

RECEIVED: October 8, 2025

ACCEPTED: November 16, 2025

PUBLISHED: December 22, 2025

Combining complex Langevin dynamics with score-based and energy-based diffusion models

Gert Aarts ,^a Diaa E. Habibi ,^a Lingxiao Wang ,^{b,c} and Kai Zhou ,^{d,e}

^aCentre for Quantum Fields and Gravity, Department of Physics, Swansea University, SA2 8PP, Swansea, U.K.

^bInterdisciplinary Theoretical and Mathematical Sciences Program (*iTHEMS*), RIKEN, Wako, Saitama 351-0198, Japan

^cInstitute for Physics of Intelligence, Graduate School of Science, The University of Tokyo, Bunkyo-ku, Tokyo 113-0033, Japan

^dSchool of Science and Engineering, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), Guangdong, 518172, China

^eSchool of Artificial Intelligence, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), Guangdong, 518172, China

E-mail: g.aarts@swansea.ac.uk, n.e.habibi@swansea.ac.uk,
lingxiao.wang@riken.jp, zhoukai@cuhk.edu.cn

ABSTRACT: Theories with a sign problem due to a complex action or Boltzmann weight can sometimes be numerically solved using a stochastic process in the complexified configuration space. However, the probability distribution effectively sampled by this complex Langevin process is not known *a priori* and notoriously hard to understand. In generative AI, diffusion models can learn distributions, or their log derivatives, from data. We explore the ability of diffusion models to learn the distributions sampled by a complex Langevin process, comparing score-based and energy-based diffusion models, and speculate about possible applications.

KEYWORDS: Algorithms and Theoretical Developments, Non-Zero Temperature and Density, Other Lattice Field Theories

ARXIV EPRINT: [2510.01328](https://arxiv.org/abs/2510.01328)

Contents

1	Introduction	1
2	Complex Langevin dynamics	2
3	Score and energy-based diffusion models	4
3.1	Score-based models	5
3.2	Energy-based models	6
4	Complex-valued quartic model	7
5	Summary and outlook	12
A	Computational details	13
B	Real-valued Gaussian mixture	14
C	Lefschetz thimbles	16
D	Estimating systematic effects	16

1 Introduction

Theories with a complex Boltzmann weight suffer from sign and overlap problems, which makes them hard to simulate using conventional numerical methods based on importance sampling [1]. A well-known example is QCD at non-zero baryon density, for which the quark determinant is complex-valued when the quark chemical potential is real,

$$[\det M(\mu)]^* = \det M(-\mu^*) \in \mathbb{C}, \quad (1.1)$$

see e.g. the reviews [2, 3]. A potential solution is to take the complexity seriously and extend the theory into the complex plane, or complexified manifold more generally. This idea underpins complex Langevin dynamics [4, 5], Lefschetz thimbles [6], holomomorphic flow [7], and variations thereof.

Complex Langevin (CL) dynamics, in which the degrees of freedom are analytically extended, does not rely on importance sampling but explores a complexified manifold via a stochastic process [4, 5]. It is an extension of stochastic quantisation [8, 9]. CL has been shown to work in lattice field theories in three [10] and four [11] Euclidean dimensions with a severe sign problem, including in QCD in four dimensions [12–18], but it may also fail, even in simple models [19–21]. This situation was clarified a few years ago [22–24] by the derivation of the formal relation between the complex distribution on the real manifold and the real and semi-positive distribution on the complexified manifold, which is effectively sampled during the CL process. This resulted in practical criteria for correctness which need to be verified *a*

posteriori [22–24]. Nevertheless, issues remain and the reliability of the method depends on a precise understanding of the behaviour of this distribution, in particular at infinity and near poles in the CL drift. Recent work can be found in e.g. refs. [25–33].

A crucial role is played by the distribution on the complexified manifold. Unfortunately, this distribution turns out to be elusive, as the Fokker-Planck equation linked to the CL process cannot be solved in general. In fact, even convergence is hard to understand, except in some simple cases, such as Gaussian models [34] and models in which the decay at infinity can be precisely understood (see below) [35]. A better characterisation of the distribution would, therefore, be welcome. Independently of which approach is used, relating averages over complex-valued Boltzmann weights to statistical averages over real, semi-positive probability distributions is an open problem worth studying more generally [36, 37].

Diffusion models [38–44] are a class of generative methods, which learn distributions from data. They rely on a stochastic process, similar to stochastic quantisation, but instead of using a known drift term derived from the underlying theory, they learn the drift from data previously collected. We have recently explored the relation between diffusion models and stochastic quantisation in scalar [45, 46] and U(1) gauge theories [47], and studied the evolution of higher-order cumulants in detail [48]. Further connections between diffusion models, field theory and the renormalisation group are pointed out in refs. [49–51]. Diffusion models are very flexible and can be formulated e.g. in real space or momentum space [52].

Since diffusion models do not require *a priori* knowledge of the distribution itself but learn directly from data, they can be employed to study the distribution sampled during the CL process [53]. When successfully trained, diffusion models can then be used to both generate additional configurations and to deepen our understanding of the real distribution on the complexified manifold, offering a fresh perspective. This is what we explore here.

The paper is organised as follows. We start with a brief summary of CL dynamics in section 2. In section 3, we first introduce diffusion models and then discuss score-based and energy-based diffusion models to construct distributions and their log derivatives. Since the purpose of this paper is to address conceptual questions, we apply the framework in section 4 to a simple and well-understood model [35]. A summary and outlook are given in section 5. Appendix A contains computational details. To ensure that our findings are not specific to the complexification required for CL dynamics, we study in appendix B an exactly solvable model with a real distribution to start with. Appendix C contains a comparison with Lefschetz thimbles. Appendix D finally contains our approach to estimating systematic uncertainties.

2 Complex Langevin dynamics

We start by summarising the idea behind stochastic quantisation and complex Langevin dynamics, considering for simplicity one degree of freedom x with a complex-valued Boltzmann weight $\rho(x)$. The setup is easily extended to many degrees of freedom and (lattice) field theory [4, 13, 54, 55].

As always, observables are defined as

$$\langle O(x) \rangle = \int dx \rho(x) O(x), \quad \rho(x) = \frac{1}{Z} \exp[-S(x)], \quad Z = \int dx \rho(x). \quad (2.1)$$

The Langevin process and drift read

$$\dot{x}(t) = K[x(t)] + \sqrt{2}\eta(t), \quad K(x) = \frac{d}{dx} \log \rho(x) = -\frac{dS(x)}{dx}, \quad (2.2)$$

where the dot indicates the Langevin time derivative and the noise $\eta \sim \mathcal{N}(0, 1)$ is Gaussian. The factor of $\sqrt{2}$ in the noise term is conventional and can be exchanged for a factor of $1/2$ in front of the drift, by rescaling Langevin time. The corresponding Fokker-Planck equation (FPE) is

$$\partial_t \rho(x; t) = \partial_x [\partial_x - K(x)] \rho(x; t). \quad (2.3)$$

For a real Boltzmann weight and drift, this process converges to the stationary solution $\rho(x) \sim \exp[-S(x)]$, typically exponentially fast [9].

When the weight is complex, one may extend $x \rightarrow z = x + iy$ into the complex plane and write

$$\dot{z}(t) = K[z(t)] + \sqrt{2}\eta(t), \quad K(z) = \frac{d}{dz} \log \rho(z) = -\frac{dS(z)}{dz}. \quad (2.4)$$

However, in this case the FPE cannot be used to show convergence as the corresponding Fokker-Planck Hamiltonian is no longer semi-positive definite [9].

Instead, one may take the real and imaginary parts of the equation above and consider the CL process,

$$\begin{aligned} \dot{x}(t) &= K_x[x(t), y(t)] + \sqrt{2N_x}\eta_x(t), & K_x(x, y) &= \text{Re} \frac{d}{dz} \log \rho(z) \Big|_{z \rightarrow x+iy}, \\ \dot{y}(t) &= K_y[x(t), y(t)] + \sqrt{2N_y}\eta_y(t), & K_y(x, y) &= \text{Im} \frac{d}{dz} \log \rho(z) \Big|_{z \rightarrow x+iy}, \end{aligned} \quad (2.5)$$

with the constraint $N_x - N_y = 1$. The FPE for this process reads

$$\partial_t p(x, y; t) = [\partial_x (N_x \partial_x - K_x) + \partial_y (N_y \partial_y - K_y)] p(x, y; t), \quad (2.6)$$

such that

$$\langle O[x(t) + iy(t)] \rangle_\eta = \int dx dy p(x, y; t) O(x + iy). \quad (2.7)$$

It is preferable to consider real noise, $N_x = 1, N_y = 0$ [22]. Unlike the original weight $\rho(z)$, $p(x, y; t)$ is real and semi-positive definite, as it represents the distribution effectively sampled by the real Langevin process in the two-dimensional plane.

This setup should be contrasted with the case of a proper two-dimensional process starting from a known (real and semi-positive) distribution $p(\mathbf{x})$, with $\mathbf{x} = (x, y)$. In that case, one can apply real Langevin dynamics, given by

$$\dot{\mathbf{x}}(t) = \mathbf{K}[\mathbf{x}(t)] + \sqrt{2}\boldsymbol{\eta}(t), \quad \mathbf{K}(\mathbf{x}) = \nabla \log p(\mathbf{x}), \quad \langle \eta_i(t) \eta_j(t') \rangle = \delta_{ij} \delta(t - t'), \quad (2.8)$$

with the corresponding FPE,

$$\partial_t p(\mathbf{x}; t) = \nabla \cdot [\nabla - \mathbf{K}(\mathbf{x})] p(\mathbf{x}; t). \quad (2.9)$$

It is easy to see that $p(\mathbf{x})$ is a stationary solution of this FPE.

Going back to the CL equation, this process yields the correct answer if a stationary solution to the FPE (2.6) exists, such that [22, 23]

$$\int dx dy p(x, y) O(x + iy) = \int dx \rho(x) O(x), \quad (2.10)$$

or, shifting the integration variables at a formal level,

$$\rho(x) = \int dy p(x - iy, y). \quad (2.11)$$

Considerable effort has been invested in deriving criteria for correctness related to the behaviour of $p(x, y)$ at infinity and near poles of the drift (if there are any), which can be used *a posteriori* to justify the results [22–27, 29, 31–33]. A better understanding of $p(x, y)$ in the stationary limit is very welcome. We turn to diffusion models to help with this question.

3 Score and energy-based diffusion models

Diffusion models are a class of probabilistic generative models, which learn from a data set representing the target distribution $p_0(\mathbf{x})$ [38–42]. During the forward process, data are gradually corrupted, by applying noise which increases incrementally in strength, and a neural network is trained to learn the change in the probability distribution, using the decomposition

$$p(\mathbf{x}, t) = \int d\mathbf{x}_0 p(\mathbf{x}, t | \mathbf{x}_0) p_0(\mathbf{x}_0). \quad (3.1)$$

The process can then be reversed and, after training, new instances representative of the original data set are generated during the backward process.

We use the description in terms of stochastic differential equations (SDEs) [41–44] and refer to refs. [45–47] for applications in lattice field theory. Ref. [48] contains a detailed analysis of the evolution of higher-order cumulants. The forward process is described by the SDE

$$\dot{\mathbf{x}}(t) = \frac{1}{2} \mathbf{K}[\mathbf{x}(t), t] + g(t) \boldsymbol{\eta}(t), \quad (3.2)$$

where $\mathbf{K}[\mathbf{x}(t), t]$ is a possible drift term and $\boldsymbol{\eta} \sim \mathcal{N}(0, 1)$ is again Gaussian noise. Compared to the previous section, we have rescaled time with a factor of 2, see eqs. (2.2, 2.8), and introduced a diffusion coefficient $g(t)$, setting the time-dependent noise strength. The initial conditions for this process are determined by the target distribution $\mathbf{x}(0) = \mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ and the process runs between $0 \leq t \leq T$.

The corresponding SDE for the denoising or backward process reads [56]

$$\dot{\mathbf{x}}(t) = \frac{1}{2} \mathbf{K}[\mathbf{x}(t), t] - g^2(t) \nabla \log p(\mathbf{x}, t) + g(t) \boldsymbol{\eta}(t), \quad (3.3)$$

where the process now starts at $t = T$ and time runs backwards to $t = 0$. Initial conditions are sampled from a simple prior distribution, such as the normal distribution with a variance comparable to the variance obtained at the end of the forward process. The additional term in the drift, the so-called score, $\nabla \log p(\mathbf{x}, t)$, is not known *a priori* and is modelled by a

neural network $\mathbf{s}_\theta(\mathbf{x}, t)$, where θ denotes all the trainable network parameters. It is learned during the forward process, by minimising a loss function imposing

$$\mathbf{s}_\theta(\mathbf{x}, t) \approx \nabla \log p(\mathbf{x}, t), \quad (3.4)$$

in conjunction with the decomposition (3.1) and Jensen's inequality. After the diffusion model has been trained, new samples from the target distribution can be generated by numerically solving the backward process (3.3) using eq. (3.4),

$$\dot{\mathbf{x}}(t) = \frac{1}{2} \mathbf{K}[\mathbf{x}(t), t] - g^2(t) \mathbf{s}_{\theta^*}(\mathbf{x}, t) + g(t) \boldsymbol{\eta}(t), \quad (3.5)$$

at a (near-)optimal set of network parameters θ^* . When the target distribution is known, it is possible to include an accept/reject step [47]. The similarities between stochastic quantisation and the dynamics in diffusion models have been noted [45].

3.1 Score-based models

Starting from the Fisher divergence [41],

$$\mathcal{F}(\theta, \lambda) = \frac{1}{2} \int_0^T dt \mathbb{E}_{p(\mathbf{x}, t)} [\lambda(t) \|\mathbf{s}_\theta(\mathbf{x}, t) - \nabla \log p(\mathbf{x}, t)\|^2], \quad (3.6)$$

score-based models are trained using the loss function

$$\mathcal{L}(\theta, \lambda) = \frac{1}{2} \int_0^T dt \mathbb{E}_{p(\mathbf{x}, t)} [\lambda(t) \|\mathbf{s}_\theta(\mathbf{x}, t) - \nabla \log p(\mathbf{x}, t|\mathbf{x}_0)\|^2], \quad (3.7)$$

where the weight $\lambda(t)$ is chosen to be the variance of the noise at time t and Jensen's inequality was used in combination with eq. (3.1). The expectation includes the average over the initial distribution $p_0(\mathbf{x})$, obtained by summing over the data set.

Given eqs. (3.4, 3.6), it is tempting to identify the integrated score at the end of the backward process with the log of the target distribution function. This identification would provide a handle on the (unnormalised) data distribution, which may be especially useful for complex Langevin dynamics, in which this distribution is elusive, as mentioned above. However, here we note (as has also been observed before, see e.g. refs. [57–59]) that there is no guarantee that the score is conservative. Indeed, as we will see below, in general one finds that the score contains both a conservative and non-conservative component,

$$\mathbf{s}_\theta(\mathbf{x}, t) = \nabla \Phi_\theta(\mathbf{x}, t) + \mathbf{r}_\theta(\mathbf{x}, t), \quad \nabla \cdot \mathbf{r}_\theta(\mathbf{x}, t) = 0, \quad (3.8)$$

and is hence not integrable. Note that this is not in conflict with the learning objective. Inserting the decomposition (3.8) in eq. (3.6), one finds for the norm under the integral

$$\|\mathbf{s}_\theta(\mathbf{x}, t) - \nabla \log p(\mathbf{x}, t)\|^2 = \|\nabla \Phi_\theta(\mathbf{x}, t) - \nabla \log p(\mathbf{x}, t)\|^2 + \|\mathbf{r}_\theta(\mathbf{x}, t)\|^2, \quad (3.9)$$

where we used partial integration in the loss function, ignoring boundary terms. Hence the (non)conservative parts are independently minimised. However, due to the non-conservative component, the integrated score,

$$\Phi_\theta(\mathbf{x}, t) = \Phi_\theta(\mathbf{0}, t) + \int_{\gamma: \mathbf{0} \rightarrow \mathbf{x}} d\mathbf{x}' \cdot \mathbf{s}_\theta(\mathbf{x}', t), \quad (3.10)$$

is path dependent and this relation cannot be used to construct $\Phi_\theta(\mathbf{x}, t)$ unambiguously. Specialising to two dimensions, while introducing a third dimension for convenience (but no z dependence), we can decompose the score in a gradient and a curl,

$$\mathbf{s}_\theta(\mathbf{x}, t) = \nabla \Phi_\theta(\mathbf{x}, t) + \nabla \times \mathbf{A}_\theta(\mathbf{x}, t) \quad \mathbf{A}_\theta(\mathbf{x}, t) = (0, 0, A_\theta(\mathbf{x}, t)), \quad (3.11)$$

or

$$s_{x,\theta}(\mathbf{x}, t) = \partial_x \Phi_\theta(\mathbf{x}, t) + \partial_y A_\theta(\mathbf{x}, t), \quad s_{y,\theta}(\mathbf{x}, t) = \partial_y \Phi_\theta(\mathbf{x}, t) - \partial_x A_\theta(\mathbf{x}, t). \quad (3.12)$$

Both scalar functions satisfy a Poisson equation

$$\nabla^2 \Phi_\theta(\mathbf{x}, t) = \nabla \cdot \mathbf{s}_\theta(\mathbf{x}, t), \quad \nabla^2 A_\theta(\mathbf{x}, t) = -(\nabla \times \mathbf{s}_\theta(\mathbf{x}, t))_z, \quad (3.13)$$

sourced by derivatives of the score. Solutions depend crucially on the boundary conditions. However, at the boundaries of the data manifold, the score is not well determined, due to the lack of data to train the model well. These Poisson equations also make clear the possibility of adding a harmonic function $h(\mathbf{x})$, with $\nabla^2 h(\mathbf{x}) = 0$, to either, without affecting the score. We emphasise that these observations do not undermine score-based diffusion models, which only require knowledge of the score itself, but complicate the explicit construction of the data distribution from the trained score.

3.2 Energy-based models

If we insist on learning the data distribution directly, it is useful to turn to energy-based models, see e.g. refs. [58–60], and impose

$$E_\theta(\mathbf{x}, t) \approx -\log p(\mathbf{x}, t), \quad (3.14)$$

directly, with an undetermined normalisation. To build on the experience of score-based models, ref. [60] proposed to use the following energy function,

$$E_\theta(\mathbf{x}, t) = \frac{1}{2} \|\mathbf{v}_\theta(\mathbf{x}, t)\|^2 \approx -\log p(\mathbf{x}, t), \quad (3.15)$$

leading to the approximate score

$$-\partial_i E_\theta(\mathbf{x}, t) = -\mathbf{v}_\theta(\mathbf{x}, t) \cdot \partial_i \mathbf{v}_\theta(\mathbf{x}, t) \approx \partial_i \log p(\mathbf{x}, t), \quad (3.16)$$

which is conservative by construction. The advantage of this formulation is that it can be used in the loss function (3.7), with training proceeding in the standard manner. It is noted that the choice of energy parametrisation is not unique and different choices may lead to different loss manifolds during training, requiring a separate treatment when choosing hyperparameters. Since the energy parametrisation $E_\theta(\mathbf{x}, t)$ is differentiated, it is important to ensure it is smooth enough by a convenient selection of activation function/nonlinearity such as SiLU or Mish, rather than ReLU or LeakyReLU. The need for additional derivatives makes this method more expensive. Finally, we note that the energy is semi-positive; due to the undetermined normalisation, this is not a restriction.

Below we refer to the two approaches as score-based models (SBMs) and energy-based models (EBMs). We are in particular interested in the learned scores at the end of the backward process, which we denote as

$$\begin{aligned} \text{SBM: } \mathbf{s}_\theta(\mathbf{x}) &= \lim_{t \rightarrow 0} \mathbf{v}_\theta(\mathbf{x}, t), \\ \text{EBM: } \mathbf{s}_\theta(\mathbf{x}) &= -\nabla E_\theta(\mathbf{x}) = -\lim_{t \rightarrow 0} \nabla E_\theta(\mathbf{x}, t) = -\lim_{t \rightarrow 0} \mathbf{v}_\theta(\mathbf{x}, t) \cdot \nabla \mathbf{v}_\theta(\mathbf{x}, t). \end{aligned} \quad (3.17)$$

Here $\mathbf{v}_\theta(\mathbf{x}, t)$ is the trained neural network, with (near-)optimal parameters θ^* . Below we drop the $*$ from θ^* to avoid a cluttering of symbols.

In the numerical implementation of the SBMs and the EBMs we used a time-conditioned feed-forward neural network, incorporating the \mathbf{Z}_2 symmetry of the target distribution (see below) and employing an exponential moving average (EMA) of the weights. More details can be found in appendix A.

4 Complex-valued quartic model

We now apply the diffusion models of the previous section to data generated using complex Langevin dynamics. Since the emphasis in this paper is on conceptual issues, we consider a simple and well-studied example with one degree of freedom. After complexification, $x \rightarrow z = x + iy$, the model has two degrees of freedom and hence we are interested in the distribution $p(x, y)$. In appendix B we verify that our findings are consistent in an exactly solvable *real-valued* model, in which the distribution $p(x, y)$ is known.

We consider the quartic model with a complex mass parameter [35]

$$S = \frac{1}{2}\sigma_0 x^2 + \frac{1}{4}\lambda x^4, \quad \sigma_0 = A + iB. \quad (4.1)$$

Exact results can be obtained by a direct evaluation of the partition function,

$$Z = \int dx e^{-S(x)} = \sqrt{\frac{4\xi}{\sigma_0}} e^\xi K_{-\frac{1}{4}}(\xi), \quad (4.2)$$

where $\xi = \sigma_0^2/(8\lambda)$ and $K_q(\xi)$ is the modified Bessel function of the second kind. Subsequently, moments $\mu_n = \mathbb{E}[x^n]$ are obtained by differentiating with respect to σ_0 . Odd moments vanish.

The model is of interest since it is possible to formulate exactly when CL dynamics works correctly and when it fails: provided that $3A^2 - B^2 \geq 0$, the CL process is contained in a strip $-y_- < y < y_+$, with [35]

$$y_-^2 = \frac{A}{2\lambda} \left(1 - \sqrt{1 - \frac{B^2}{3A^2}} \right), \quad y_- > 0. \quad (4.3)$$

In this case CL dynamics yields the correct results [35]. When $3A^2 - B^2 < 0$, the distribution is no longer contained and decays with a power law, $p(x, y) \sim (x^2 + y^2)^{-3}$ for large x, y . Higher-order moments are then no longer well-defined and CL dynamics fails. In the numerical experiments we take $A = B = \lambda = 1$, for which $y_- = 0.3029$. In figure 1 (left) we show the CL drift. The boundaries at $|y| = y_-$ are clearly visible: all arrows point inwards and hence the process cannot escape the strip $|y| < y_-$ in the case that real noise is used.

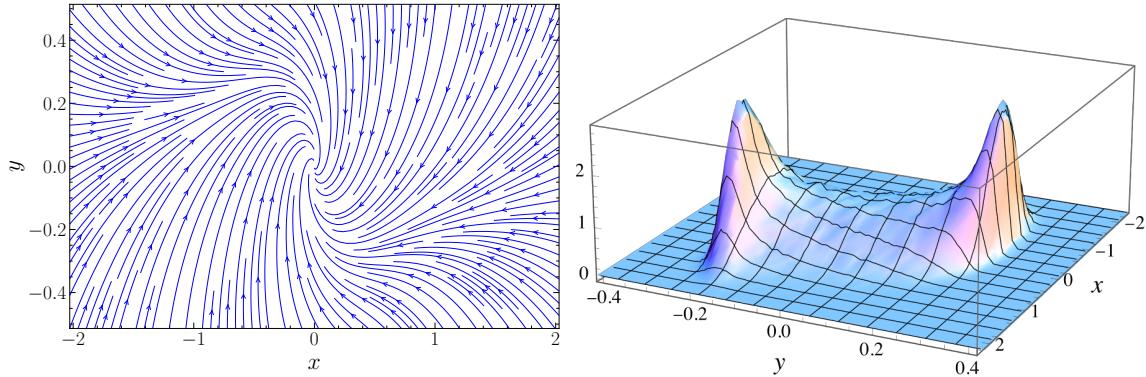


Figure 1. Complex-valued quartic model with parameters $\sigma_0 = 1 + i$ and $\lambda = 1$: CL drift in the complex plane (left) and histogram $P(x, y)$ obtained by sampling the CL process (right) [35].

We solve the CL process with real noise by discretising eq. (2.5) with stepsize ϵ using the higher-order algorithm of ref. [61], first applied to CL dynamics in ref. [10]. This algorithm improves stepsize corrections from $\mathcal{O}(\epsilon)$ to $\mathcal{O}(\epsilon^{3/2})$. Collecting the data in a histogram yields the ‘empirical’ histogram shown in figure 1 (right). As stated above, no analytical expression for this distribution is available.¹ The distribution is strictly zero when $|y| > y_-$, consistent with the argument based on the drift, given above.

Without direct access to the distribution, it is hard to assess the reliability of the approach. Diffusion models, however, can learn (the log-derivative of) this distribution. Notably, for a diffusion model it is irrelevant what the origin of the configurations is. We have trained the diffusion models using the CL data as the training set. Details can be found in appendix A. In figure 2 we show the learned scores at the end of the backward process in the score-based (left) and energy-based (right) formulations, see eq. (3.17). These vector fields should be contrasted with the CL drift of figure 1 (left). The first observation is that they are different, as they should be. Recall that the CL drift is used in the CL equation with noise in the x direction only, whereas the score is used in the diffusion model with noise applied in both directions. Moreover, the CL drift is not integrable, since $\partial_x K_y(\mathbf{x}) \neq \partial_y K_x(\mathbf{x})$. In particular it has an attractive fixed point at the origin, while the origin in the diffusion models is a saddle point. Reversely, in the diffusion models, the two peaks, at $\mathbf{x}_{\text{peak}} \sim \pm(0.6, -0.25)$, are attractive. To contain the data within the strip, also in the diffusion models the drift is pointing inwards at $|y| = y_-$. For completeness, we compare these drifts with Lefschetz thimbles in appendix C.

While the two drifts in figure 2 look quite similar, the one obtained in the score-based formulation is not conservative. This is demonstrated in figure 3 (left), where the curl of the score, $\partial_x s_y(\mathbf{x}) - \partial_y s_x(\mathbf{x})$, is shown, averaged over 10 independently trained models. Hence we cannot integrate the score directly. We can, however, obtain the distribution by sampling the process, as in the CL case. The result is shown in figure 3 (right). Comparing the histograms in figures 1 and 3, we observe that the trained score-based model captures the two peaks characteristic of this model as well as the boundary restrictions from the

¹In ref. [35] an approximate solution of the FPE was given using a double expansion in terms of Hermite functions.

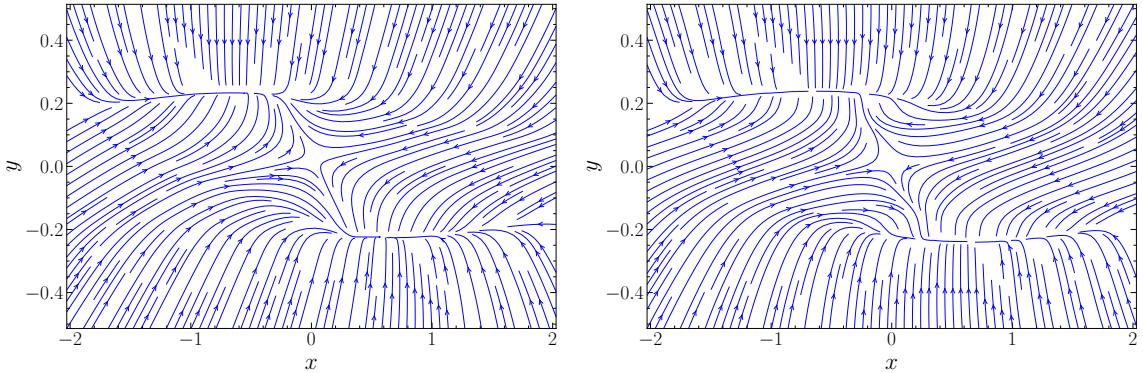


Figure 2. Learned scores in the quartic model at the end of the backward process, using a score-based (left) and energy-based (right) diffusion model.

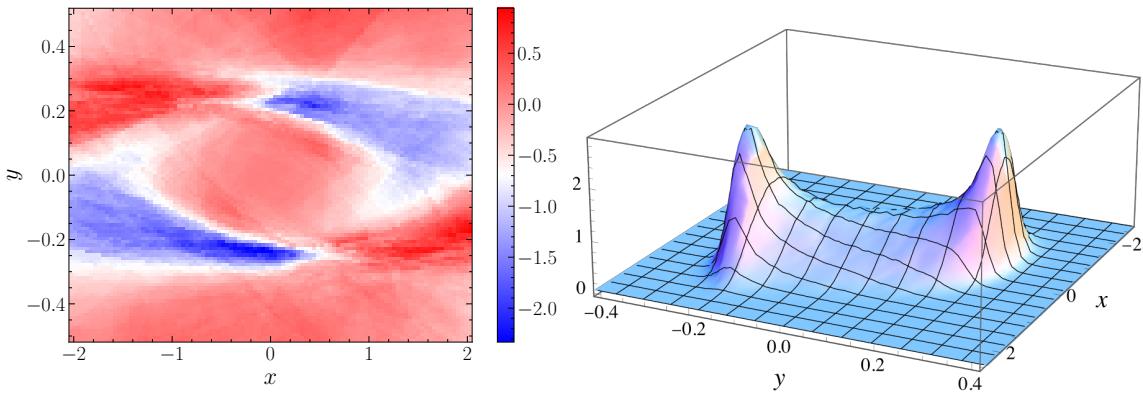


Figure 3. Quartic model using the score-based formulation: curl of the score, averaged over 10 independently trained models (left) and histogram obtained by sampling data using the process learnt by the score-based model (right).

training data, although the edges at $|y| = y_-$ seem less sharp. We discuss a more quantitative comparison below.

In the energy-based formulation, the energy $E_\theta(\mathbf{x})$ is learned directly. In figure 4 (left) we show a three-dimensional plot, restricted to $E_\theta(\mathbf{x}) \leq 10$ for clarity. In figure 5 we show the same energy using a contour plot. We note the two (shallow) minima around \mathbf{x}_{peak} and a quickly rising energy outside the main region of interest. The score is given by the gradient of the energy and is conservative by construction; we have verified that this is the case within numerical precision when the derivatives are computed numerically.

In contrast to the cases above, we now have direct access to the distribution, by exponentiation and without any further sampling. The result is shown figure 4 (right). We have determined the undetermined prefactor by normalising the distribution to 1. It should not come as a surprise that the distribution again resembles the ones obtained previously. The edges at $|y| = y_-$ seem as sharp as in the original CL distribution. We emphasise that this is the first time a parametrisation of the distribution sampled by a CL process is obtained in a non-trivial case without explicitly constructing a histogram by further sampling.

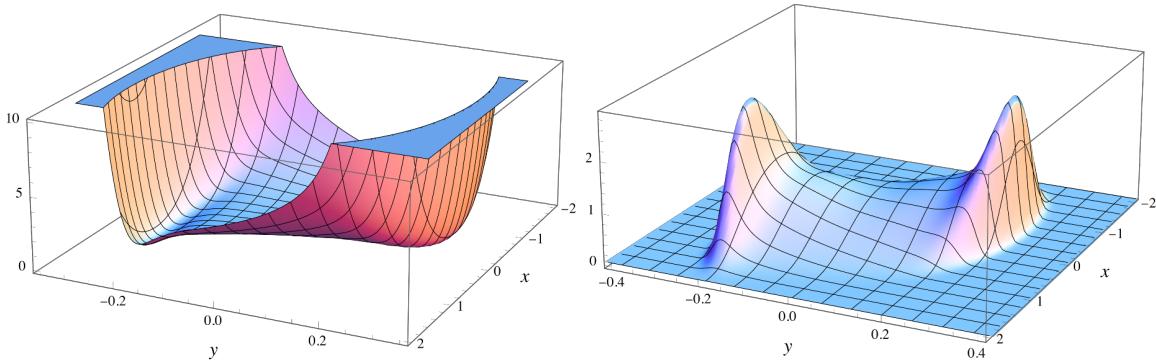


Figure 4. Quartic model: energy $E_\theta(\mathbf{x})$ learned in the energy-based diffusion model (left) and the corresponding distribution $p_\theta(\mathbf{x}) \sim \exp[-E_\theta(\mathbf{x})]$ (right).

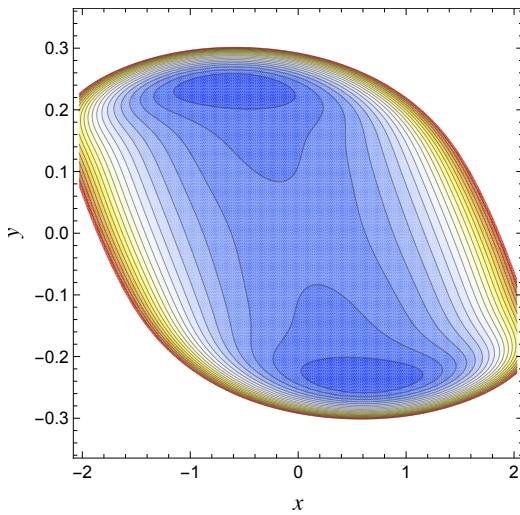


Figure 5. Quartic model: contour plot of the energy learned in the energy-based model.

To make a semi-quantitative comparison, we construct the marginal distributions

$$p_x(x) = \int dy p_\theta(\mathbf{x}), \quad p_y(y) = \int dx p_\theta(\mathbf{x}), \quad p_\theta(\mathbf{x}) = \frac{\exp[-E_\theta(\mathbf{x})]}{\int dx dy \exp[-E_\theta(\mathbf{x})]}, \quad (4.4)$$

by numerical integration. In figure 6 we compare these distribution with those obtained by the CL process and by direct sampling of the energy-based model, collecting the data in histograms in the latter cases. We observe that the three distributions are in good agreement, with small deviations reflecting fluctuations in ensembles with a finite sample size.

To make a final quantitative comparison, we evaluate moments μ_n and cumulants κ_n , starting from

$$\mu_n = \mathbb{E}[z^n] = \int dz \rho(z) z^n \stackrel{?}{=} \int dx dy p(x, y) (x + iy)^n. \quad (4.5)$$

Cumulants follow in the usual way [48]. In tables 1 and 2 we present the results for μ_n and κ_n respectively, with $n = 2, 4, 6, 8$. In each case the first line is the exact result. The second line is obtained using CL sampling and includes a statistical error. We remind the

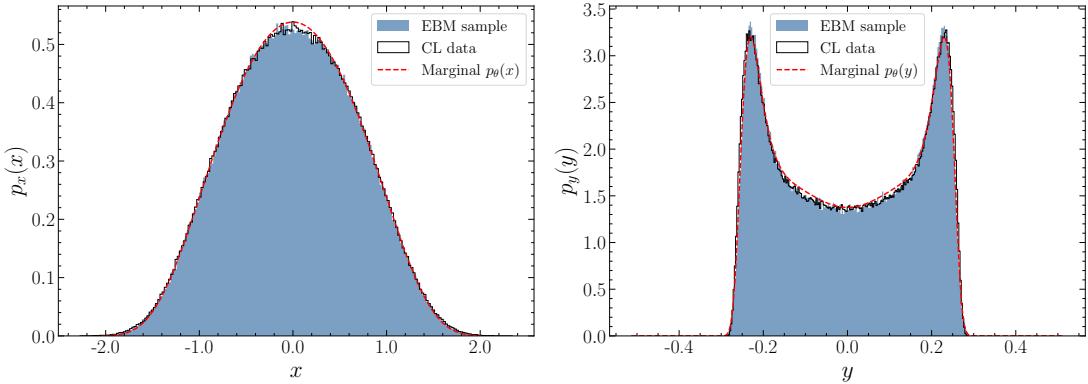


Figure 6. Marginalised distributions $p_x(x)$ and $p_y(y)$ in the quartic model obtained from $E_\theta(\mathbf{x})$ and by direct sampling using the CL process (data) and the energy-based model (EBM sample).

	Re μ_2	Im μ_2	Re μ_4	Im μ_4
Exact	0.428142	-0.148010	0.423848	-0.280132
CL	0.4281(5)	-0.1481(2)	0.4232(11)	-0.2798(6)
SBM	0.4259(2)(5)	-0.1473(1)(3)	0.4237(4)(14)	-0.2777(2)(9)
EBM	0.4264(1)(37)	-0.1487(1)(15)	0.4192(2)(61)	-0.2795(1)(39)
MCMC-EBM	0.4254(2)(75)	-0.1497(1)(31)	0.4169(3)(122)	-0.2802(2)(80)
	Re μ_6	Im μ_6	Re μ_8	Im μ_8
Exact	0.580445	-0.587746	0.95105	-1.39336
CL	0.5787(26)	-0.5866(18)	0.9482(87)	-1.3901(84)
SBM	0.594(1)(4)	-0.5882(1)(3)	1.031(2)(14)	-1.435(3)(14)
EBM	0.569(1)(12)	-0.5834(4)(11)	0.918(2)(26)	-1.374(2)(32)
MCMC-EBM	0.565(1)(24)	-0.584(1)(22)	0.912(2)(49)	-1.377(2)(61)

Table 1. Moments μ_n for the quartic model with parameters $\sigma_0 = 1 + i$, $\lambda = 1$, evaluated using complex Langevin (CL) dynamics, with statistical errors, score-based (SBM) and energy-based (EBM) diffusion models, with statistical and systematic errors, and an MCMC computation employing the energy learned using the EBM, with statistical errors and systematic errors.

reader that the CL process provides the training data and is therefore the benchmark for the comparison. The third and fourth lines are obtained using trained score-based and energy-based diffusion models. To estimate the systematic uncertainty, we follow the procedure outlined in appendix D, creating an ensemble of $N = 10$ independently trained models, with the same set of (near-)optimal hyperparameters. The first (second) number between the brackets represents the statistical (systematic) uncertainty. The systematic uncertainty dominates, reflecting the stochastic nature of the training procedure. Errors in the moments and cumulants are determined independently, using an equal number of bootstrap samples, and cumulants have been constructed from central moments.

The final line is obtained by taking the energy parametrisation $E_\theta(\mathbf{x})$ in the energy-based model and perform a Metropolis-Hastings Markov Chain Monte Carlo (MCMC) simulation using the learned energy in the accept/reject criterion. Again we used 10 trained EBMs, and combine the outcomes to yield an estimate of the systematic uncertainty due to the learned energies. We emphasise that this approach was not feasible before, as the probability distribution was not expressed in terms of an energy function $E_\theta(\mathbf{x})$, which can be evaluated

	Re κ_2	Im κ_2	Re κ_4	Im κ_4
Exact	0.428142	-0.148010	-0.060347	0.100083
CL	0.4280(5)	-0.1480(2)	-0.0606(6)	0.1003(5)
SBM	0.4259(2)(6)	-0.1473(1)(3)	-0.0554(2)(4)	0.0986(1)(4)
EBM	0.4273(2)(9)	-0.1478(1)(2)	-0.0607(2)(3)	0.1001(1)(4)
MCMC-EBM	0.4254(2)(76)	-0.1497(1)(31)	-0.0602(2)(44)	0.1030(1)(62)
	Re κ_6	Im κ_6	Re κ_8	Im κ_8
Exact	-0.00934	-0.19222	0.41578	0.5923
CL	-0.009(1)	-0.194(2)	0.414(5)	0.60(1)
SBM	-0.0131(4)(7)	-0.1863(6)(11)	0.423(2)(6)	0.557(3)(4)
EBM	-0.0102(4)(6)	-0.193(1)(2)	0.422(2)(5)	0.594(3)(8)
MCMC-EBM	-0.0124(4)(32)	-0.205(1)(20)	0.468(2)(51)	0.661(4)(81)

Table 2. As in table 1, for the cumulants κ_n .

at arbitrary \mathbf{x} within the data domain. We note that in this approach, after creating the data set and training the energy-based model, new configurations are created without any reference to the CL dynamics or the diffusion model, i.e., only the energy function is required. This opens the door to apply improved MCMC algorithms to theories with a sign problems eventually and should be explored further.

Overall we observe good agreement between the various methods, with fluctuations representing mostly systematic uncertainties. It goes without saying that since the original weight $\rho(z)$ is complex-valued, we cannot employ a Metropolis step based on the original theory, as in the Metropolis-adjusted Langevin algorithm (MALA) [62].

5 Summary and outlook

We have used diffusion models to study the distribution sampled by a complex Langevin process. Since this distribution is not known *a priori*, this provides a new perspective on understanding CL dynamics for theories with a sign problem. We have noted that both score-based (SBM) and energy-based (EBM) diffusion models can learn the score, i.e., the log derivative of the distribution, and can subsequently be used to generate additional configurations. However, in score-based models the learned score is not conservative. While this does not affect the generative power, it means that the distribution itself cannot be obtained by integration. We found that the CL drift and the SBM and EBM scores all differ (the latter only due to the non-conservative component in the SBM), but yield consistent values for observables.

Energy-based models give direct access to the (unnormalised) distribution via $p_\theta(\mathbf{x}) \sim \exp[-E_\theta(\mathbf{x})]$. One possible application is a detailed study of properties of the distribution in particular in regions where CL encounters problems. A second application, which we have explored here, is to use the learned energy in an MCMC simulation. In such a setup, CL is used to provide training data and the diffusion model is used to learn the energy. Afterwards, an MCMC simulation is used to generate additional configurations, without reference to the original CL process or diffusion model. This setup should be explored further. It goes without saying that such an approach is only useful when the original CL process converges correctly, i.e., the diffusion models as employed here will not solve the sign problem when

CL fails. So far we have explored these aspects in a simple model. The obvious next step is to extend this to field theories, for which all ingredients are in place.

Acknowledgments

GA thanks KZ and his group for the kind hospitality at CUHK-Shenzhen during the completion of this work. This visit was supported in part by the Royal Society International Exchanges 2024 Global Round 2 IES\R2\242215. GA is further supported by STFC Consolidated Grant ST/X000648/1. DEH is supported by the UKRI AIMLAC CDT EP/S023992/1. LW thanks the DEEP-IN working group at RIKEN-iTHEMS for its support in the preparation of this paper. LW is also supported by the RIKEN TRIP initiative (RIKEN Quantum), JSPS KAKENHI Grant No. 25H01560, and JST-BOOST Grant No. JPMJBY24H9. KZ is supported by the CUHK-Shenzhen university development fund under grant No. UDF01003041 and UDF03003041, Shenzhen Peacock fund under No. 2023TC0179 and NSFC grant under No. 92570117. We acknowledge the support of the Supercomputing Wales and AccelerateAI projects, which are part-funded by the European Regional Development Fund (ERDF) via Welsh Government.

A Computational details

Training data is generated by solving the discretised Langevin process with stepsize ϵ using the higher-order algorithm of ref. [61], first applied to CL dynamics in ref. [10]. This algorithm improves stepsize corrections from $\mathcal{O}(\epsilon)$ to $\mathcal{O}(\epsilon^{3/2})$. For both the complex-valued quartic model and the real-valued Gaussian mixture, we have generated an ensemble of 10^6 configurations for training, which are preprocessed by scaling to zero mean and unit variance.

In the diffusion model, we employ a variance-expanding scheme, in which the drift term in eq. (3.2) is set to zero, $K(\mathbf{x}, t) = 0$. Note that due to the complexification, we have two degrees of freedom to consider. We choose the diffusion coefficient $g(t) = \sigma^{t/T}$ and pick $\sigma = 10$ and $T = 1$. We choose to run the backward process using 1000 steps for 10^6 trajectories to obtain samples. Our choice of hyperparameters is summarised in table 3. More details can be found in ref. [48].

We denote by $\mathbf{v}_\theta(\mathbf{x}, t) : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ a time-conditioned neural network for which the output has the same dimensions as the input data for every $t \in [0, T]$. In score-based diffusion models, $\mathbf{v}_\theta(\mathbf{x}, t) = \mathbf{s}_\theta(\mathbf{x}, t)$ directly, while in energy-based models, \mathbf{v}_θ is a trainable function which parametrises an energy functional [60] to approximate the log-likelihood of the target distribution, see eqs. (3.15, 3.17). Using this formulation we may follow a similar prescription in training the two schemes, while being careful with the choice of hyperparameters. In particular, since in energy-based models $E_\theta(\mathbf{x}, t)$ is differentiated to obtain the score, $\mathbf{v}_\theta(\mathbf{x}, t)$ should be sufficiently smooth. Hence we have used the SiLU activation function in EBMs and the LeakyReLU for the SBMs. Concerning the architecture for \mathbf{v}_θ , we employed time-conditioned feed-forward neural networks using a Gaussian Fourier feature mapping [64] and incorporating the \mathbb{Z}_2 symmetry of the data/target distribution, following ref. [65]. For the weights, we keep exponential moving averages (EMA) [66], which are used during sampling.

Finally, in the energy-based model, we perform a Metropolis-Hastings MCMC using the energy parametrisation $E_\theta(\mathbf{x}, 0)$ in the accept/reject criterion. The update was tuned to

Hyperparameter	Value SBM / EBM	Hyperparameter	Value
Layers	[64, 64, 64] / [64, 64]	Learning Rate	1e-4
Time Embedding dims	256 / 256	Batch Size	512
Activation Function	Leaky ReLU / SiLU	Optimizer	Adam
Weight Initialization	LeCun Uniform [63]	Max Epochs	300

Table 3. Model and hyperparameters used in training the score-based (SBM) and energy-based (EBM) diffusion models for the complex-valued quartic theory. Weights with the best loss are saved during the training process and we employ early stopping. For the Gaussian mixture, the same setup is used, with two layers of 32 nodes and a batch size of 1024.

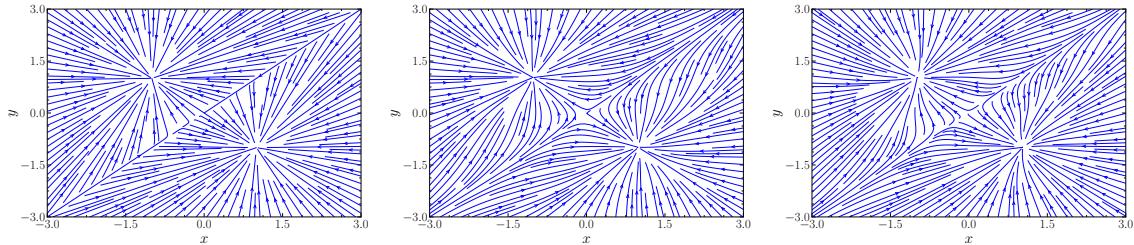


Figure 7. Score in the Gaussian mixture with parameters $\mathbf{x}_0 = (1, -1)$ and $\sigma_0^2 = 1/16$: exact (left), and using a score-based (middle) and energy-based (right) diffusion model.

70–80% acceptance rate, yielding about 10^6 independent configurations. The same approach is followed for direct MCMC simulations in the Gaussian mixture.

B Real-valued Gaussian mixture

To verify that our findings are not due to the complexification, we consider a real target distribution in two dimensions, namely the following Gaussian mixture,

$$p_0(\mathbf{x}) = \frac{1}{2} \left[\mathcal{N}(\mathbf{x}; \mathbf{x}_0, \sigma_0^2 \mathbf{1}) + \mathcal{N}(\mathbf{x}; -\mathbf{x}_0, \sigma_0^2 \mathbf{1}) \right]. \quad (\text{B.1})$$

Note that $p_0(\mathbf{x}) = p_0(-\mathbf{x})$, as in the case considered above. In the numerical experiments, we take $\mathbf{x}_0 = (x_0, y_0) = (1, -1)$ and $\sigma_0^2 = 1/16$ throughout. We use the real-valued Langevin drift

$$\mathbf{K}(\mathbf{x}) = \nabla \log p_0(\mathbf{x}) = -\frac{\mathbf{x}}{\sigma_0^2} + \frac{\mathbf{x}_0}{\sigma_0^2} \tanh\left(\frac{x_0 x + y_0 y}{\sigma^2}\right) \quad (\text{B.2})$$

to generate configurations using Langevin dynamics. Subsequently, we train the score-based and energy-based diffusion models on this data distribution and generate new configurations.

In figure 7 we show the exact score (B.2), and the approximate scores obtained by the score-based and energy-based diffusion models at the end of the backward process. Unlike in the case of the complex distribution, these vector fields have the same structure, as expected, with some deviations from the exact one visible by eye. In figure 8 we present the exact distribution $p(\mathbf{x})$ (left) and the one obtained in the energy-based model (middle), by exponentiating the energy, $p_\theta(\mathbf{x}) \sim \exp[-E_\theta(\mathbf{x})]$, as before. The prefactor is obtained by normalising the distribution to 1 between the boundaries shown, i.e., $-3 \leq x, y \leq 3$. The score obtained in the score-based model is again non-conservative, and its curl is shown in figure 8 (right). Hence it cannot be integrated to obtain the distribution itself.

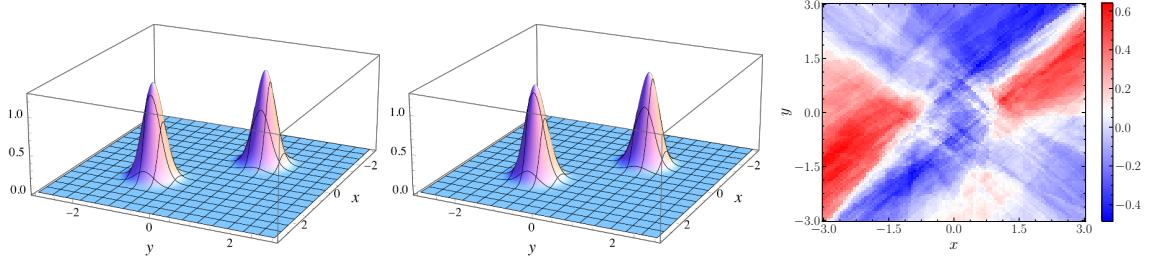


Figure 8. Gaussian mixture: exact distribution (left) and distribution obtained in the energy-based diffusion model (middle). Curl of the score in the score-based formulation, averaged over 10 independently trained models.

	$\langle x^2 \rangle$	$\langle y^2 \rangle$	$\langle x^4 \rangle$	$\langle y^4 \rangle$
Exact	1.0625	1.0625	1.386719	1.386719
Data	1.0620(5)	1.06215(5)	1.3860(13)	1.3846(13)
SBM	1.0632(2)(7)	1.0638(2)(10)	1.3954(4)(27)	1.3974(4)(28)
EBM	1.0621(2)(9)	1.0621(2)(6)	1.3846(4)(22)	1.3849(4)(15)
MCMC-EBM	1.0721(2)(122)	1.0762(2)(155)	1.4237(4)(310)	1.4350(4)(438)
MCMC-MH	1.0630(5)	1.0620(4)	1.3886(14)	1.3861(12)
	$\langle x^6 \rangle$	$\langle y^6 \rangle$	$\langle x^8 \rangle$	$\langle y^8 \rangle$
Exact	2.11694	2.11694	3.67445	3.67445
Data	2.116(4)	2.113(3)	3.674(8)	3.667(9)
SBM	2.149(1)(8)	2.155(1)(7)	3.778(3)(21)	3.792(3)(19)
EBM	2.110(1)(5)	2.111(1)(4)	3.655(3)(13)	3.659(3)(10)
MCMC-EBM	2.218(1)(71)	2.246(1)(96)	3.933(3)(161)	4.003(3)(231))
MCMC-MH	2.122(3)	2.116(4)	3.692(8)	3.672(8)

	$\langle xy \rangle$	$\langle x^2 y^2 \rangle$	$\langle x^3 y + xy^3 \rangle$	$\langle x^3 y^3 \rangle$	$\langle x^4 y^2 + x^2 y^4 \rangle$
Exact	-1	1.12891	-2.375	-1.41016	2.94678
Data	-1.0007(4)	1.1306(8)	-2.3785(17)	-1.4133(15)	2.9535(35)
SBM	-0.9997(1)(6)	1.1327(3)(16)	-2.3848(5)(34)	-1.4238(5)(36)	2.977(1)(7)
EBM	-1.0000(1)(6)	1.1285(3)(12)	-2.3737(5)(24)	-1.4087(5)(21)	2.943(1)(5)
MCMC-EBM	-1.0082(1)(116)	1.1543(3)(255)	-2.4331(5)(558)	-1.4652(5)(483)	3.066(1)(102)
MCMC-MH	-0.9999(3)	1.1290(7)	-2.3755(17)	-1.4110(16)	2.9492(31)
	$\langle x^5 y + xy^5 \rangle$	$\langle x^4 y^4 \rangle$	$\langle x^5 y^3 + x^3 y^5 \rangle$	$\langle x^6 y^2 + x^2 y^6 \rangle$	$\langle xy^7 + x^7 y \rangle$
Exact	-3.36719	1.92299	-3.99854	4.4985	-5.49658
Data	-3.373(4)	1.928(3)	-4.009(7)	4.510(8)	-5.509(9)
SBM	-3.409(1)(9)	1.960(1)(7)	-4.076(2)(15)	4.593(2)(19)	-5.630(3)(25)
EBM	-3.361(1)(5)	1.920(1)(4)	-3.991(2)(8)	4.488(2)(8)	-5.479(3)(11)
MCMC-EBM	-3.521(1)(120)	2.031(1)(87)	-4.228(2)(181)	4.776(2)(209)	-5.876(3)(266)
MCMC-MH	-3.370(4)	1.925(3)	-4.004(6)	4.504(7)	-5.506(8)

Table 4. Various moments for the Gaussian mixture with parameters $\sigma_0 = 1/4, \mu_0 = 1$. As in table 1, with the addition of the final line, using direct MCMC Metropolis-Hastings sampling from the target distribution, with statistical error.

In table 4 we show the results for the moments, observing good agreement. In this case it is possible to perform an MCMC simulation directly from the real probability distribution

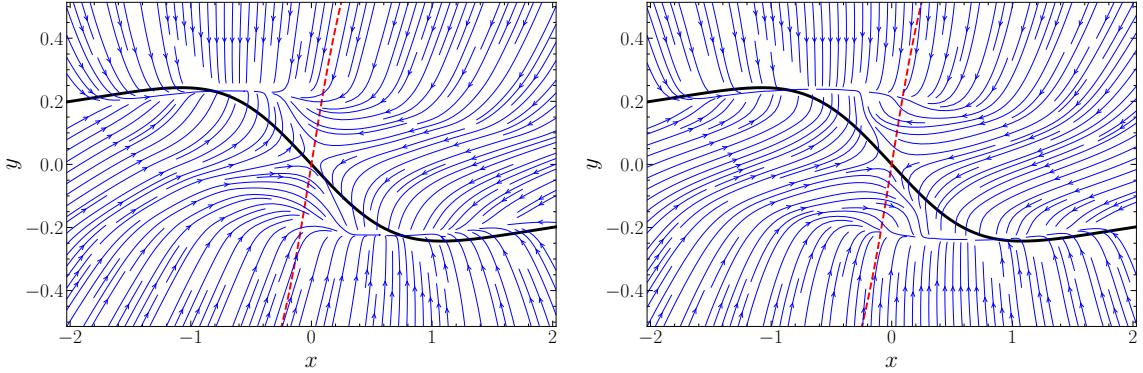


Figure 9. Stable (full) and unstable (thin dashed) thimbles superimposed on the drift determined in the score-based (left) and energy-based (right) diffusion models.

$p_0(\mathbf{x})$, allowing a comparison with the MCMC simulation based on the EBM. These results are given in the final line.

C Lefschetz thimbles

Lefschetz thimbles are an alternative approach to enter the complexified manifold and potentially evade the sign problem [6]. For the quartic model considered here, the stable and unstable thimbles were derived analytically in ref. [67]. For completeness, we show in figure 9 the thimbles superimposed on the drifts obtained in the score-based and energy-based diffusion models. We observe that the stable thimble follows in particular the attractive region near the top of the ridges. As already discussed in ref. [67], in the thimble case the component of the drift in the imaginary direction is reversed and the critical point at the origin is attractive along the stable thimble. Hence the weight along the thimble peaks at the origin and not at the ridges.

D Estimating systematic effects

After determining the (near-)optimal choice of hyperparameters, we evaluate the uncertainty for the diffusion models using $N = 10$ models, which are seeded differently and trained independently, but with the same choice of hyperparameters. Consider an observable O . For each model we estimate O_i , labelled by the index of the model, $i = 1, \dots, N$, with some statistical uncertainty δO_i determined using a jackknife analysis. The goal is to quantify the contribution of per-run statistical errors and estimate a (between-run) systematic error.

For this we apply a random-effects meta-analysis known as DerSimonian-Laird [68, 69]. Each model is weighted by

$$w_i = \frac{1}{\delta O_i^2}, \quad (\text{D.1})$$

with a fixed-effect weighted mean,

$$\bar{O} = \frac{\sum_i w_i O_i}{\sum_i w_i}. \quad (\text{D.2})$$

The excess scatter across models is quantified by a heterogeneity statistic such as Cochran's Q [70],

$$Q = \sum_i w_i \left(O_i - \bar{O} \right)^2. \quad (\text{D.3})$$

A large Q indicates variability beyond statistical deviation.

If we define the effective information quantity C by

$$C = \sum_i w_i - \frac{\sum_i w_i^2}{\sum_i w_i}, \quad (\text{D.4})$$

then the between-model variance is estimated by

$$\tau^2 = \max \left(0, \frac{Q - (N - 1)}{C} \right). \quad (\text{D.5})$$

Adjusting the weights for random effects,

$$w_i^* = \frac{1}{\delta O_i^2 + \tau^2}, \quad (\text{D.6})$$

the adjusted weighted mean is

$$O^* = \frac{\sum_i w_i^* O_i}{\sum_i w_i^*}. \quad (\text{D.7})$$

Let

$$a_i = \frac{w_i^*}{\sum_i w_i^*}. \quad (\text{D.8})$$

The contributions to the error are then finally broken down as

- Statistical contribution: $\sigma_{\text{stat}}^2 = \sum_i (a_i \delta O_i)^2$.
- Systematic contribution: $\sigma_{\text{sys}}^2 = \tau^2 \sum_i a_i^2$.
- Total random-effects uncertainty: $\sigma_{\text{total}}^2 = 1 / \sum_i w_i^*$.

In tables 1, 2 and 4 we quote $O^*(\sigma_{\text{stat}})(\sigma_{\text{sys}})$ for the score- and energy-based models.

Data Availability Statement. The data used for this manuscript is available in ref. [71].

Code Availability Statement. The code used for this manuscript is available in ref. [71].

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] M. Troyer and U.-J. Wiese, *Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations*, *Phys. Rev. Lett.* **94** (2005) 170201 [[cond-mat/0408370](#)] [[INSPIRE](#)].
- [2] P. de Forcrand, *Simulating QCD at finite density*, *PoS LAT2009* (2009) 010 [[arXiv:1005.0539](#)] [[INSPIRE](#)].
- [3] G. Aarts, *Introductory lectures on lattice QCD at nonzero baryon number*, *J. Phys. Conf. Ser.* **706** (2016) 022004 [[arXiv:1512.05145](#)] [[INSPIRE](#)].
- [4] G. Parisi, *On complex probabilities*, *Phys. Lett. B* **131** (1983) 393 [[INSPIRE](#)].
- [5] J.R. Klauder, *Stochastic quantization*, in *Recent Developments in High-Energy Physics*, H. Mitter and C.B. Lang eds., Springer, Vienna, Austria (1983), pg. 251.
- [6] AURORASCIENCE collaboration, *New approach to the sign problem in quantum field theories: High density QCD on a Lefschetz thimble*, *Phys. Rev. D* **86** (2012) 074506 [[arXiv:1205.3996](#)] [[INSPIRE](#)].
- [7] A. Alexandru et al., *Sign problem and Monte Carlo calculations beyond Lefschetz thimbles*, *JHEP* **05** (2016) 053 [[arXiv:1512.08764](#)] [[INSPIRE](#)].
- [8] G. Parisi and Y.-S. Wu, *Perturbation Theory Without Gauge Fixing*, *Sci. Sin.* **24** (1981) 483 [[INSPIRE](#)].
- [9] P.H. Damgaard and H. Huffel, *Stochastic Quantization*, *Phys. Rept.* **152** (1987) 227 [[INSPIRE](#)].
- [10] G. Aarts and F.A. James, *Complex Langevin dynamics in the $SU(3)$ spin model at nonzero chemical potential revisited*, *JHEP* **01** (2012) 118 [[arXiv:1112.4655](#)] [[INSPIRE](#)].
- [11] 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](#)].
- [12] 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](#)].
- [13] D. Sexty, *Simulating full QCD at nonzero density using the complex Langevin equation*, *Phys. Lett. B* **729** (2014) 108 [[arXiv:1307.7748](#)] [[INSPIRE](#)].
- [14] G. Aarts, F. Attanasio, B. Jäger and D. Sexty, *The QCD phase diagram in the limit of heavy quarks using complex Langevin dynamics*, *JHEP* **09** (2016) 087 [[arXiv:1606.05561](#)] [[INSPIRE](#)].
- [15] D. Sexty, *Calculating the equation of state of dense quark-gluon plasma using the complex Langevin equation*, *Phys. Rev. D* **100** (2019) 074503 [[arXiv:1907.08712](#)] [[INSPIRE](#)].
- [16] M. Scherzer, D. Sexty and I.O. Stamatescu, *Deconfinement transition line with the complex Langevin equation up to $\mu/T \sim 5$* , *Phys. Rev. D* **102** (2020) 014515 [[arXiv:2004.05372](#)] [[INSPIRE](#)].
- [17] Y. Ito et al., *Complex Langevin calculations in QCD at finite density*, *JHEP* **10** (2020) 144 [[arXiv:2007.08778](#)] [[INSPIRE](#)].
- [18] S. Tsutsui et al., *On the validity of the complex Langevin method near the deconfining phase transition in QCD at finite density*, *JHEP* **10** (2025) 108 [[arXiv:2505.06551](#)] [[INSPIRE](#)].
- [19] J. Ambjorn and S.K. Yang, *Numerical Problems in Applying the Langevin Equation to Complex Effective Actions*, *Phys. Lett. B* **165** (1985) 140 [[INSPIRE](#)].

[20] G. Aarts and F.A. James, *On the convergence of complex Langevin dynamics: The Three-dimensional XY model at finite chemical potential*, *JHEP* **08** (2010) 020 [[arXiv:1005.3468](https://arxiv.org/abs/1005.3468)] [[INSPIRE](#)].

[21] G. Aarts et al., *Stability of complex Langevin dynamics in effective models*, *JHEP* **03** (2013) 073 [[arXiv:1212.5231](https://arxiv.org/abs/1212.5231)] [[INSPIRE](#)].

[22] 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](https://arxiv.org/abs/0912.3360)] [[INSPIRE](#)].

[23] 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](https://arxiv.org/abs/1101.3270)] [[INSPIRE](#)].

[24] J. Nishimura and S. Shimasaki, *New Insights into the Problem with a Singular Drift Term in the Complex Langevin Method*, *Phys. Rev. D* **92** (2015) 011501 [[arXiv:1504.08359](https://arxiv.org/abs/1504.08359)] [[INSPIRE](#)].

[25] M. Scherzer, E. Seiler, D. Sexty and I.-O. Stamatescu, *Complex Langevin and boundary terms*, *Phys. Rev. D* **99** (2019) 014512 [[arXiv:1808.05187](https://arxiv.org/abs/1808.05187)] [[INSPIRE](#)].

[26] M. Scherzer, E. Seiler, D. Sexty and I.-O. Stamatescu, *Controlling Complex Langevin simulations of lattice models by boundary term analysis*, *Phys. Rev. D* **101** (2020) 014501 [[arXiv:1910.09427](https://arxiv.org/abs/1910.09427)] [[INSPIRE](#)].

[27] D. Alvestad, R. Larsen and A. Rothkopf, *Towards learning optimized kernels for complex Langevin*, *JHEP* **04** (2023) 057 [[arXiv:2211.15625](https://arxiv.org/abs/2211.15625)] [[INSPIRE](#)].

[28] K. Boguslavski, P. Hotzy and D.I. Müller, *Stabilizing complex Langevin for real-time gauge theories with an anisotropic kernel*, *JHEP* **06** (2023) 011 [[arXiv:2212.08602](https://arxiv.org/abs/2212.08602)] [[INSPIRE](#)].

[29] E. Seiler, D. Sexty and I.-O. Stamatescu, *Complex Langevin: Correctness criteria, boundary terms, and spectrum*, *Phys. Rev. D* **109** (2024) 014509 [[arXiv:2304.00563](https://arxiv.org/abs/2304.00563)] [[INSPIRE](#)].

[30] K. Boguslavski, P. Hotzy and D.I. Müller, *Real-time correlators in 3+1D thermal lattice gauge theory*, *Phys. Rev. D* **109** (2024) 094518 [[arXiv:2312.03063](https://arxiv.org/abs/2312.03063)] [[INSPIRE](#)].

[31] M.W. Hansen and D. Sexty, *Testing dynamical stabilization of complex Langevin simulations of QCD*, *Phys. Rev. D* **111** (2025) 054508 [[arXiv:2405.20709](https://arxiv.org/abs/2405.20709)] [[INSPIRE](#)].

[32] K. Boguslavski, P. Hotzy and D.I. Müller, *Lefschetz thimble-inspired weight regularizations for complex Langevin simulations*, *SciPost Phys.* **18** (2025) 092 [[arXiv:2412.02396](https://arxiv.org/abs/2412.02396)] [[INSPIRE](#)].

[33] M. Mandl, E. Seiler and D. Sexty, *Necessary and sufficient conditions for correctness of complex Langevin*, *J. Phys. A* **58** (2025) 495202 [[arXiv:2508.14512](https://arxiv.org/abs/2508.14512)] [[INSPIRE](#)].

[34] G. Aarts, *Complex Langevin dynamics at finite chemical potential: Mean field analysis in the relativistic Bose gas*, *JHEP* **05** (2009) 052 [[arXiv:0902.4686](https://arxiv.org/abs/0902.4686)] [[INSPIRE](#)].

[35] G. Aarts, P. Giudice and E. Seiler, *Localised distributions and criteria for correctness in complex Langevin dynamics*, *Annals Phys.* **337** (2013) 238 [[arXiv:1306.3075](https://arxiv.org/abs/1306.3075)] [[INSPIRE](#)].

[36] D. Weingarten, *Complex probabilities on R^*N as real probabilities on $C^{**}N$ and an application to path integrals*, *Phys. Rev. Lett.* **89** (2002) 240201 [[quant-ph/0210195](https://arxiv.org/abs/quant-ph/0210195)] [[INSPIRE](#)].

[37] L.L. Salcedo and E. Seiler, *Schwinger-Dyson equations and line integrals*, *J. Phys. A* **52** (2019) 035201 [[arXiv:1809.06888](https://arxiv.org/abs/1809.06888)] [[INSPIRE](#)].

[38] J. Sohl-Dickstein, E.A. Weiss, N. Maheswaranathan and S. Ganguli, *Deep Unsupervised Learning using Nonequilibrium Thermodynamics*, [arXiv:1503.03585](https://arxiv.org/abs/1503.03585) [[INSPIRE](#)].

[39] J. Ho, A. Jain and P. Abbeel, *Denoising Diffusion Probabilistic Models*, [arXiv:2006.11239](https://arxiv.org/abs/2006.11239) [[INSPIRE](#)].

[40] Y. Song and S. Ermon, *Generative Modeling by Estimating Gradients of the Data Distribution*, [arXiv:1907.05600](https://arxiv.org/abs/1907.05600).

[41] Y. Song et al., *Score-Based Generative Modeling through Stochastic Differential Equations*, [arXiv:2011.13456](https://arxiv.org/abs/2011.13456) [INSPIRE].

[42] L. Yang et al., *Diffusion Models: A Comprehensive Survey of Methods and Applications*, [arXiv:2209.00796](https://arxiv.org/abs/2209.00796).

[43] Y. Song, C. Durkan, I. Murray and S. Ermon, *Maximum Likelihood Training of Score-Based Diffusion Models*, [arXiv:2101.09258](https://arxiv.org/abs/2101.09258).

[44] T. Karras, M. Aittala, T. Aila and S. Laine, *Elucidating the Design Space of Diffusion-Based Generative Models*, [arXiv:2206.00364](https://arxiv.org/abs/2206.00364).

[45] L. Wang, G. Aarts and K. Zhou, *Diffusion models as stochastic quantization in lattice field theory*, *JHEP* **05** (2024) 060 [[arXiv:2309.17082](https://arxiv.org/abs/2309.17082)] [INSPIRE].

[46] L. Wang, G. Aarts and K. Zhou, *Generative Diffusion Models for Lattice Field Theory*, in the proceedings of the *37th Conference on Neural Information Processing Systems*, New Orleans, U.S.A. (2023) [[arXiv:2311.03578](https://arxiv.org/abs/2311.03578)] [INSPIRE].

[47] Q. Zhu et al., *Physics-Conditioned Diffusion Models for Lattice Gauge Theory*, [arXiv:2502.05504](https://arxiv.org/abs/2502.05504) [INSPIRE].

[48] G. Aarts, D.E. Habibi, L. Wang and K. Zhou, *On learning higher-order cumulants in diffusion models*, *Mach. Learn. Sci. Tech.* **6** (2025) 025004 [[arXiv:2410.21212](https://arxiv.org/abs/2410.21212)] [INSPIRE].

[49] J. Cotler and S. Reznikov, *Renormalizing Diffusion Models*, [arXiv:2308.12355](https://arxiv.org/abs/2308.12355) [INSPIRE].

[50] Y. Hirono, A. Tanaka and K. Fukushima, *Understanding Diffusion Models by Feynman’s Path Integral*, [arXiv:2403.11262](https://arxiv.org/abs/2403.11262) [INSPIRE].

[51] K. Fukushima, S. Kamata and Y. Hirono, *Stochastic Quantization and Diffusion Models*, *J. Phys. Soc. Jap.* **94** (2025) 031010 [[arXiv:2411.11297](https://arxiv.org/abs/2411.11297)] [INSPIRE].

[52] M. Gerdes, M. Welling and M.C.N. Cheng, *GUD: Generation with Unified Diffusion*, [arXiv:2410.02667](https://arxiv.org/abs/2410.02667) [INSPIRE].

[53] D.E. Habibi, G. Aarts, L. Wang and K. Zhou, *Diffusion models learn distributions generated by complex Langevin dynamics*, *PoS LATTICE2024* (2025) 039 [[arXiv:2412.01919](https://arxiv.org/abs/2412.01919)] [INSPIRE].

[54] J. Berges and D. Sexty, *Real-time gauge theory simulations from stochastic quantization with optimized updating*, *Nucl. Phys. B* **799** (2008) 306 [[arXiv:0708.0779](https://arxiv.org/abs/0708.0779)] [INSPIRE].

[55] G. Aarts and I.-O. Stamatescu, *Stochastic quantization at finite chemical potential*, *JHEP* **09** (2008) 018 [[arXiv:0807.1597](https://arxiv.org/abs/0807.1597)] [INSPIRE].

[56] B.D.O. Anderson, *Reverse-time diffusion equation models*, *Stochastic Processes and their Applications* **12** (1982) 313.

[57] S. Saremi, *On approximating ∇f with neural networks*, [arXiv:1910.12744](https://arxiv.org/abs/1910.12744).

[58] T. Salimans and J. Ho, *Should EBMs model the energy or the score?*, in the proceedings of the *International Conference on Learning Representations (ICLR 2021)*, Vienna, Austria (2021).

[59] C. Horvat and J.-P. Pfister, *On gauge freedom, conservativity and intrinsic dimensionality estimation in diffusion models*, [arXiv:2402.03845](https://arxiv.org/abs/2402.03845).

[60] Y. Du et al., *Reduce, Reuse, Recycle: Compositional Generation with Energy-Based Diffusion Models and MCMC*, in the proceedings of the *International Conference on Learning Representations (ICLR 2023)*, Kigali, Rwanda (2023) [[arXiv:2302.11552](https://arxiv.org/abs/2302.11552)].

- [61] C.C. Chang, *Numerical solution of stochastic differential equations with constant diffusion coefficients*, *Math. Comput.* **49** (1987) 523.
- [62] G.O. Roberts and J.S. Rosenthal, *Optimal scaling of discrete approximations to Langevin diffusions*, *J. Roy. Statist. Soc. B* **60** (1998) 255.
- [63] Y. LeCun, L. Bottou, G.B. Orr and K.R. Müller, *Efficient BackProp*, Springer, Heidelberg, Germany (1998) [[DOI:10.1007/3-540-49430-8_2](https://doi.org/10.1007/3-540-49430-8_2)].
- [64] M. Tancik et al., *Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains*, [arXiv:2006.10739](https://arxiv.org/abs/2006.10739).
- [65] M. Mattheakis et al., *Physical Symmetries Embedded in Neural Networks*, [arXiv:1904.08991](https://arxiv.org/abs/1904.08991) [[INSPIRE](https://inspirehep.net/search?p=find+EPRINT+arXiv:1904.08991)].
- [66] Y. Song and S. Ermon, *Improved Techniques for Training Score-Based Generative Models*, [arXiv:2006.09011](https://arxiv.org/abs/2006.09011).
- [67] G. Aarts, *Lefschetz thimbles and stochastic quantization: Complex actions in the complex plane*, *Phys. Rev. D* **88** (2013) 094501 [[arXiv:1308.4811](https://arxiv.org/abs/1308.4811)] [[INSPIRE](https://inspirehep.net/search?p=find+EPRINT+arXiv:1308.4811)].
- [68] R. DerSimonian and N. Laird, *Meta-analysis in clinical trials*, *Control. Clin. Trials* **7** (1986) 177.
- [69] R. DerSimonian and N. Laird, *Meta-analysis in clinical trials revisited*, *Contemp. Clin. Trials* **45** (2015) 139.
- [70] W. Cochran, *The combination of estimates from different experiments*, *Biometrics* **10** (1954) 101.
- [71] D.E. Habibi, G. Aarts, L. Wang and K. Zhou, *Combining complex Langevin dynamics with score-based and energy-based diffusion models — DATA RELEASE & WORKFLOW*, Zenodo, <https://doi.org/10.5281/zenodo.17725665> (2025).