard to evaluate in practice for general vector-valued functions  $f$ , but we can use an unbiased estimator, named Skilling-Hutchinson estimator [71], to approximate the trace. Let  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . The Skilling-Hutchinson estimator is based on the fact that

$$\nabla \cdot f(\phi) = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\epsilon^\top J_f(\phi) \epsilon]. \tag{B.1}$$

Therefore, we can simply sample a random vector  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , and then use  $\epsilon^\top J_f(\phi) \epsilon$  to estimate the divergence of  $f(\phi)$ . This estimator only requires computing the Jacobian-vector product  $J_f(\phi) \epsilon$ , which is typically efficient.

As a result, for our probability flow ODE, we can compute the (log) data likelihood with the following,

$$\log p_0(\phi(0)) = \log p_T(\phi(T)) - \int_0^T \frac{d[\sigma^2(t)]}{dt} \nabla \cdot s_\theta(\phi(t), t) dt. \tag{B.2}$$

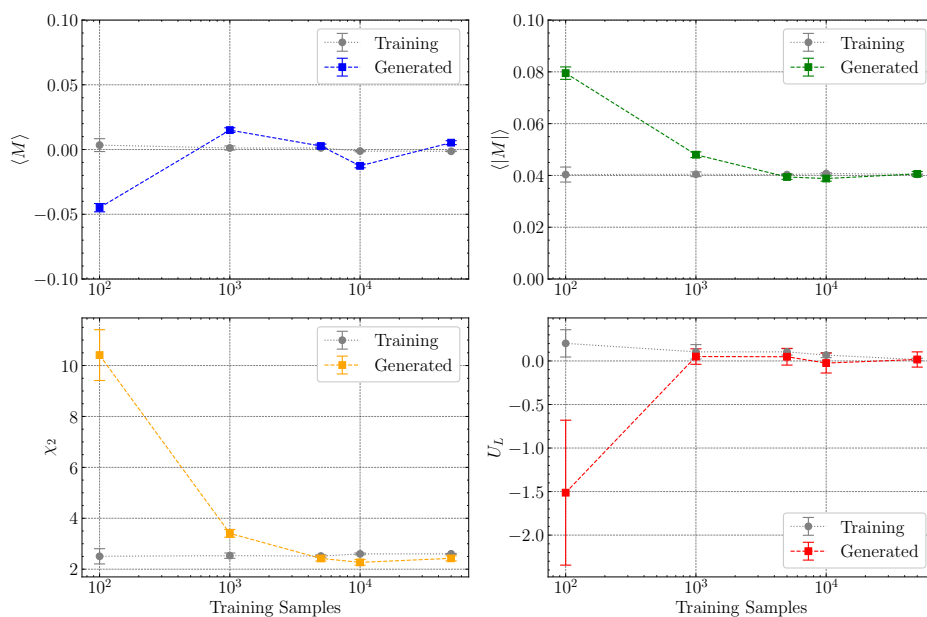
With the Skilling-Hutchinson estimator, we can compute the divergence via

$$\nabla \cdot s_\theta(\phi(t), t) = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\epsilon^\top J_{s_\theta}(\phi(t), t) \epsilon]. \tag{B.3}$$

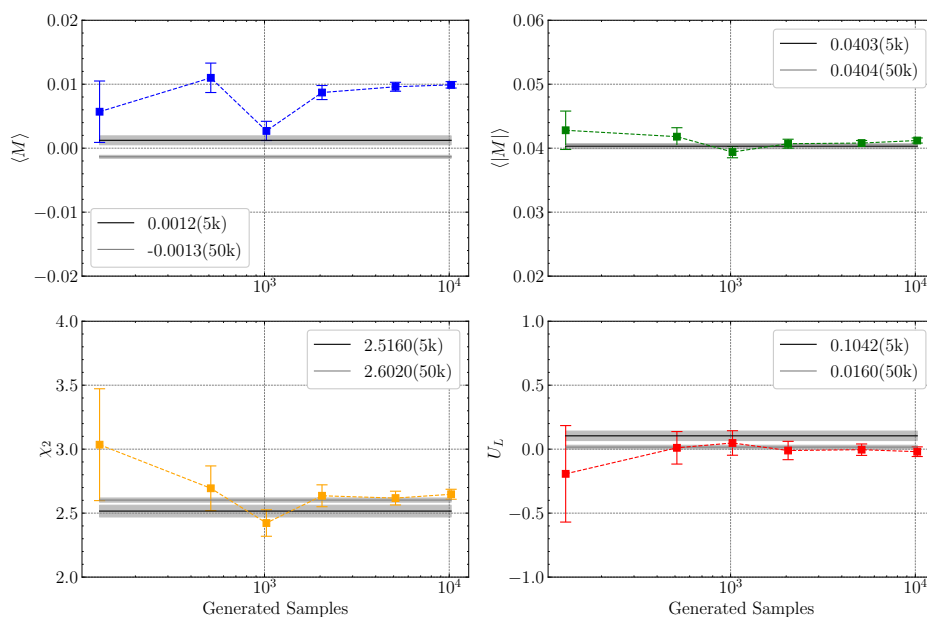
Subsequently, we can compute the integral using numerical integrators. This provides us with an unbiased estimate of the true likelihood of the data, and we can make this estimate increasingly accurate by running the calculation multiple times and taking the average. The numerical integrator requires  $\phi(t)$  as a function of  $t$ , which can be obtained by the probability flow ODE sampler. In our calculations, we tested different sampling sizes ranging from 1 to 100 and found that it converges when the size is larger than 10. Therefore, we chose the sample size  $\epsilon$  as 10 to ensure the accuracy of the estimation.

## C Statistics examination


In figure 11 (with the same physical parameter setup as in section 4.3), we demonstrate the dependence of four key physical observables including their uncertainties on the size of the training datasets, while the size of the generated dataset remains constant (1024 configurations). The illustration uses gray markers to denote values derived from various



**Figure 11.** Dependence of four observables and their uncertainties on the size of the training dataset within the symmetric phase ( $\kappa = 0.21, \lambda = 0.022$ ), using a fixed number of 1024 generated configurations.



**Figure 12.** Dependence of observables on the number of generated samples obtained from a DM trained on a dataset of 5120 configurations within the symmetric phase ( $\kappa = 0.21, \lambda = 0.022$ ). The DM results are compared with predictions from the training set and with those from a larger unused dataset with 51200 configurations.



**Figure 13.** Dependence of the acceptance rate on the hopping parameter  $\kappa$  on a  $32 \times 32$  lattice.

training datasets, while colored markers represent values obtained from the generated datasets. The analysis reveals a trend towards convergence in the estimations of observables as the size of the training dataset increases, even with a relatively small generated dataset size.

In figure 12, we present the dependence of the observables on the number of generated samples obtained from a DM trained on a dataset of 5120 configurations. The observables' values calculated on the training dataset are traced with black lines, contrasting with gray lines that represent calculations over a larger, unused dataset (51200 configurations). The results show that observables from the generated datasets tend to align with those from the larger dataset as their size increases.

Although there is a noticeable deviation in the  $\langle M \rangle$  estimations towards the training dataset values, it can be conceivably improved by introducing the  $Z_2$  symmetry into the training data-sets explicitly, e.g., data augmentation with flipping all configurations.

## D Acceptance rate


We investigate the dependence of the acceptance rate on the value of the hopping parameter, the system size and the number of training epochs.

We have calculated the dependence of the acceptance rate on the hopping parameter, specifically at  $\kappa = 0.1, 0.21, 0.27, 0.4$ , using an identical setup as in the main part of the paper, i.e. with 250 training epochs,  $T = 100$  MC time-steps and 1024 configurations. In figure 13, we observe that the acceptance rate lies between 0.4 and 0.6. This observation indicates that the efficacy of the DM-MC method remains robust against the variations in interaction couplings, demonstrating its utility as an efficient sampling technique across a broad range of parameter values.

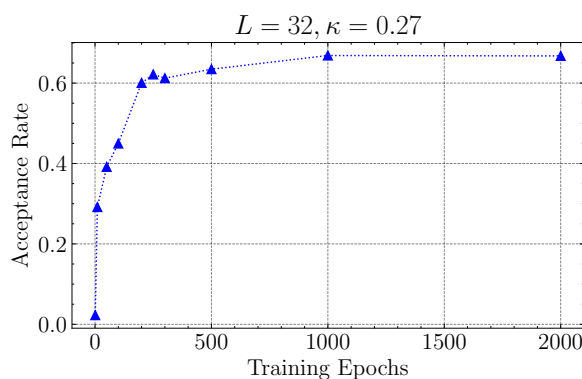
Next we vary the linear size of the lattice,  $L = 16, 24, 32, 48, 64$ , at fixed  $\kappa = 0.27$ , using an otherwise identical setup as above. In figure 13, we observe that the acceptance rate converges to around 0.6 for larger sizes, suggesting that the efficacy of the DM-MC will not be affected by the system size.

The dependence of the acceptance rate of DM-MC on the number of training epochs is shown in figure 15 for parameter values closer to the critical point. It shows that the acceptance rate increases rapidly within the first 200 epochs and converges slowly over the remaining 1800 epochs. Although our choice of 250 epochs used in the main part of the paper





**Figure 14.** Dependence of the acceptance rate on the linear lattice size at a fixed hopping parameter  $\kappa = 0.27$ .



**Figure 15.** Dependence of the acceptance rate on the number of epochs used for training, at  $\kappa = 0.27, L = 32, \lambda = 0.022$ .

yields an acceptance rate slightly below the asymptotic value, it is indeed an acceptable balance between training cost and performance.

Concerning the dependence of the training cost on the various parameters, we have found that

- for different  $\kappa$  values on a  $32 \times 32$  lattice the training cost is almost constant (at around 6.54 seconds per epoch on our infrastructure);
- increasing the volume results in a longer training time; for linear size  $L = 16, 24, 32, 48, 64$ , the training cost is 5.80, 6.19, 6.54, 6.97, 8.63 seconds per epoch on our infrastructure. A rough fit then yields a training time of  $L^{0.43} + 2.31$  seconds per epoch.

**Open Access.** This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

## References

- [1] F. Knechtli, M. Günther and M. Peardon, *Lattice Quantum Chromodynamics: Practical Essentials*, Springer (2017) [[DOI:10.1007/978-94-024-0999-4](https://doi.org/10.1007/978-94-024-0999-4)] [[INSPIRE](#)].
- [2] U. Wolff, *Critical Slowing Down*, *Nucl. Phys. B Proc. Suppl.* **17** (1990) 93 [[INSPIRE](#)].

- [3] J.M. Tomczak, *Deep Generative Modeling*, Springer International Publishing (2022) [[DOI:10.1007/978-3-030-93158-2](#)].
- [4] D. Boyda et al., *Applications of Machine Learning to Lattice Quantum Field Theory*, in the proceedings of the *Snowmass 2021*, Seattle, U.S.A., July 17–26 (2022) [[arXiv:2202.05838](#)] [[INSPIRE](#)].
- [5] K. Zhou, G. Endrődi, L.-G. Pang and H. Stöcker, *Regressive and generative neural networks for scalar field theory*, *Phys. Rev. D* **100** (2019) 011501 [[arXiv:1810.12879](#)] [[INSPIRE](#)].
- [6] J.M. Pawłowski and J.M. Urban, *Reducing Autocorrelation Times in Lattice Simulations with Generative Adversarial Networks*, *Mach. Learn. Sci. Tech.* **1** (2020) 045011 [[arXiv:1811.03533](#)] [[INSPIRE](#)].
- [7] 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 [[arXiv:1904.12072](#)] [[INSPIRE](#)].
- [8] G. Kanwar et al., *Equivariant flow-based sampling for lattice gauge theory*, *Phys. Rev. Lett.* **125** (2020) 121601 [[arXiv:2003.06413](#)] [[INSPIRE](#)].
- [9] K.A. Nicoli et al., *Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models*, *Phys. Rev. Lett.* **126** (2021) 032001 [[arXiv:2007.07115](#)] [[INSPIRE](#)].
- [10] M.S. Albergo et al., *Flow-based sampling for fermionic lattice field theories*, *Phys. Rev. D* **104** (2021) 114507 [[arXiv:2106.05934](#)] [[INSPIRE](#)].
- [11] M.S. Albergo et al., *Introduction to Normalizing Flows for Lattice Field Theory*, [arXiv:2101.08176](#) [[INSPIRE](#)].
- [12] L. Del Debbio, J.M. Rossney and M. Wilson, *Efficient modeling of trivializing maps for lattice  $\phi^4$  theory using normalizing flows: A first look at scalability*, *Phys. Rev. D* **104** (2021) 094507 [[arXiv:2105.12481](#)] [[INSPIRE](#)].
- [13] D.C. Hackett et al., *Flow-based sampling for multimodal distributions in lattice field theory*, [arXiv:2107.00734](#) [[INSPIRE](#)].
- [14] M.S. Albergo et al., *Flow-based sampling in the lattice Schwinger model at criticality*, *Phys. Rev. D* **106** (2022) 014514 [[arXiv:2202.11712](#)] [[INSPIRE](#)].
- [15] S. Bacchio, P. Kessel, S. Schaefer and L. Vaitl, *Learning trivializing gradient flows for lattice gauge theories*, *Phys. Rev. D* **107** (2023) L051504 [[arXiv:2212.08469](#)] [[INSPIRE](#)].
- [16] M. Caselle, E. Cellini, A. Nada and M. Panero, *Stochastic normalizing flows as non-equilibrium transformations*, *JHEP* **07** (2022) 015 [[arXiv:2201.08862](#)] [[INSPIRE](#)].
- [17] S. Chen et al., *Fourier-flow model generating Feynman paths*, *Phys. Rev. D* **107** (2023) 056001 [[arXiv:2211.03470](#)] [[INSPIRE](#)].
- [18] M. Gerdes et al., *Learning lattice quantum field theories with equivariant continuous flows*, *SciPost Phys.* **15** (2023) 238 [[arXiv:2207.00283](#)] [[INSPIRE](#)].
- [19] D. Albandea et al., *Learning trivializing flows*, *Eur. Phys. J. C* **83** (2023) 676 [[arXiv:2302.08408](#)] [[INSPIRE](#)].
- [20] 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 [[arXiv:2207.00980](#)] [[INSPIRE](#)].
- [21] R. Abbott et al., *Aspects of scaling and scalability for flow-based sampling of lattice QCD*, *Eur. Phys. J. A* **59** (2023) 257 [[arXiv:2211.07541](#)] [[INSPIRE](#)].

- [22] K.A. Nicoli et al., *Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories*, *Phys. Rev. D* **108** (2023) 114501 [[arXiv:2302.14082](#)] [[INSPIRE](#)].
- [23] L. Wang, Y. Jiang, L. He and K. Zhou, *Continuous-Mixture Autoregressive Networks Learning the Kosterlitz-Thouless Transition*, *Chin. Phys. Lett.* **39** (2022) 120502 [[arXiv:2005.04857](#)] [[INSPIRE](#)].
- [24] D. Luo et al., *Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models*, *Phys. Rev. Res.* **5** (2023) 013216 [[INSPIRE](#)].
- [25] M. Favoni, A. Ipp, D.I. Müller and D. Schuh, *Lattice Gauge Equivariant Convolutional Neural Networks*, *Phys. Rev. Lett.* **128** (2022) 032003 [[arXiv:2012.12901](#)] [[INSPIRE](#)].
- [26] J. Aronsson, D.I. Müller and D. Schuh, *Geometrical aspects of lattice gauge equivariant convolutional neural networks*, [arXiv:2303.11448](#) [[INSPIRE](#)].
- [27] R. Abbott et al., *Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions*, *Phys. Rev. D* **106** (2022) 074506 [[arXiv:2207.08945](#)] [[INSPIRE](#)].
- [28] K. Zhou, L. Wang, L.-G. Pang and S. Shi, *Exploring QCD matter in extreme conditions with Machine Learning*, *Prog. Part. Nucl. Phys.* **135** (2024) 104084 [[arXiv:2303.15136](#)] [[INSPIRE](#)].
- [29] W.-B. He et al., *High-energy nuclear physics meets machine learning*, *Nucl. Sci. Tech.* **34** (2023) 88 [[arXiv:2303.06752](#)] [[INSPIRE](#)].
- [30] L. Yang et al., *Diffusion Models: A Comprehensive Survey of Methods and Applications*, *ACM Comput. Surv.* **56** (2023) 1. [[arXiv:2209.00796](#)].
- [31] F.-A. Croitoru, V. Hondru, R.T. Ionescu and M. Shah, *Diffusion Models in Vision: A Survey*, *IEEE Trans. Pattern Anal. Machine Intell.* **45** (2023) 10850.
- [32] A. Ramesh et al., *Hierarchical Text-Conditional Image Generation with CLIP Latents*, [arXiv:2204.06125](#) [[INSPIRE](#)].
- [33] R. Rombach et al., *High-Resolution Image Synthesis with Latent Diffusion Models*, in the proceedings of the *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, U.S.A., June 18–24 (2022) [[DOI:10.1109/cvpr52688.2022.01042](#)].
- [34] V. Mikuni and B. Nachman, *Score-based generative models for calorimeter shower simulation*, *Phys. Rev. D* **106** (2022) 092009 [[arXiv:2206.11898](#)] [[INSPIRE](#)].
- [35] V. Mikuni, B. Nachman and M. Pettee, *Fast point cloud generation with diffusion models in high energy physics*, *Phys. Rev. D* **108** (2023) 036025 [[arXiv:2304.01266](#)] [[INSPIRE](#)].
- [36] G. Parisi and Y.S. Wu, *Perturbation theory without gauge fixing*, *Sci. China, A* **24** (1980) 483.
- [37] P.H. Damgaard and H. Hüffel, *Stochastic Quantization*, *Phys. Rept.* **152** (1987) 227 [[INSPIRE](#)].
- [38] M. Namiki, *Basic ideas of stochastic quantization*, *Prog. Theor. Phys. Suppl.* **111** (1993) 1 [[INSPIRE](#)].
- [39] G. Parisi, *On Complex Probabilities*, *Phys. Lett. B* **131** (1983) 393 [[INSPIRE](#)].
- [40] J. Berges, S. Borsanyi, D. Sexty and I.-O. Stamatescu, *Lattice simulations of real-time quantum fields*, *Phys. Rev. D* **75** (2007) 045007 [[hep-lat/0609058](#)] [[INSPIRE](#)].
- [41] G. Aarts and I.-O. Stamatescu, *Stochastic quantization at finite chemical potential*, *JHEP* **09** (2008) 018 [[arXiv:0807.1597](#)] [[INSPIRE](#)].
- [42] G. Aarts, *Can stochastic quantization evade the sign problem? The relativistic Bose gas at finite chemical potential*, *Phys. Rev. Lett.* **102** (2009) 131601 [[arXiv:0810.2089](#)] [[INSPIRE](#)].

- [43] E. Seiler, D. Sexty and I.-O. Stamatescu, *Gauge cooling in complex Langevin for QCD with heavy quarks*, *Phys. Lett. B* **723** (2013) 213 [[arXiv:1211.3709](#)] [[INSPIRE](#)].
- [44] D. Sexty, *Simulating full QCD at nonzero density using the complex Langevin equation*, *Phys. Lett. B* **729** (2014) 108 [[arXiv:1307.7748](#)] [[INSPIRE](#)].
- [45] G. Aarts, *Introductory lectures on lattice QCD at nonzero baryon number*, *J. Phys. Conf. Ser.* **706** (2016) 022004 [[arXiv:1512.05145](#)] [[INSPIRE](#)].
- [46] F. Attanasio, B. Jäger and F.P.G. Ziegler, *Complex Langevin simulations and the QCD phase diagram: Recent developments*, *Eur. Phys. J. A* **56** (2020) 251 [[arXiv:2006.00476](#)] [[INSPIRE](#)].
- [47] C.E. Berger et al., *Complex Langevin and other approaches to the sign problem in quantum many-body physics*, *Phys. Rept.* **892** (2021) 1 [[arXiv:1907.10183](#)] [[INSPIRE](#)].
- [48] K. Nagata, *Finite-density lattice QCD and sign problem: Current status and open problems*, *Prog. Part. Nucl. Phys.* **127** (2022) 103991 [[arXiv:2108.12423](#)] [[INSPIRE](#)].
- [49] G. Aarts, E. Seiler and I.-O. Stamatescu, *The Complex Langevin method: When can it be trusted?*, *Phys. Rev. D* **81** (2010) 054508 [[arXiv:0912.3360](#)] [[INSPIRE](#)].
- [50] G. Aarts, F.A. James, E. Seiler and I.-O. Stamatescu, *Complex Langevin: Etiology and Diagnostics of its Main Problem*, *Eur. Phys. J. C* **71** (2011) 1756 [[arXiv:1101.3270](#)] [[INSPIRE](#)].
- [51] K. Nagata, J. Nishimura and S. Shimasaki, *Argument for justification of the complex Langevin method and the condition for correct convergence*, *Phys. Rev. D* **94** (2016) 114515 [[arXiv:1606.07627](#)] [[INSPIRE](#)].
- [52] G. Aarts, E. Seiler, D. Sexty and I.-O. Stamatescu, *Complex Langevin dynamics and zeroes of the fermion determinant*, *JHEP* **05** (2017) 044 [*Erratum ibid.* **01** (2018) 128] [[arXiv:1701.02322](#)] [[INSPIRE](#)].
- [53] M. Scherzer, E. Seiler, D. Sexty and I.-O. Stamatescu, *Complex Langevin and boundary terms*, *Phys. Rev. D* **99** (2019) 014512 [[arXiv:1808.05187](#)] [[INSPIRE](#)].
- [54] M. Westh Hansen and D. Sexty, *Complex Langevin boundary terms in full QCD*, *PoS LATTICE2022* (2023) 163 [[arXiv:2212.12029](#)] [[INSPIRE](#)].
- [55] D. Alvestad, R. Larsen and A. Rothkopf, *Stable solvers for real-time Complex Langevin*, *JHEP* **08** (2021) 138 [[arXiv:2105.02735](#)] [[INSPIRE](#)].
- [56] D. Alvestad, R. Larsen and A. Rothkopf, *Towards learning optimized kernels for complex Langevin*, *JHEP* **04** (2023) 057 [[arXiv:2211.15625](#)] [[INSPIRE](#)].
- [57] N.M. Lampl and D. Sexty, *Real time evolution of scalar fields with kernelled Complex Langevin equation*, [[arXiv:2309.06103](#)] [[INSPIRE](#)].
- [58] D.J.E. Callaway and A. Rahman, *The Microcanonical Ensemble: A New Formulation of Lattice Gauge Theory*, *Phys. Rev. Lett.* **49** (1982) 613 [[INSPIRE](#)].
- [59] H. Risken, *The Fokker-Planck equation: Methods of solution and application*, Springer (1996) [[DOI:10.1007/978-3-642-61544-3](#)].
- [60] M. Namiki et al., *Stochastic quantization*, Springer Berlin, Heidelberg (1992) [[DOI:10.1007/978-3-540-47217-9](#)] [[INSPIRE](#)].
- [61] G.G. Batrouni et al., *Langevin Simulations of Lattice Field Theories*, *Phys. Rev. D* **32** (1985) 2736 [[INSPIRE](#)].
- [62] J. Sohl-Dickstein, E.A. Weiss, N. Maheswaranathan and S. Ganguli, *Deep Unsupervised Learning using Nonequilibrium Thermodynamics*, [[arXiv:1503.03585](#)] [[INSPIRE](#)].

- [63] Y. Song and S. Ermon, *Generative modeling by estimating gradients of the data distribution*, in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox and R. Garnett eds., Curran Associates Inc. (2019) [ISBN: 9781713807933].
- [64] Y. Song et al., *Score-Based Generative Modeling through Stochastic Differential Equations*, [arXiv:2011.13456](#) [INSPIRE].
- [65] J. Ho, A. Jain and P. Abbeel, *Denoising diffusion probabilistic models*, in *NIPS '20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, U.S.A., Curran Associates Inc. (2020), pp. 6840–6851.
- [66] B.D.O. Anderson, *Reverse-time diffusion equation models*, *Stochastic Processes and their Applications* **12** (1982) 313.
- [67] A. Hyvärinen, *Estimation of non-normalized statistical models by score matching*, *J. Mach. Learn. Res.* **6** (2005) 695.
- [68] P. Vincent, *A Connection Between Score Matching and Denoising Autoencoders*, *Neural Comput.* **23** (2011) 1661.
- [69] M. Welling and Y.W. Teh, *Bayesian learning via stochastic gradient langevin dynamics*, in the proceedings of the *28th International Conference on Machine Learning*, Madison, U.S.A., June 28 – July 2 (2011) [Omnipress (2011), pp. 681–688].
- [70] D. Maoutsa, S. Reich and M. Opper, *Interacting Particle Solutions of Fokker-Planck Equations Through Gradient-Log-Density Estimation*, *Entropy* **22** (2020) 802 [[arXiv:2006.00702](#)].
- [71] W. Grathwohl, R.T.Q. Chen, J. Bettencourt and D. Duvenaud, *Scalable reversible generative models with free-form continuous dynamics*, in the proceedings of the *ICLR 2019 — The Seventh International Conference on Learning Representations International Conference on Learning Representations*, New Orleans, U.S.A. May 6–9 (2019).
- [72] M. Luscher, *Trivializing maps, the Wilson flow and the HMC algorithm*, *Commun. Math. Phys.* **293** (2010) 899 [[arXiv:0907.5491](#)] [INSPIRE].
- [73] Z. Bern, M.B. Halpern and L. Sadun, *Continuum Regularization of Quantum Field Theory. IV. Langevin Renormalization*, *Z. Phys. C* **35** (1987) 255 [INSPIRE].
- [74] J.M. Pawłowski, I.-O. Stamatescu and F.P.G. Ziegler, *Cooling Stochastic Quantization with colored noise*, *Phys. Rev. D* **96** (2017) 114505 [[arXiv:1705.06231](#)] [INSPIRE].
- [75] J. Smit, *Introduction to Quantum Fields on a Lattice*, Cambridge University Press (2023) [[DOI:10.1017/9781009402705](#)] [INSPIRE].
- [76] S. Akiyama, Y. Kuramashi and Y. Yoshimura, *Phase transition of four-dimensional lattice  $\phi_4$  theory with tensor renormalization group*, *Phys. Rev. D* **104** (2021) 034507 [[arXiv:2101.06953](#)] [INSPIRE].
- [77] R.M. Neal, *MCMC using Hamiltonian dynamics*, [arXiv:1206.1901](#) [[DOI:10.1201/b10905](#)].
- [78] O. Ronneberger, P. Fischer and T. Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, [arXiv:1505.04597](#) [INSPIRE].
- [79] K. Binder, *Critical Properties from Monte Carlo Coarse Graining and Renormalization*, *Phys. Rev. Lett.* **47** (1981) 693 [INSPIRE].
- [80] D. Bachtis, G. Aarts and B. Lucini, *Mapping distinct phase transitions to a neural network*, *Phys. Rev. E* **102** (2020) 053306 [[arXiv:2007.00355](#)] [INSPIRE].

- [81] ALPHA collaboration, *Critical slowing down and error analysis in lattice QCD simulations*, *Nucl. Phys. B* **845** (2011) 93 [[arXiv:1009.5228](#)] [[INSPIRE](#)].
- [82] ALPHA collaboration, *Monte Carlo errors with less errors*, *Comput. Phys. Commun.* **156** (2004) 143 [*Erratum ibid.* **176** (2007) 383] [[hep-lat/0306017](#)] [[INSPIRE](#)].
- [83] F. Joswig, S. Kuberski, J.T. Kuhlmann and J. Neuendorf, *pyerrors: A python framework for error analysis of Monte Carlo data*, *Comput. Phys. Commun.* **288** (2023) 108750 [[arXiv:2209.14371](#)] [[INSPIRE](#)].
- [84] J. Robnik and U. Seljak, *Fluctuation without dissipation: Microcanonical Langevin Monte Carlo*, [arXiv:2303.18221](#) [[INSPIRE](#)].
- [85] J. Cotler and S. Rezchikov, *Renormalizing Diffusion Models*, [arXiv:2308.12355](#) [[INSPIRE](#)].