

Adaptive Monte Carlo augmented with normalizing flows

Marylou Gabrié

*Flatiron Institute, New York, NY and
Center for Data Science, New York University, New York, NY**

Grant M. Rotskoff

Dept. of Chemistry, Stanford University, Stanford, CA 94305†

Eric Vanden-Eijnden

Courant Institute, New York University, New York, NY 10012‡

Many problems in the physical sciences, machine learning, and statistical inference necessitate sampling from a high-dimensional, multi-modal probability distribution. Markov Chain Monte Carlo (MCMC) algorithms, the ubiquitous tool for this task, typically rely on random, reversible, and local updates to propagate configurations of a given system in a way that ensures that generated configurations will be distributed according to a target probability distribution asymptotically. In high-dimensional settings with multiple relevant metastable basins, local approaches require either immense computational effort or intricately designed importance sampling strategies to capture information about, for example, the relative populations of such basins. Here we analyze a framework for augmenting MCMC sampling with nonlocal transition kernels parameterized with generative models known as normalizing flows. We focus on a setting where there is no preexisting data, as is commonly the case for problems in which MCMC is used. Our results emphasize that the implementation of the normalizing flow must be adapted to the structure of the target distribution in order to preserve the statistics of the target at all scales. Furthermore, we analyze the propensity of our algorithm to discover new states and demonstrate the importance of initializing the training with some *a priori* knowledge of the relevant modes. We show that our algorithm can sample effectively across large free energy barriers, providing dramatic accelerations relative to traditional MCMC algorithms.

Neural networks approximate high-dimensional functions with a robustness unparalleled by classical function approximation techniques [1, 2]. These favorable approximation properties have motivated a veritable explosion of interest in applying machine learning to otherwise intractable problems in the physical sciences [3]. Monte Carlo sampling methods seem poised to benefit from these tools: slow relaxation and metastability plague sampling problems that arise in chemistry and biophysics [4]. And, indeed, these challenges have stimulated a flurry of work seeking to accelerate sampling using machine learning. How robust are these approaches? What are the prospects for sampling *new* configurations far from any previously known?

There are several distinct strategies that deploy machine learning to accelerate sampling [5–8]. For example, the approach introduced in Ref. [6] parameterizes a transition kernel by propagating samples using Hamiltonian dynamics with a learned Hamiltonian. Here we focus on a direct approach in which a neural network model is trained to generate data that overlaps with data sampled from a given target distribution; this network can then serve as a transition operator in a Metropolis-Hastings MCMC.

The prospect of enhancing sampling with suitable generative models is an active area of inquiry [5–9]. While the term “generative neural network” includes a wide variety of models such as generative adversarial networks [10], variational auto-encoders [11], and energy based models such as restricted Boltzmann machines [12], sampling via Metropolis-Hastings MCMC requires the computation of each transition generation probability of and its inverse. As a result, the model architectures on which most generative neural networks rely are not conducive to Metropolis-Hastings MCMC. However, specific classes of neural networks have been designed with this in mind, allowing for efficient estimates of the probability of a generated sample, including auto-regressive models [13] and normalizing flows [14, 15]. Described in detail below, normalizing flows are expressive invertible function representations that can be trained to transform samples from a simple base measure to mimic samples of a given empirical distribution. At this point, normalizing flow models have been investigated as transition operators in MCMC algorithms in a variety of contexts in the physical sci-

* mgabrie@nyu.edu

† rotskoff@stanford.edu

‡ eve2@cims.nyu.edu

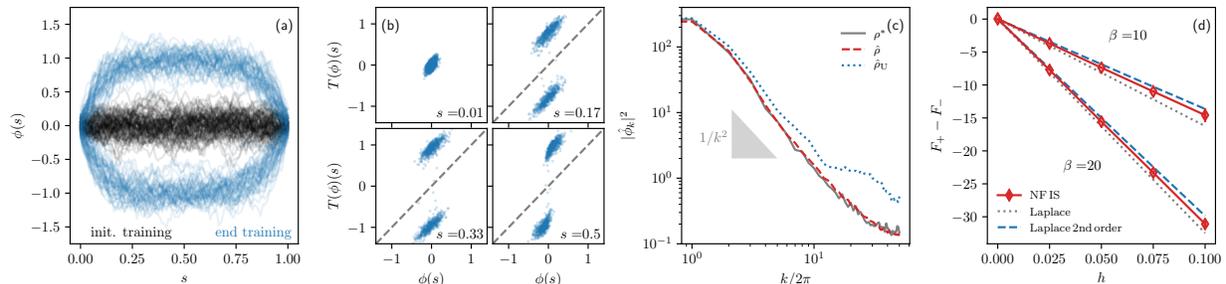


FIG. 1. **Sampling metastable states of the stochastic Allen-Cahn model with Langevin dynamics augmented with a normalizing flow.** (a) Configurations obtained by pushing independent samples from the informed base measure Eq. 19 through the flow T at the beginning (black) and at the end of training (blue). Around $\sim 60\%$ of generated configurations are accepted according to the Metropolis-Hasting criteria. (b) The learned map T is local in space. (c) Fourier spectrum of the target samples, samples from a flow with informed base measure and uniformed base measure. An informed base measure is necessary to capture the higher frequency features of the target density. (d) Computation of the free energy differences between positive and negative modes with importance sampling from the normalizing flow as a function of a local biasing field added in the Hamiltonian Eq. 20. Results are reported for inverse temperature $\beta = 20$, as in the rest of the plots, and for the same experiment repeated at temperature $\beta = 10$. Errors bars computed from estimator variance are smaller than marker.

ences [16–22]. These techniques show significant promise, dramatically reducing the computational expense of sampling when compared to local methods used to generate data, such as molecular dynamics.

The conventional strategy for training a generative model requires a large, representative data set. Arguably, these models have succeeded most dramatically in domains where the cost of generating and curating data is comparatively low (e.g., image recognition) [7, 10, 23]. In scientific computing applications, sampling is necessary to obtain data in the first place. For instance, when evaluating free energies differences of molecular conformations or marginalizing over a Bayesian posterior, the starting point is an unnormalized density and no *a priori* data set exists. In this setting, it is natural to ask whether sampling strategies employing generative models can meaningfully augment our current sampling capabilities which are driven primarily by local dynamics such as Hamiltonian Monte Carlo or Langevin dynamics.

We aim to critically evaluate sampling strategies based on normalizing flows and assess their prospects for accelerating sampling in cases where no extensive preexisting data set is available. We describe an algorithm where training and sampling are performed simultaneously and synergistically. The algorithm augments a local dynamics with the generative sampling and learns how to mix samples efficiently among modes regardless of the dimensionality. We show that finding unknown metastable states is statistically unlikely when using generative Monte Carlo schemes. However, in the setting where configurations representative of the modes of interest in the target distribution are known, we find this augmented approach to training can remedy statistical imbalances among the initial states, leading to accurate estimates of the relative statistical weights of metastable states without constructing a specific pathway between the basins of interest. We analyze this algorithm and show that, in the continuous time limit, it converges asymptotically with an exponential rate that can be explicitly estimated. Our analysis and numerical experiments also emphasize the importance of the normalizing flow base distribution, which must be adapted to the properties of the target distribution in order to ensure robust training.

I. DESIGN CHALLENGES IN MCMC METHODS

The goal of sampling is to generate configurations $x \in \Omega \subset \mathbb{R}^d$ in proportion to some probability measure $\nu_*(dx) = \rho_*(x)dx$ which we assume has probability density function ρ_* . In physical systems, we typically write this in Boltzmann form

$$\rho_*(x) = Z_*^{-1} e^{-U_*(x)} \quad (1)$$

where $U_* \propto -\log \rho_*$ is the potential energy function for the system. Here we assume that we have an explicit model for U_* and can efficiently evaluate this energy function, though we may have little *a priori*

information about the distribution of configurations associated with this energy and in general do not know the normalization constant Z_* .

Markov chain Monte Carlo algorithms avoid computing Z_* by generating a sequence $\{x(k)\}_{k \in \mathbb{N}}$ of configurations with a transition kernel $\pi(x, y)$ with $\int_{\Omega} \pi(x, y) dy = 1$ for all $x \in \Omega$, which quantifies the conditional probability density of a transition from state x into state y . Assume that the kernel $\pi(x, y)$ is irreducible and aperiodic [24], and satisfies the detailed balance relation

$$\rho_*(x)\pi(x, y) = \rho_*(y)\pi(y, x). \quad (2)$$

Then the sequence $\{x(k)\}_{k \in \mathbb{N}}$ will sample the target density ρ_* in the sense that the empirical average of any suitable observable ϕ converges to its expectation over ρ_* , i.e.,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N \phi(x(k)) = \int_{\Omega} \phi(x) \rho_*(x) dx. \quad (3)$$

Designing a transition kernel π leading to fast convergence of the series in [3] is a generically challenging task for MCMC algorithms. In Metropolis-Hastings MCMC one constructs a proposal distribution that creates new samples that are then accepted or rejected according to a criterion that maintains [2]. For example, in the Metropolis adjusted Langevin algorithm (MALA) [25] new configurations are proposed by approximating the solution of the Langevin equation propagated on a fixed time interval.

Metropolis-Hastings MCMC algorithms, however, involve a trade-off between two requirements that are hard to fulfill simultaneously. Proposal distributions using local dynamics like MALA suffer from long decorrelation times when there is metastability in the target density ρ_* . At the same time, seeking faster mixing times with non-local proposal distributions requires careful design to avoid high rejection rates. Recent work in the machine learning literature has suggested a *data-driven* approach to constructing the transition kernel [6–8] that aids in this design challenge; these approaches originally were pioneered in the context of adaptive and nonlinear Monte Carlo algorithms [26–29]. Here, we explore the use of normalizing flows to adaptively parameterize the transition kernel.

II. SAMPLING WITH NORMALIZING FLOWS

A normalizing flow is an invertible map T that is optimized to transport samples from a base measure $\nu_{\text{B}}(dx) = \rho_{\text{B}}(x)dx$ (typically a Gaussian with unit variance) to a given target distribution [15]. The goal is to produce a map T_* with inverse \bar{T}_* such that an expectation of an observable with respect to ρ_* can be estimated by transforming samples from the base density to the target, i.e. if x_{B} is drawn from $\rho_{\text{B}}(x)$ then $T_*(x_{\text{B}})$ is a sample from $\rho_*(x)$ so that for any suitable observable \mathcal{O} we have

$$\int_{\Omega} \mathcal{O}(T_*(x)) \rho_{\text{B}}(x) dx = \int_{\Omega} \mathcal{O}(x) \rho_*(x) dx. \quad (4)$$

Of course, in practice we do not have direct access to this ideal map T_* , and must learn some approximation T . Throughout, we denote the push-forward of ρ_{B} under the map T simply by $\hat{\rho}$: it has the explicit form

$$\hat{\rho}(x) = \rho_{\text{B}}(\bar{T}(x)) \det |\nabla_x \bar{T}|, \quad (5)$$

where \bar{T} denotes the inverse map, i.e. $\bar{T}(T(x)) = T(\bar{T}(x)) = x$. In practice, the parametrization of the map T must be designed carefully to evaluate this density efficiently, requiring easily estimable Jacobian determinants and inverses. This issue has been one of the main foci in the normalizing flow literature [15] and is for instance solved using coupling layers [30, 31]. Even if the map T is not the optimal T_* , i.e. $\hat{\rho}(x) \neq \rho_*(x)$, as long as $\hat{\rho}$ and ρ_* are either both positive or both zero at any point $x \in \Omega$, we can still generate configurations using T with the correct statistical weight in the target distribution by using a Metropolis-Hastings MCMC algorithm with an accept-reject step: a proposed configuration $y = T(x_{\text{B}})$ from a given configuration x is accepted with probability

$$\text{acc}(x, y) = \min \left[1, \frac{\hat{\rho}(x) \rho_*(y)}{\rho_*(x) \hat{\rho}(y)} \right]. \quad (6)$$

Algorithm1 Concurrent MCMC sampling and map training

```

1: SAMPLETRAIN( $U_*$ ,  $T$ ,  $\{x_i(0)\}_{i=1}^n$ ,  $\tau$ ,  $k_{\max}$ ,  $k_{\text{Lang}}$ ,  $\epsilon$ )
2: Inputs:  $U_*$  target energy,  $T$  initial map,  $\{x_i(0)\}_{i=1}^n$  initial data,  $\tau > 0$  time step,  $k_{\max} \in \mathbb{N}$  total duration,
    $k_{\text{Lang}} \in \mathbb{N}$  number of Langevin steps per resampling step,  $\epsilon > 0$  map training time step
3:  $k = 0$ 
4: while  $k < k_{\max}$  do
5:   for  $i = 1, \dots, n$  do
6:     if  $k \bmod k_{\text{Lang}} + 1 = 0$  then
7:        $x'_{\text{B},i} \sim \rho_{\text{B}}$ 
8:        $x'_i = T(x'_{\text{B},i}) \triangleright$  push-forward via  $T$ 
9:        $x_i(k+1) = x'_i$  with probability  $\text{acc}(x_i(k), x'_i)$ , otherwise  $x_i(k+1) = x_i(k) \triangleright$  resampling step
10:    else
11:       $x'_i = x_i(k) - \tau \nabla U_*(x_i(k)) + \sqrt{2\tau} \eta_i$  with  $\eta_i \sim \mathcal{N}(0_d, I_d) \triangleright$  discretized Langevin step
12:       $x_i(k+1) = x'_i$  with MALA acceptance probability or ULA, otherwise  $x_i(k+1) = x_i(k)$ 
13:     $k \leftarrow k + 1$ 
14:     $\mathcal{L}[T] = -\frac{1}{n} \sum_{i=1}^n \log \hat{\rho}(x_i(k+1)) \triangleright$  evaluate  $D_{\text{KL}}(\rho_t \parallel \hat{\rho})$  on sampled data
15:     $T \leftarrow T - \epsilon \nabla \mathcal{L}[T] \triangleright$  Update the map
16: return:  $\{x_i(k)\}_{k=0, i=1}^{k_{\max}, n}$ ,  $T$ 

```

This procedure is equivalent to using the transition kernel

$$\pi_T(x, y) = \text{acc}(x, y) \hat{\rho}(y) + (1 - r(x)) \delta(x - y) \quad (7)$$

where

$$r(x) = \int_{\Omega} \text{acc}(x, y) \hat{\rho}(y) dy. \quad (8)$$

The formula in [6] for the acceptance probability emphasizes that if the generated configurations do not have appreciable statistical weight in the target distribution (i.e. $\rho_*(y)$ is very small) few configurations will be accepted. This problem can become fundamental in high dimensional spaces (e.g. for infinite dimensional measures associated to the solutions of stochastic partial differential equations) because, unless care is taken to ensure otherwise, the push-forward measure and the target will not overlap. We discuss this issue and the precise measure-theoretic formulation of the discussion above in Appendix A and the implications for the applications in Secs. IV and V. In contrast, when the map is good enough to ensure a satisfactory acceptance rate, the flow based proposals may mix much faster than proposals based on local moves as independent configurations y can be directly sampled from $\hat{\rho}$.

Improving the map T requires that we optimize some objective function measuring the discrepancy between the $\hat{\rho}(x)$ and $\rho_*(x)$: for example the Kullback-Leibler divergence of ρ_* with respect to $\hat{\rho}$, which is given by an expectation over ρ_*

$$D_{\text{KL}}(\rho_* \parallel \hat{\rho}) = C_* - \int_{\Omega} \log \hat{\rho}(x) \rho_*(x) dx, \quad (9)$$

where $C_* = \int_{\Omega} \log \rho_*(x) \rho_*(x) dx$ is a constant irrelevant for the optimization of T . Typically, this procedure is used in situations where a data set from ρ_* is available beforehand [7, 23] and can be used to construct an empirical approximation of [9]; in contrast, we are focused on situations where only limited data exists initially [6, 18].

III. CONCURRENT SAMPLING AND TRAINING

We concurrently acquire new data with a local MCMC algorithm augmented with a normalizing flow and use these data to optimize the normalizing flow. This procedure is summarized in Algorithm 1 with MALA as the local MCMC algorithm. The sampling component of the algorithm amounts to making a few steps using a local MCMC kernel π followed by one using the normalizing flow kernel π_T . As such it is consistent

with the compounded transition kernel (assuming for simplicity of notation that we make one step which each kernel)

$$\hat{\pi}(x, y) = \int_{\Omega} \pi(x, z) \pi_T(z, y) dz \quad (10)$$

which satisfies the detailed balance relation [2] because the transitions kernels π and π_T individually do. The local kernel π robustly explores the fine-scale structure of modes of the target, while the flow based kernel π_T allows global mixing between modes (once T is sufficiently optimized). The convergence rate of a chain using $\hat{\pi}(x, y)$ is necessarily faster than that of MCMC using $\pi(x, z)$ or $\pi_T(z, y)$ individually: if we assume existence of a spectral gap for both π and π_T and denote the leading eigenvalues of these kernels by $\hat{\lambda} < 1$, $\lambda < 1$, and $\lambda_T < 1$, respectively, we have $\hat{\lambda} \leq \lambda \lambda_T$. Still, if the initial value of the map T is non-optimal, the acceleration will not be substantial, meaning that we must optimize T which becomes possible as we acquire samples. While we employ MALA here, any detailed balance MCMC method could be used in [10]. The transition kernel π does not need to be local, it should, however, have satisfactory acceptance rates early in the training procedure. Note that in the experiments that follow we used unadjusted Langevin dynamics (ULA) because the time steps were sufficiently small to ensure a high acceptance rate.

The training component of Algorithm 1 uses the newly sampled configurations as data with which we can optimize the parameters of the normalizing flow T . Denoting by ρ_k the probability density of the chain with kernel $\hat{\pi}$ after $k \in \mathbb{N}$ steps from initialization ρ_0 , we minimize the KL-divergence of ρ_k with respect to $\hat{\rho}$, $D_{\text{KL}}(\rho_k \parallel \hat{\rho})$, instead of the unknown ρ_* with respect to $\hat{\rho}$ as in [9]. Denoting by $\{x_i(k)\}_{i=1}^n$ the sample of n chains after $k \in \mathbb{N}$ steps of MCMC, this amounts to using the following consistent estimator for $D_{\text{KL}}(\rho_k \parallel \hat{\rho})$, up to an irrelevant constant:

$$\begin{aligned} \mathcal{L}_n[T] &= -\frac{1}{n} \sum_{i=1}^n \log \hat{\rho}(x_i(k)) \\ &= \frac{1}{n} \sum_{i=1}^n (U_{\text{B}}(\bar{T}(x_i(k)) - \log \det |\nabla \bar{T}(x_i(k))|). \end{aligned} \quad (11)$$

In practice, we use stochastic gradient descent on this loss function to update the parameters of the normalizing flow (Algorithm 1 line 11). Details of the architecture and training procedure are described in Appendix E. While the expression for the loss is written at iteration k , we can average gradients over multiple MCMC steps. One key feature of Algorithm 1 is that it performs the two operations of sampling of new configurations via MCMC using $\hat{\pi}(x, y)$, and learning T by minimization of $\mathcal{L}_n[T]$ concurrently.

To initialize MCMC chains, we assume that we have initial data lying in each of the metastable basins of interest, but require no additional information. That is, we take

$$\rho_0(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - x_i(0)), \quad (12)$$

where the $x_i(0)$ are initially in the different basins of attraction of the target but they are not necessarily drawn from ρ_* . We emphasize that the method applies in situations where the location of metastable states of interest are known a priori and one should not expect the procedure to find states far away from initialization.

In Appendix D, we show that it is unlikely that a normalizing flow will discover new metastable basins without any initial information about their location, a fact which is transparently illustrated by experiments on a Gaussian mixture model in Fig. A1. We initialize the map T as the identity transformation and propagate the initial data using $\hat{\pi}$. The initial sampling is essentially driven by the local MCMC, here Langevin dynamics, as the map is not adapted to the target. As the map improves, nonlocal moves start to be accepted, the autocorrelation time drops (Fig. 2) and the Markov chains reallocate mass in proportion to the statistical weights of the different basins, Fig. A1 and Fig. 1.

“Self-training” has been proposed as an alternative to the approach described here. This alternative scheme uses only samples from the model push-forward to approximate and minimize the reverse Kullback-Leibler divergence $D_{\text{KL}}(\hat{\rho} \parallel \rho_*)$. This strategy sees the normalizing flow (or autoregressive model) as a powerful ansatz for variational inference [21, 32]. While successful and highly effective for unimodal distributions [17],

self-training easily leads to “mode collapse” where some of the modes are missed by the model [33], though annealing of the target can help in some cases.

Two important questions arise regarding Algorithm 1: first, does this scheme produce samples that converge in distribution towards the target and, if so, does the adaptive training of the map T improve the rate of convergence to the target distribution? To analyze the properties of a transition operator that combines nonlocal moves with the normalizing flow and a local MCMC algorithm, we consider our approach in the continuous-time limit and show that in this limit the density of the evolving ρ_t with respect to the target ρ_* , defined as $g_t = \rho_t/\rho_*$, satisfies

$$\begin{aligned} \partial_t g_t &= -\nabla U_* \cdot \nabla g_t + \Delta g_t \\ &+ \alpha \int_{\Omega} \min(\hat{g}_t(x), \hat{g}_t(y)) (g_t(y) - g_t(x)) \rho_*(y) dy \end{aligned} \quad (13)$$

where $\hat{g}_t = \hat{\rho}_t/\rho_*$ and $\alpha \geq 0$ is an adjustable parameter that measure the balance between the Langevin and the resampling parts of the dynamics. Setting $\alpha = 0$ amounts to using the Langevin dynamics alone: in that case, for any initial condition ρ_0 , we have that $\rho_t \rightarrow \rho_*$ (i.e. $g_t \rightarrow 1$) as $t \rightarrow \infty$, but this convergence will be exponentially slow in general [34]. The situation changes if we include the resampling step, i.e. consider [13] with $\alpha > 0$. In Appendix B, under various assumptions about \hat{g}_t we derive convergence rates under the dynamics in [13] for the Pearson χ^2 -divergence of ρ_t with respect to ρ_* , which we denote as

$$D_t = \int_{\Omega} \frac{\rho_t^2}{\rho_*} dx - 1 = \int_{\Omega} g_t^2 \rho_* dx - 1 \geq 0. \quad (14)$$

In particular, we study the situation where T learns the instantaneous distribution at all times, that is, $\hat{\rho} = \rho_t$ (and hence $\hat{g}_t = g_t$) for all $t \geq 0$. While this is certainly a significant approximation, we observe in numerical experiments that there is a dramatic improvement in sampling once there is some mixing between metastable basins, which motivates this limiting scenario. In this case, under the assumptions that there exists some $t_0 \geq 0$ such that $D_{t_0} < \infty$ and

$$G_{t_0} = \inf_{x \in \Omega} \frac{\rho_{t_0}(x)}{\rho_*(x)} = \inf_{x \in \Omega} g_{t_0}(x) > 0, \quad (15)$$

we show that

$$\forall t \geq t_0 \quad : \quad D_t \leq \frac{D_{t_0}}{(G_{t_0}(e^{\alpha(t-t_0)} - 1) + 1)^2}. \quad (16)$$

This equation indicates that $D_t \leq D_{t_0}$ remains approximately constant for $\alpha(t - t_0) \leq \log G_{t_0}^{-1}$, then decays exponentially with constant rate $2\alpha > 0$ subsequently. The derivation of [16] also shows that the exponential rate is controlled by the resampling step of the MCMC algorithm that relies on the normalizing flow, and this rate can only improve when we concurrently use Langevin dynamics steps. In Appendix C, we connect the sampling scheme we use to a birth-death Fokker-Planck equation [35], which could also be implemented in practice as a Markov jump process; again this analysis emphasizes the favorable convergence properties of the scheme.

IV. FAST-MIXING AUGMENTED MCMC FOR RANDOM FIELDS

To illustrate the efficacy of augmented sampling, we first consider a stochastic Allen-Cahn model, a canonical and ubiquitous model for the microscopic physics of phase transitions in condensed matter systems [36]. The stochastic Allen-Cahn equation is defined in terms of a random field $\phi : [0, 1] \rightarrow \mathbb{R}$ that satisfies

$$\partial_t \phi = a \partial_s^2 \phi + a^{-1} (\phi - \phi)^3 + \sqrt{2\beta^{-1}} \eta(t, s) \quad (17)$$

where $a > 0$ is a parameter, β is the inverse temperature, $s \in [0, 1]$ denotes the spatial variable, η is a spatio-temporal white noise and we impose Dirichlet boundary conditions in which $\phi(s = 0) = \phi(s = 1) = 0$

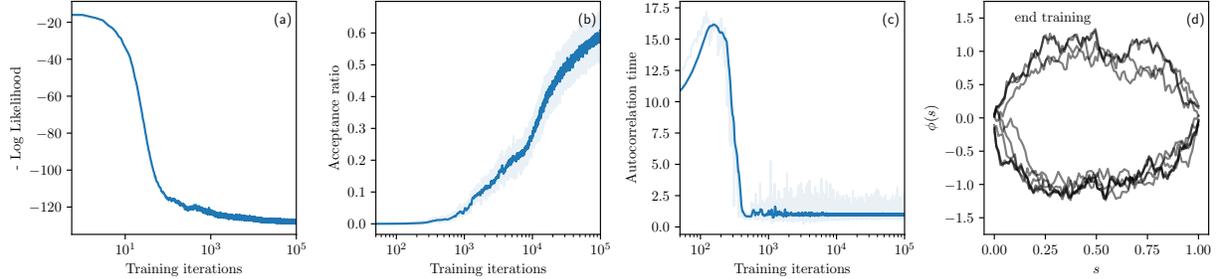


FIG. 2. **Concurrent training and sampling of the stochastic Allen-Cahn model with a Real-NVP normalizing flow.** (a) The stochastic gradient descent using samples generated by the procedure decreases the negative log-likelihood gradually. (b) As the training progresses the acceptance rate in the Metropolis-Hasting using proposals from the normalizing flow improves gradually, reaching levels well beyond 50%. Rolling average over the last 50 time steps is plotted in darker color. (c) As independent proposals from the flow starts getting accepted the Markov Chain autocorrelation times drops abruptly. (d) Fast mixing is illustrated by looking at the consecutive states of one walker updated with the transition kernel combining local Langevin updates and resampling with the push-forward. In 10 steps, the single walker has jumped between ϕ_+ and ϕ_- .

throughout. This stochastic partial differential equation (SPDE) is well-posed in one spatial dimension [37, 38], and its invariant measure is the Gibbs measure associated with the Hamiltonian

$$U_*[\phi] = \beta \int_0^1 \left[\frac{a}{2} (\partial_s \phi)^2 + \frac{1}{4a} (1 - \phi^2(s))^2 \right] ds. \quad (18)$$

The first term in the Hamiltonian [18] is a spatial coupling that penalizes changes in ϕ and hence, at low temperature, has the effect of aligning the field in positive or negative direction. As a result the Hamiltonian [18] has two global minima, denoted by ϕ^+ and ϕ^- , in which the typical values of ϕ are ± 1 (see Fig. 1 (a)). Because there is a free energy barrier between ϕ^+ and ϕ^- , local updates via traditional MCMC based e.g. on using the stochastic Allen-Cahn equation [17] will not mix, even on very long timescales. Indeed, if we wanted to compute the free energy difference between these basins, we would need to construct a pathway through configuration space and use importance sampling techniques along the path [39]. Our normalizing flow augmented MCMC algorithm offers an alternative approach. Fig. 2 demonstrates that a map T can be trained to efficiently generate samples with high statistical weight in the target distribution enabling rapid mixing across the free energy barrier.

In order to learn the map robustly, a standard implementation of a normalizing flow model, with a standard Gaussian field with uncoupled spins as base measure, does not suffice in this instance. As shown in Fig. A3, samples generated by the standard scheme are not credible enough to be accepted in the Metropolis-Hasting step of the concurrent sampling and training (line 9 of Alg. 1), even after a subsequent amount of training.

Using a base measure that is “informed” alleviates this issue. Explicitly, we sample the base measure corresponding to a Gaussian random field with a local coupling (a “Ornstein-Uhlenbeck bridge”), which corresponds to a system with Hamiltonian

$$U_B[\phi] = \beta \int_0^1 \left[\frac{a}{2} (\partial_s \phi)^2 + \frac{1}{2a} \phi^2 \right] ds. \quad (19)$$

Importantly, this measure does not have any metastability and remains easy to sample. As discussed in Appendix A, we must choose this measure to ensure that the push-forward distribution has a non-vanishing statistical weight in the target distribution; this can be achieved by adapting the base measure to the statistical properties of the target measure or by adapting the samples from the target to the base measure using a “whitening” transformation [15].

In practice, we must discretize the field on a grid, and throughout we take $N = 100$ with a lattice spacing $\Delta s = 1/N$ meaning that the map we must learn is high-dimensional $T : \mathbb{R}^N \rightarrow \mathbb{R}^N$. We also use the associated Langevin equation as discretized version of the SPDE [17] to generate samples in the local component of our compounded MCMC scheme. Typical configurations $\phi(x)$, in this case generated by the normalizing flow, are shown in Fig. 1 (a).

We trained maps T and T_U using our algorithm with the informed base measure [19] and an uninformed Gaussian measure that lacked coupling term (Eq. E3 in the Appendix), respectively. While T generates samples which are accepted in the MCMC procedure with average acceptance rate approaching 60% (Fig. 2), T_U fails to produce samples that have appreciable statistical weight in the target distribution. The evident difference is in the local structure of the random fields that are produced. Fig. 1 (c) shows the Fourier spectrum of field ϕ computed with samples from the target measure (obtained using the proposed MCMC method after convergence) as well as from the push-forwards in the informed $\hat{\rho}$ and uninformed $\hat{\rho}_U$ case. While $\hat{\rho}$ accurately captures the decay of the Fourier components at all scales, T_U fails to compensate for the uncoupled base measure and $\hat{\rho}_U$ does not accurately capture high-frequency oscillations of the field ϕ , which subsequently leads to high rejection rates in the MCMC procedure.

Examining the learned map T reveals its simple underlying structure. As shown in Fig. 1 (b),

the map is spatially local, transporting spins near the center of the domain to ± 1 while spins near the boundary are mapped closer to 0. It is again useful to examine the properties of a mapped configuration in Fourier space; the $k = 0$ mode reveals that the mean value is transported substantially: $T(\hat{\phi}_0)$ is approximately ± 1 , as shown in Fig. A2 (c). However, higher frequency modes are left invariant by the map, see Fig. A2 (d,e).

Perhaps most remarkably, the learned map T can be used to evaluate free energy differences between the metastable basins ϕ^- and ϕ^+ , even in thermodynamic conditions distinct from those in which the map was trained. Fig. 1 (d) shows an estimate of the free energy difference between the positive and negative metastable basins as a function of an external field h , which enters the Hamiltonian as

$$U_{*,h}[\phi] = \beta \int_0^1 \left[\frac{a}{2} (\partial_s \phi)^2 + \frac{1}{4a} (1 - \phi^2(s))^2 + h\phi(s) \right] ds. \quad (20)$$

These estimates are produced with importance sampling using $\hat{\rho}$ as described in Appendix F. Analytical estimates at low temperature via a Laplace approximation reveal that the normalizing flow accurately recapitulates the free energy difference despite the fact that the map was optimized only with samples where $h = 0$. Similar generalization properties were observed in Refs. [18, 32], where a map was used at temperatures distinct from the temperature at which training data was collected. This approach is valid in cases where the modified parameter, here the field h , distorts the relative populations of the metastable basins, but has a mild effect on the local structure of the field, which can be controlled by monitoring the variance of the estimator. Finally, we also show that the method can be useful to sample configurations with domain walls by tilting the Hamiltonian (see Appendix E and Fig. A4).

V. NONEQUILIBRIUM TRANSITION PATH SAMPLING

The success of the algorithm on the Allen-Cahn SPDE model suggests that it could be a useful strategy for sampling transition paths on complex potential energy surfaces even in the presence of nonconservative driving forces. Transition path theory [40, 41] and path sampling techniques [4, 42] allow for space-time local sampling, but it is generically difficult to compute relative weights between transition paths. Nevertheless, these problems are important for determining reaction mechanisms and reaction rates for chemical reactions, a topic which increasingly is being studied in nonequilibrium settings [43, 44]. We consider the continuous-time dynamics

$$dX_t = b(X_t)dt + \sqrt{2\beta^{-1}}dW_t \quad (21)$$

where b is a drift. When b is non-conservative, the dynamics is not in detailed balance and the stationary distribution of the dynamics is non-Boltzmann. Few robust methods are currently available to compare relative path weights when there is metastability in space-time [45].

Following the perspective from potential theory [46], we investigated the ability of our augmented Monte Carlo dynamics to sample multiple transition paths.

Denoting by $x_{[0,t_m]}$ a trajectory on $t \in [0, t_m]$, we seek to sample such trajectories conditioning on the end points $x_0 = x_A$ and $x_{t_m} = x_B$. Using a path integral type expression for the probability density in path space, we can write

$$\mathbb{P}_*(x_{[0,t_m]}) \propto \exp \left[-\frac{\beta}{4} \int_0^{t_m} |\dot{x}_t - b(x_t)|^2 dt \right]. \quad (22)$$

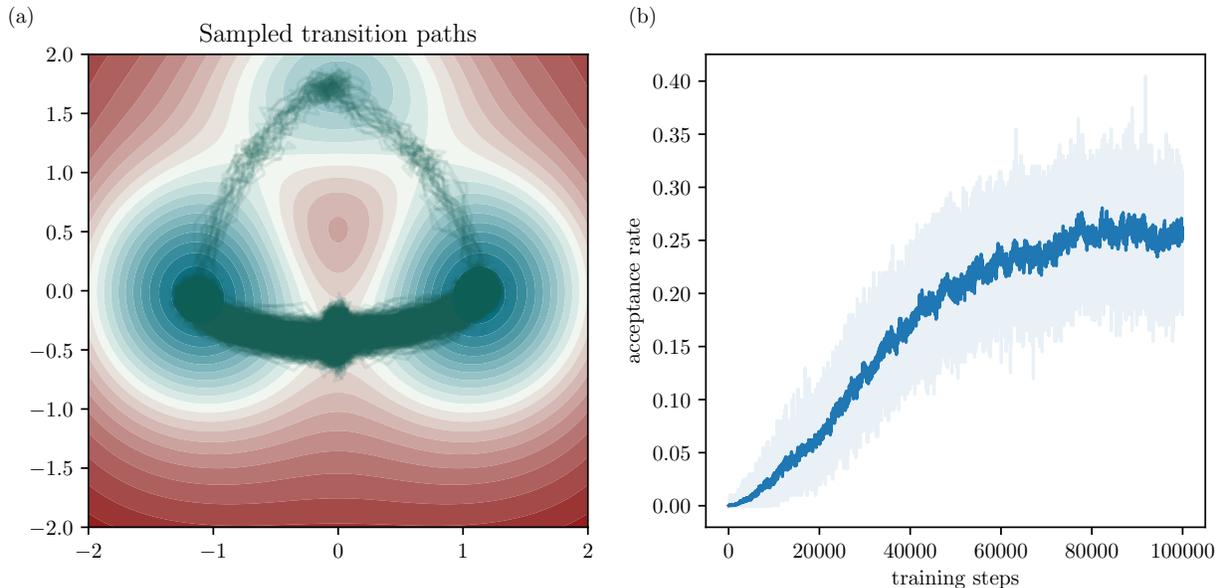


FIG. 3. **Sampling nonequilibrium transition paths over an entropic path-space barrier with augmented path-space MCMC.** (a) 1000 paths (solutions to [21] with end-point conditioning) sampled with $\hat{\pi}$. The potential energy surface is shown as contours. (b) The acceptance rate of samples generated with the flow map as a function of optimization time is shown with the rolling average over the last 50 time steps plotted above.

We emphasize here that while the above expression is only formal, the relative path measures that we work with can be made rigorous using the Girsanov theorem, cf. Appendix A.

As in the case of the stochastic Allen-Cahn model, we must take care to ensure that the base measure is well-adapted to the target. In the infinite-dimensional Hilbert space on which the path measure is defined, the naive base measure typically used for a normalizing flow will be orthogonal to [22]. As a consequence, the relative entropy used to train the map T [11] can be infinite [47].

The probability measure associated with the Brownian bridge process offers a simple solution to this problem. Again, noting that the expression is only formal, we write the path integral expression for the statistical weight of a Brownian bridge path $x_{[0,T]}$:

$$\mathbb{P}_B(x_{[0,t_m]}) \propto \exp \left[-\frac{\beta}{4} \int_0^{t_m} |\dot{x}_t|^2 dt \right], \quad (23)$$

with the boundary conditions $x_0 = x_A$ and $x_{t_m} = x_B$. The choice fulfills the key requirement that the nonequilibrium path measure is absolutely continuous with respect to the base measure (cf. Appendix A).

To carry out the simulations, we must discretize the path into N points with a time step Δt . We then obtain the target path action

$$S_* = \beta \Delta t \sum_{i=1}^{N-1} \left| \frac{x_{i+1} - x_i}{\Delta t} - b(x_i) \right|^2 \quad (24)$$

where the factor of Δt in the effective temperature ensures proper scaling in the limit $\Delta t \rightarrow 0$; the action S_* plays the role of U_* in the example above. We impose the boundary conditions $x_0 = x_A$ and $x_N = x_B$. An analogous expression for the path action associated with the base measure simply lacks the contribution from the drift, $b(x_i)$. We update the paths locally in space-time by carrying out Langevin dynamics on the path action, which can be written pointwise for $i = 1, \dots, N-1$ as

$$x_i(k+1) = x_i(k) - \frac{\partial S_*(x(k))}{\partial x_i} \tau + \sqrt{2\tau} \eta_i(k), \quad (25)$$

where $\tau > 0$ is the time step of the path-space Langevin dynamics. Just as in the examples considered above, if there is metastability in the space of paths, the space-time local Langevin dynamics is metastable and inter-conversion between transition paths will be prohibitively slow for a purely local sampling algorithm.

Fig. 3 shows that the augmented MCMC algorithm can be used to efficiently and directly sample transition paths from separated metastable basins. This potential energy surface has been used as a model of entropic switching, originally in Ref. [48] (cf. Appendix E). The drift term b is the sum of a conservative force $-\nabla V$ and a nonconservative drift given by the vector field $c(-x_2, x_1)^T$, leading to a nonequilibrium dynamics. While techniques based on transition path theory typically require an equilibrium dynamics, the present method applies equally well to problems in which there is a nonequilibrium drift. The dimensionality of the underlying potential is low in this example, nevertheless, the normalizing flow accepts a $2N$ dimensional input, where N is the number of discretization points, leading to an effective dimension of 200 in this example. The results emphasize that metastability in path space can be probed directly even when there is an appreciable gap between the two basins and a substantially different local structure to the typical transition paths.

The paths shown in Fig. 3 were obtained by training the normalizing flow with Langevin alone, the limit of slow resampling during the training. This was done to mitigate a potential failure mechanism of the algorithm in which a higher probability mode A is learned by the normalizing flow before a lower probability mode B (for a detailed discussion cf. Appendix E and Fig. A5). With a finite batch size, if a A is learned initially, the Langevin walkers from B can be transported to A , potentially eliminating all data from B . This failure mechanism is both easy to monitor and easy to correct, simply by ensuring that the resampling does not eliminate a basin of interest.

As in the case of the stochastic Allen-Cahn model, the map learns low frequency information and preserves high-frequency. We illustrate this fact by examining the behavior of the map on a truncated Karhunen-Loève expansion of the Brownian bridge, shown in Fig. A6. The low order terms of this expansion capture the large scale structure of the bridge process whereas higher order terms include higher frequency oscillations. The map recapitulates the overall structure of the paths for an expansion of the bridge process truncated at $k_{\max} = 2$, but the paths themselves are comparatively smooth compared to solutions of the conditioned SDE [21]. Including higher order terms does not change the overall shape of the paths, furthermore, doing so restores the higher frequency oscillations in the transition paths.

VI. CONCLUSIONS

As the use of data-driven methods from machine learning become increasingly routine in the physical sciences, we must carefully assess the cost of data acquisition and training to ensure that we can leverage ML methods in a productive fashion. Sampling systems with complex local structure and multiple metastable basins is a generically challenging task in high-dimensional systems, and we have already seen that neural networks can contend with this challenge in nontrivial settings [16–18, 21, 23]. Nonlocal transport in MCMC algorithms can significantly enhance mixing and normalizing flows provide a compelling framework for designing adaptive schemes, even in cases where no statistically representative data set is available at first. Nevertheless, we do not find that these methods enable discovery of unknown modes of a target distribution, emphasizing the importance of having some *a priori* information about the metastable states of the system. Many questions remain about how to best adapt the base measure to different targets to ensure efficient learning and encourage desirable properties of the map, such as locality and transferability.

Acknowledgements: GMR acknowledges generous support from the Terman Faculty Fellowship. EVE acknowledges partial support from the National Science Foundation (NSF) Materials Research Science and Engineering Center Program grant DMR-1420073, NSF DMS-1522767, and a Vannevar Bush Faculty Fellowship.

[1] A. R. Barron, Universal approximation bounds for superpositions of a sigmoidal function, *IEEE Transactions on Information Theory* **39**, 930 (1993).

- [2] F. Bach, Breaking the Curse of Dimensionality with Convex Neural Networks, *Journal of Machine Learning Research* **18**, 1 (2017).
- [3] G. Carleo, I. Cirac, K. Cranmer, and L. Daudet, Machine learning and the physical sciences*, *Rev. Mod. Phys.* **91**, 39 (2019).
- [4] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, Transition Path Sampling: Throwing Ropes Over Rough Mountain Passes, in the Dark, *Annual Review of Physical Chemistry* **53**, 291 (2002).
- [5] D. Sejdinovic, H. Strathmann, M. L. Garcia, C. Andrieu, and A. Gretton, Kernel Adaptive Metropolis-Hastings, in *International Conference on Machine Learning* (PMLR, 2014) pp. 1665–1673.
- [6] D. Levy, M. D. Hoffman, and J. Sohl-Dickstein, Generalizing Hamiltonian Monte Carlo with Neural Networks, in *International Conference on Learning Representations* (2018).
- [7] J. Song, S. Zhao, and S. Ermon, A-nice-mc: Adversarial training for MCMC, in *Advances in Neural Information Processing Systems*, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017).
- [8] M. K. Titsias, Learning Model Reparametrizations: Implicit Variational Inference by Fitting MCMC distributions, arXiv:1708.01529 [stat] (2017), arXiv:1708.01529 [stat].
- [9] T. A. Le, M. Igl, T. Rainforth, T. Jin, and F. Wood, Auto-encoding sequential monte carlo, in *International Conference on Learning Representations* (2018).
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative Adversarial Nets, in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014) pp. 2672–2680.
- [11] D. P. Kingma and M. Welling, Auto-Encoding Variational Bayes, arXiv (2013), arXiv:1312.6114v10.
- [12] R. Salakhutdinov, A. Mnih, and G. Hinton, Restricted Boltzmann machines for collaborative filtering, in *ICML '07* (ACM, New York, New York, USA, 2007) pp. 791–798.
- [13] M. Germain, K. Gregor, I. Murray, and H. Larochelle, MADE: Masked autoencoder for distribution estimation, in *32nd International Conference on Machine Learning, ICML 2015*, Vol. 2 (2015) pp. 881–889.
- [14] D. Rezende and S. Mohamed, Variational Inference with Normalizing Flows, in *International Conference on Machine Learning* (PMLR, 2015) pp. 1530–1538.
- [15] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *Journal of Machine Learning Research* **22**, 1 (2021).
- [16] L. Huang and L. Wang, Accelerated Monte Carlo simulations with restricted Boltzmann machines, *Physical Review B* **95**, 035105 (2017).
- [17] M. S. Albergo, G. Kanwar, and P. E. Shanahan, Flow-based generative models for Markov chain Monte Carlo in lattice field theory, *Physical Review D* **100**, 034515 (2019).
- [18] F. Noé, S. Olsson, J. Köhler, and H. Wu, Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning, *Science* **365**, eaaw1147 (2019).
- [19] H. Sidky, W. Chen, and A. L. Ferguson, Molecular latent space simulators, *Chemical Science* **11**, 9459 (2020).
- [20] L. Sbaïlò, M. Dibak, and F. Noé, Neural mode jump Monte Carlo, *Journal of Chemical Physics* **154**, 074101 (2021).
- [21] D. Wu, L. Wang, and P. Zhang, Solving Statistical Mechanics Using Variational Autoregressive Networks, *Physical Review Letters* **122**, 080602 (2019).
- [22] K. A. Nicoli, C. J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel, S. Nakajima, and P. Stornati, Estimation of Thermodynamic Observables in Lattice Field Theories with Deep Generative Models, *Physical Review Letters* **126**, 032001 (2021).
- [23] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, Score-based generative modeling through stochastic differential equations, in *International Conference on Learning Representations* (2021).
- [24] S. P. Meyn and R. L. Tweedie, *Markov chains and stochastic stability* (Springer Science & Business Media, 2012).
- [25] G. O. Roberts and R. L. Tweedie, Exponential convergence of Langevin distributions and their discrete approximations, *Bernoulli* **2**, 341 (1996).
- [26] C. Andrieu and É. Moulines, On the ergodicity properties of some adaptive MCMC algorithms, *Annals of Applied Probability* **16**, 1462 (2006).
- [27] H. Haario, E. Saksman, and J. Tamminen, An adaptive Metropolis algorithm, *Bernoulli* **7**, 223 (2001).
- [28] C. Andrieu, A. Jasra, A. Doucet, and P. D. Moral, On nonlinear Markov chain Monte Carlo, *Bernoulli* **17**, 987 (2011).
- [29] A. Jasra, D. A. Stephens, and C. C. Holmes, On population-based simulation for static inference, *Statistics and Computing* **17**, 263 (2007).
- [30] L. Dinh, D. Krueger, and Y. Bengio, NICE: Non-linear independent components estimation, 3rd International Conference on Learning Representations, ICLR 2015 - Workshop Track Proceedings **1**, 1 (2015), 1410.8516.
- [31] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density Estimation Using Real NVP, in *International Conference on Learning Representations* (2017) p. 32.
- [32] K. A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K. R. Müller, and P. Kessel, Asymptotically unbiased estimation of physical observables with neural samplers, *Physical Review E* **101**, 10.1103/PhysRevE.101.023304 (2020).

- (2020).
- [33] G. S. Hartnett and M. Mohseni, Self-Supervised Learning of Generative Spin-Glasses with Normalizing Flows, [arXiv:2001.00585 \[cond-mat, physics:quant-ph, stat\]](#) (2020), [arXiv:2001.00585 \[cond-mat, physics:quant-ph, stat\]](#).
 - [34] D. W. Stroock, Logarithmic Sobolev inequalities for gibbs states, in *Dirichlet Forms: Lectures given at the 1st Session of the Centro Internazionale Matematico Estivo (C.I.M.E.) Held in Varenna, Italy, June 8–19, 1992*, Lecture Notes in Mathematics, edited by E. Fabes, M. Fukushima, L. Gross, C. Kenig, M. Röckner, D. W. Stroock, G. Dell’Antonio, and U. Mosco (Springer, Berlin, Heidelberg, 1993) pp. 194–228.
 - [35] Y. Lu, J. Lu, and J. Nolen, Accelerating Langevin Sampling with Birth-death, [arXiv:1905.09863 \[cs, math, stat\]](#) (2019), [arXiv:1905.09863 \[cs, math, stat\]](#).
 - [36] N. Berglund, G. D. Gesù, and H. Weber, An Eyring–Kramers law for the stochastic Allen–Cahn equation in dimension two, [Electronic Journal of Probability](#) **22**, 1 (2017).
 - [37] W. G. Faris and G. Jona-Lasinio, Large fluctuations for a nonlinear heat equation with noise, [Journal of Physics A: Mathematical and General](#) **15**, 3025 (1982).
 - [38] R. Marcus, Parabolic Ito Equations, [Transactions of the American Mathematical Society](#) **198**, 177 (1974).
 - [39] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier, 2001).
 - [40] W. E, W. Ren, and E. Vanden-Eijnden, Transition pathways in complex systems: Reaction coordinates, isocommittor surfaces, and transition tubes, [Chem. Phys. Lett.](#) **413**, 242 (2005).
 - [41] W. E and E. Vanden-Eijnden, Transition-Path Theory and Path-Finding Algorithms for the Study of Rare Events, [Annual Review of Physical Chemistry](#) **61**, 391 (2010).
 - [42] R. J. Allen, C. Valeriani, and P. R. ten Wolde, Forward flux sampling for rare event simulations, [Journal of Physics: Condensed Matter](#) **21**, 463102 (2009).
 - [43] G. Falasco and M. Esposito, Dissipation-Time Uncertainty Relation, [Physical Review Letters](#) **125**, 120604 (2020).
 - [44] B. Kuznets-Speck and D. T. Limmer, Dissipation bounds the amplification of transition rates far from equilibrium, [Proceedings of the National Academy of Sciences](#) **118**, e2020863118 (2021).
 - [45] P. Metzner, C. Schütte, and E. Vanden-Eijnden, Illustration of transition path theory on a collection of simple examples, [The Journal of Chemical Physics](#) **125**, 084110 (2006).
 - [46] A. Bovier, M. Eckhoff, V. Gaynard, and M. Klein, Metastability and Low Lying Spectra in Reversible Markov Chains, [Communications in Mathematical Physics](#) **228**, 219 (2002).
 - [47] M. Hairer, *An Introduction to Stochastic PDEs* (2009) p. 78.
 - [48] S. Park, M. K. Sener, D. Lu, and K. Schulten, Reaction paths based on mean first-passage times, [The Journal of Chemical Physics](#) **119**, 1313 (2003).
 - [49] H. Wu, J. Köhler, and F. Noé, Stochastic normalizing flows, in *Advances in Neural Information Processing Systems*, Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., 2020) pp. 5933–5944.

Appendix A: Measure Theoretic Formulation of the Normalizing Flow augmented sampling algorithm

In the continuous limit, the models discussed in the main text do not have a well-defined probability density function. For the sake of precision and completeness, we reformulate the discussion Sec. II in terms of arbitrary probability measures on a set Ω with σ -algebra \mathcal{F} . Let $(\Omega, \mathcal{F}, \nu_*)$ be the probability space associated with the target probability measure ν_* . Let us denote the transition kernel by $P(x, dy)$ (so that $P(x, dy) = \pi(x, y)dy$ if P has a density) and write the detailed balance relation as (compare [2])

$$\nu_*(dx)P(x, dy) = \nu_*(dy)P(y, dx), \quad (\text{A1})$$

meaning that for all suitable test functions $\chi : \Omega \times \Omega \rightarrow \mathbb{R}$ we have (compare [3])

$$\int_{\Omega \times \Omega} \chi(x, y) (\nu_*(dx)P(x, dy) - \nu_*(dy)P(y, dx)) = 0. \quad (\text{A2})$$

In this general setup, the nonlocal map that we use to augment the local dynamics is an invertible transformation from Ω to itself, i.e $T : \Omega \rightarrow \Omega$. The map T is used to transport a base measure ν_B from which we can sample efficiently into the target measure ν_* . In the ideal case, the map T_* is constructed so that

$$\nu_* = T_{*\#}\nu_B \quad (\text{A3})$$

where the equality above holds in the weak sense, i.e. for a test function $\phi : \Omega \rightarrow \mathbb{R}$,

$$\int_{\Omega} \phi(x)\nu_*(dx) = \int_{\Omega} \phi(T_*(x))\nu_B(dx). \quad (\text{A4})$$

In practice, instead of T_* we only have at our disposal some approximation $T : \Omega \rightarrow \Omega$, also invertible. In order to use this imperfect map in a Metropolis-Hastings procedure, we must require that the target measure, ν_* , and the push-forward of the base measure measure, $T_{\#}\nu_B$, be mutually absolutely continuous with respect to one another. If that is the case, denoting by $(d\nu_*/dT_{\#}\nu_B)(x)$ and $(dT_{\#}\nu_B/d\nu_*)(x) = 1/(d\nu_*/dT_{\#}\nu_B)(x)$ the Radon-Nikodym derivatives of $T_{\#}\nu_B$ with respect to ν_* and of ν_* with respect to $T_{\#}\nu_B$, the MCMC algorithm that resamples from $T_{\#}\nu_B$ proceeds as follows

1. Given $x(k)$, sample $x_B \sim \nu_B$ and let $y = T(x_B)$
2. Set $x(k+1) = y$ with probability

$$\text{acc}(x(k), y) = \min \left[1, \frac{d\nu_*(y)}{dT_{\#}\nu_B(y)} \frac{dT_{\#}\nu_B}{d\nu_*(x(k))} \right] \quad (\text{A5})$$

and otherwise set $x(k+1) = x(k)$.

In our approach, the map T is evolved by minimizing some objective function, for example the KL divergence, measuring the discrepancy between the pushforward $T_{\#}\nu_B$ and the target ν_* or its approximation after k step of MCMC sampling, ν_k . In this case it is crucial to choose a base measure ν_B such that the KL divergence is well-defined initially: indeed if it is finite at the start of the procedure, it will remain so during optimization since the KL divergence can only decrease during gradient descent. Such a choice is not automatic for the SPDE examples we consider in text, because the infinite-dimensional measures associated with the solution of these equations have a strong propensity to be mutually orthogonal [47].

1. Stochastic AC equation

In the case of the stochastic Allen-Cahn equation, as base measure it is convenient to use the measure associated with the Hamiltonian [19]. This is the Gaussian measure of an Ornstein-Uhlenbeck bridge, i.e. the

process with mean zero and covariance (this is the inverse of the operator $-\beta a \partial_s^2 + \beta a^{-1}$ with zero Dirichlet boundary conditions):

$$C_B(s, s') = \frac{a}{2\beta} \frac{e^{-|s-s'|/a+1} + e^{-|s-s'|/a-1} - e^{-(s+s')/a+1} - e^{(s+s')/a-1}}{(e^1 - e^{-1})}. \quad (\text{A6})$$

Denoting a realization of this bridge by $B \in C([0, 1], \mathbb{R})$ we then have

$$\frac{d\nu_*}{d\nu_B}(B) = Z_V^{-1} \exp\left(-\beta \int_0^1 V(B(s)) ds\right), \quad V(B) = \frac{1}{4a}(1 - B^2)^2 \quad (\text{A7})$$

where

$$Z_V = \mathbb{E}_{\nu_B} \exp\left(-\beta \int_0^1 V(B(s)) ds\right) \quad (\text{A8})$$

The Radon-Nikodym derivative [A7] is well-defined if $Z_V < \infty$, which is the case for the potential energy $V(B)$ [37]. The KL divergence of ν_* with respect to ν_B can therefore be expressed as

$$D_{\text{KL}}(\nu_* \| \nu_B) = -\log Z_V - \beta Z_V^{-1} \mathbb{E}_{\nu_B} \left(\bar{V}(B) e^{-\beta \bar{V}(B)} \right) \quad (\text{A9})$$

where we denote $\bar{V}(B) = \int_0^1 V(B(s)) ds$. We can also invert these relations and express them in terms of expectation over the target ν_* . Denoting by $\phi \in C([0, 1], \mathbb{R})$ a random sample of this measure, we have

$$Z_V = \left(\mathbb{E}_{\nu_*} \exp\left(\beta \int_0^1 V(\phi(s)) ds\right) \right)^{-1} \quad (\text{A10})$$

and

$$D_{\text{KL}}(\nu_* \| \nu_B) = -\log Z_V - \beta \mathbb{E}_{\nu_*} \left(\int_0^1 V(\phi(s)) ds \right). \quad (\text{A11})$$

The quantities in these last two formulas can be estimated by replacing the expectation over ν_* by an empirical average over samples drawn from this measure or, more generally, drawn from the current approximation of this measure in the MCMC, ν_k . This allows us to start the training procedure of the map T . In practice, of course, the fields in the formula above are discretized on a grid with N points and the integrals are replaced by Riemann sums. Correspondingly, the map T is approximated by a map from \mathbb{R}^N to \mathbb{R}^N , using the RealNVP described in Appendix E.

2. Transition Paths

A similar procedure can be used to construct the base measure in the context of transition path sampling. In this case, as base measure we can take the Gaussian measure associated with the bridge process connecting x_A to x_B : this process can be expressed as

$$X_B(t) = x_A(1 - t/t_m) + x_B t/t_m + 2\beta^{-1} B(t/t_m) \quad (\text{A12})$$

where $B \in C([0, 1], \mathbb{R}^d)$ is the standard Brownian bridge in \mathbb{R}^d , i.e. the Gaussian process with mean zero and covariance

$$C(u, u') = \begin{cases} u'(1 - u) \text{Id} & \text{if } u' \leq u \\ u(1 - u') \text{Id} & \text{if } u \leq u' \end{cases} \quad (\text{A13})$$

Using this bridge process as base measure leads to Radon-Nikodym derivative

$$\frac{d\nu_*}{d\nu_B}(X_B) = \exp\left(\frac{1}{2\beta} \int_0^{t_m} b(X_B(t)) dX_B(t) - \frac{1}{4\beta} \int_0^{t_m} |b(X_B(t))|^2 dt\right) \quad (\text{A14})$$

which is a Martingale with respect to the filtration \mathcal{F}_{t_m} whenever the drift satisfies the Novikov condition, namely,

$$\mathbb{E} \exp\left(-\frac{1}{4\beta} \int_0^{t_m} |b(X_B(t))|^2 dt\right) < +\infty. \quad (\text{A15})$$

Appendix B: Continuous limit of the MCMC scheme

1. Chapman-Kolmogorov equation

Written in terms of the densities $\rho_*(x)$ and $\hat{\rho}(y)$ (assumed to be fixed for now) the transition kernel in [10] reads

$$\pi_T(x, y) = a(x, y)\hat{\rho}(y) + (1 - b(x))\delta(x - y) \quad (\text{B1})$$

where

$$a(x, y) = \min\left(\frac{\hat{\rho}(x)\rho_*(y)}{\hat{\rho}(y)\rho_*(x)}, 1\right), \quad (\text{B2})$$

$$b(x) = \int_{\Omega} a(x, y)\hat{\rho}(y)dy.$$

Denoting as $\{\rho_k(x)\}_{k \in \mathbb{N}}$ the updated probability density of the walker in the Markov chain associated with the kernel $\pi_T(x, y)$ alone, this density satisfies the Chapman-Kolmogorov equation

$$\rho_{k+1}(x) = \int_{\Omega} \rho_k(y)\pi_T(y, x)dy. \quad (\text{B3})$$

Using the explicit form of $\pi_T(x, y)$ in [B1], after some simple reorganization this equation can be written as

$$\rho_{k+1}(x) = \rho_k(x) + \int_{\Omega} R(x, y) (\rho_*(x)\rho_k(y) - \rho_k(x)\rho_*(y)) dy \quad (\text{B4})$$

where we defined

$$R(x, y) = R(y, x) = \min\left(\frac{\hat{\rho}(x)}{\rho_*(x)}, \frac{\hat{\rho}(y)}{\rho_*(y)}\right). \quad (\text{B5})$$

Note that if we had $\hat{\rho} = \rho_*$, then $R(x, y) = 1$ and [B4] would reach equilibrium in one step, $\rho_{k+1} = \rho_*$ whatever ρ_k .

Algorithm2 MCMC with partial resampling steps, given an evolving map

```

1: SAMPLE( $U_*$ ,  $\{T_t\}_{t \geq 0}$ ,  $\{x_i(0)\}_{i=1}^n$ ,  $\tau$ ,  $k_{\max}$ ,  $\alpha$ )
2: Inputs:  $U_*$  target energy,  $\{T_t\}_{t \geq 0}$  evolving maps,  $\{x_i(0)\}_{i=1}^n$  initial data,  $\tau > 0$  time step,  $T_{\max} > 0$  total duration,  $\alpha > 0$ 
3:  $k = 0$ 
4: while  $k < T_{\max}/\tau$  do
5:   for  $i = 1, \dots, n$  do
6:     if  $k \bmod 2 = 0$  then
7:        $x'_{B,i} \sim \rho_B$ 
8:        $x'_i = T_{k\tau}(x'_{B,i}) \triangleright$  push-forward via  $T$ 
9:        $x_i(k+1) = x'_i$  with probability  $\min(\alpha\tau \text{acc}(x_i(k), x'_i), 1)$ , otherwise  $x_i(k+1) = x_i(k) \triangleright$  partial resampling step
10:    else
11:       $x_i(k+1) = x_i(k) - \tau \nabla U_*(x_i(k)) + \sqrt{2\tau} \eta_i$  with  $\eta_i \sim \mathcal{N}(0_d, I_d) \triangleright$  discretized Langevin step
12:     $k \leftarrow k + 1$ 
13: return:  $\{x_i(k)\}_{k=0, i=1}^{k_{\max}, n}$ 

```

2. Continuous limit

To take the continuous limit of [B4], we modify this equation in a way that the update of the density is only partial. Specifically, denoting ρ_t the value of the density at time $t \geq 0$, we turn this equation into

$$\rho_{t+\tau}(x) = \rho_t(x) + \alpha\tau \int_{\Omega} R(x, y) (\rho_*(x)\rho_t(y) - \rho_t(x)\rho_*(y)) dy \quad (\text{B6})$$

where $\alpha > 0$ and $\tau > 0$ are parameters. This will allow us to make the MCMC resampling updates on par with those of MALA, using $\tau > 0$ as timestep in both (see Algorithm 2). Subtracting $\rho_t(x)$ from both sides of [B6], dividing by τ , and letting $\tau \rightarrow 0$ gives

$$\partial_t \rho_t(x) = \alpha \int_{\Omega} R(x, y) (\rho_*(x) \rho_t(y) - \rho_t(x) \rho_*(y)) dy. \quad (\text{B7})$$

We can now add the Langevin terms that arise in the continuous limit of the compounded MCMC scheme that we use, to arrive at

$$\partial_t \rho_t = \nabla \cdot (\rho_t \nabla U_* + \nabla \rho_t) + \alpha \int_{\Omega} R(x, y) (\rho_*(x) \rho_t(y) - \rho_t(x) \rho_*(y)) dy \quad (\text{B8})$$

where $\alpha > 0$ measures the separation of time scale between the Langevin and the resampling terms. This equation arises in the limit as $\tau \rightarrow 0$ for the evolution specified in Algorithm 2. Written in term of $g_t = \rho_t / \rho_*$ and $\hat{g}_t = \hat{\rho}_t / \rho_*$ (now also allowed to vary with time) Eq. [B8] reads

$$\partial_t g_t = -\nabla U_* \cdot \nabla g_t + \Delta g_t + \alpha \int_{\Omega} \min(\hat{g}_t(x), \hat{g}_t(y)) (g_t(y) - g_t(x)) \rho_*(y) dy \quad (\text{B9})$$

3. Convergence rate

Consider the evolution of the Pearson χ^2 -divergence of ρ_t with respect to ρ_* defined in [14] assuming that $D_0 < \infty$. Using [B9] we deduce

$$\begin{aligned} \frac{dD_t}{dt} &= 2 \int_{\Omega} g_t(x) \partial_t g_t(x) \rho_*(x) dx \\ &= 2 \int_{\Omega} g_t(x) \nabla \cdot (\rho_*(x) \nabla g_t(x)) dx + 2\alpha \int_{\Omega^2} \min(\hat{g}_t(x), \hat{g}_t(y)) (g_t(y) - g_t(x)) g_t(x) \rho_*(x) \rho_*(y) dx dy \\ &= -2 \int_{\Omega} |\nabla g_t(x)|^2 \rho_*(x) dx - \alpha \int_{\Omega^2} \min(\hat{g}_t(x), \hat{g}_t(y)) |g_t(y) - g_t(x)|^2 \rho_*(x) \rho_*(y) dx dy \\ &\leq -\alpha \int_{\Omega^2} \min(\hat{g}_t(x), \hat{g}_t(y)) |g_t(y) - g_t(x)|^2 \rho_*(x) \rho_*(y) dx dy \end{aligned} \quad (\text{B10})$$

where we used $(-\nabla U_* \cdot \nabla g_t + \Delta g_t) \rho_* = \nabla \cdot (\rho_* \nabla g_t)$ to reexpress the first integral in the second equality. If we denote $\hat{G}_t = \inf_{x \in \Omega} \hat{g}_t(x) \in [0, 1]$, [B10] implies

$$\frac{dD_t}{dt} \leq -\alpha \hat{G}_t \int_{\Omega^2} |g_t(y) - g_t(x)|^2 \rho_*(x) \rho_*(y) dx dy = -2\alpha \hat{G}_t D_t, \quad (\text{B11})$$

where we used the normalization conditions $\int_{\Omega} g_t(x) \rho_*(x) dx = \int_{\Omega} \hat{\rho}_t(x) dx = 1$. As a result, using Gronwall inequality we deduce

$$D_t \leq D_0 e^{-\alpha \int_0^t \hat{G}_s ds}. \quad (\text{B12})$$

This equation indicates that $D_t \rightarrow 0$ as $t \rightarrow \infty$ as long as $\int_0^t \hat{G}_s ds \rightarrow \infty$. That is, convergence can only fail if $\hat{G}_t = o(t^{-1})$ as $t \rightarrow \infty$, and it is guaranteed otherwise. Convergence is also exponential asymptotically, as long as \hat{G}_t remains bounded away from 0 as $t \rightarrow \infty$.

To get a more explicit convergence rate, let us analyze [B12] in two subcases. First let us assume that the map is not trained, i.e. $\hat{g}_t(x) = \hat{g}(x)$ is fixed, and denote $\hat{G} = \inf_{x \in \Omega} \hat{g}(x) \in [0, 1]$. In this case, [B12] reduces to

$$D_t \leq D_0 e^{-2\alpha \hat{G} t} \quad (\hat{g}_t = \hat{g} \text{ fixed}). \quad (\text{B13})$$

Note that this bound is only nontrivial if $\hat{G} > 0$. Even if that is the case, the rate in [B13] can be pretty poor if \hat{G} is very small (e.g exponentially small in the input dimension d), which is to be expected if the map

is not trained. The best case scenario is of course the idealized situation when $\hat{G} = 1$, which requires that $\hat{g} = 1$ (i.e. $\hat{\rho} = \rho_*$) because of the normalization conditions $\int_{\Omega} \hat{g}(x) \rho_*(x) dx = \int_{\Omega} \hat{\rho}(x) dx = 1$: this case is the continuous equivalent of the one step convergence of the discrete MCMC scheme with resampling from ρ_* .

Second let us assume that $\hat{g}_t = g_t$, that is the trained distribution instantaneously follows the walkers distribution at all times. In this case, [B12] reduces to

$$D_t \leq D_0 e^{-2\alpha \int_0^t G_s ds} \quad (\hat{g} = g_t), \quad (\text{B14})$$

where we denote

$$G_t = \inf_{x \in \Omega} \left(\frac{\rho_t(x)}{\rho_*(x)} \right) = \inf_{x \in \Omega} g_t(x) \in [0, 1]. \quad (\text{B15})$$

To make this bound explicit, let us consider the evolution of G_t . Denoting $x_t = \operatorname{argmin}_{x \in \Omega} g_t(x)$ so that $G_t = g_t(x_t)$, and using $\min(g_t(x_t), g_t(y)) = g_t(x_t) = G_t$, $\nabla g_t(x_t) = 0$, and $\Delta g_t(x_t) \geq 0$ by definition of x_t , from [B9] we have

$$\begin{aligned} \frac{dG_t}{dt} &= \partial_t g_t(x_t) + \dot{x}_t \cdot \nabla g_t(x_t) \\ &= \Delta g_t(x_t) + \alpha G_t \int_{\Omega} (g_t(y) - G_t) \rho_*(y) dy \\ &= \Delta g_t(x_t) + \alpha G_t - \alpha G_t^2 \\ &\geq \alpha G_t - \alpha G_t^2 \end{aligned} \quad (\text{B16})$$

where we used again the normalization conditions $\int_{\Omega} g_t(y) \rho_*(y) dy = \int_{\Omega} \rho_*(y) dy = 1$. Eq. [B16] implies that

$$\frac{1}{G_t - G_t^2} \frac{dG_t}{dt} \geq \alpha \quad (\text{B17})$$

which after integration gives

$$\log \left(\frac{G_t(1 - G_0)}{G_0(1 - G_t)} \right) \geq \alpha t \quad (\text{B18})$$

This means that we have

$$G_t \geq \frac{G_0}{G_0 + (1 - G_0)e^{-\alpha t}}. \quad (\text{B19})$$

Inserting this equation in [B14] and performing the integral explicitly gives

$$D_t \leq \frac{D_0}{(G_0(e^{\alpha t} - 1) + 1)^2}. \quad (\text{B20})$$

This bound is only nontrivial if $G_0 \in (0, 1]$. By performing all the time integration on $[t_0, t]$ we can obtain [16], which is now nontrivial if $G_{t_0} \in (0, 1]$.

Appendix C: Augmentation by Resampling via Birth-Death

In the augmented MCMC procedure described in main text, we mix Langevin and Metropolized resampling with the normalizing flow. Alternatively, we can perform this resampling via a birth-death process where mass is continuously transported non-locally by resampling configurations with the push-forward $\hat{\rho}$. This yields the following evolution equation for the sample density ρ_t

$$\partial_t \rho_t = \nabla \cdot (\rho_t \nabla U_* + \nabla \rho_t) - \alpha (V(\rho_t) - \bar{V}_t) \hat{\rho}_t \quad (\text{C1})$$

where we defined

$$V(\rho_t) = \frac{\rho_t}{\rho_*}, \quad \text{and} \quad \bar{V}_t = \int_{\Omega} V(\rho_t) \hat{\rho}_t dx = \int_{\Omega} \frac{\rho_t \hat{\rho}_t}{\rho_*} dx. \quad (\text{C2})$$

The interpretation of the added term proportional to α is somewhat subtle: this term amounts to adding or removing mass from ρ_t when ρ_t/ρ_* is respectively above or below its mean with respect to $\hat{\rho}_t$, a process that only stops when $V(\rho_t) = \bar{V}_t$, i.e. $\rho_t = \rho_*$. In practice, this process can be implemented by:

1. removing samples from our empirical representation of ρ_t with rate $\hat{\rho}_t/\rho_*$, which accounts for the effect of the term involving $V(\rho_t)$;
2. reinjecting into ρ_t the mass lost this way by generating samples from $\hat{\rho}_t$, which accounts for the effect of the term involving \bar{V}_t .

In terms of $g_t = \rho_t/\rho_*$ and $\hat{g}_t = \hat{\rho}_t/\rho_*$, [C1] reads

$$\partial_t g_t = -\nabla U_* \cdot \nabla g_t + \Delta g_t - \alpha(g_t - \bar{V}_t)\hat{g}_t, \quad \bar{V}_t = \int_{\Omega} \hat{g}_t g_t \rho_* dx, \quad (\text{C3})$$

and we can derive a convergence rate for the solution to this equation by looking at the evolution of the Pearson χ^2 -divergence of ρ_t with respect to ρ_* defined in [14]. Treating the Langevin terms as we did in Eq. [B10] we obtain

$$\frac{dD_t}{dt} \leq -2\alpha \int_{\Omega} (g_t - \bar{V}_t) g_t \hat{g}_t \rho_* dx = -2\alpha \int_{\Omega} g_t^2 \hat{g}_t \rho_* dx + 2\alpha \bar{V}_t^2 = -2\alpha \int_{\Omega} |g_t - \bar{V}_t|^2 \hat{g}_t \rho_* dx. \quad (\text{C4})$$

Since $\int_{\Omega} \bar{V}_t (g_t - 1) \rho_* dx = 0$, the Pearson χ^2 -divergence can also be expressed as

$$D_t = \int_{\Omega} (g_t - \bar{V}_t)(g_t - 1) \rho_* dx \quad (\text{C5})$$

using Cauchy-Schwartz inequality we obtain

$$D_t^2 \leq E_t \int_{\Omega} |g_t - \bar{V}_t|^2 \hat{g}_t \rho_* dx \quad \text{where} \quad E_t = \int_{\Omega} \frac{|g_t - 1|^2}{\hat{g}_t} \rho_* dx. \quad (\text{C6})$$

Using this inequality in [C4] we deduce

$$\frac{dD_t}{dt} \leq -2\alpha E_t^{-1} D_t^2 \quad (\text{C7})$$

which implies

$$D_t \leq \frac{D_0}{1 + 2\alpha D_0 \int_0^t E_s^{-1} ds}. \quad (\text{C8})$$

This equation implies convergence as long as $\int_0^t E_s^{-1} ds \rightarrow \infty$ as $t \rightarrow \infty$. In particular, if $\sup_{t \geq 0} E_t \leq M < \infty$, assuming $M > 0$, we have

$$D_t \leq \frac{D_0}{1 + 2\alpha D_0 M^{-1} t}. \quad (\text{C9})$$

Since

$$E_t \leq \int_{\Omega} \frac{g_t^2}{\hat{g}_t} \rho_* dx + \int_{\Omega} \frac{1}{\hat{g}_t} \rho_* dx = \int_{\Omega} \frac{\rho_t^2}{\hat{\rho}_t} dx + \int_{\Omega} \frac{\rho_*^2}{\hat{\rho}_t} dx \quad (\text{C10})$$

we see that Eq. [C9] holds as long as the Pearson χ^2 -divergences of ρ_t with respect to $\hat{\rho}_t$ and ρ_* with respect to $\hat{\rho}_t$ remain bounded at all times.

A more concrete convergence rate can be obtained if we assume again that $\hat{\rho}_t = \rho_t$ (i.e. $\hat{g}_t = g_t$). In this case

$$E_t = \int_{\Omega} \frac{\rho_*}{g_t} dx - 1 = \int_{\Omega} \frac{\rho_*^2}{\rho_t} dx - 1 \quad (\text{C11})$$

i.e. it is the Pearson χ^2 -divergences of ρ_* with respect to ρ_t . Since this divergence is known to be non-increasing if ρ_t solves a master equation with ρ_* as fixed point (see below for a proof), we have $E_t \leq E_0$ for $t \geq 0$, which implies that

$$D_t \leq \frac{D_0}{1 + 2\alpha D_0 E_0^{-1} t}. \quad (\text{C12})$$

This bound is only nontrivial if $D_0 < \infty$ and $E_0 < \infty$. More generally if there exists a $t_0 \geq 0$ such that $D_{t_0} < \infty$ and $E_{t_0} < \infty$, a similar argument gives

$$\forall t \geq t_0 \quad : \quad D_t \leq \frac{D_{t_0}}{1 + 2\alpha D_{t_0} E_{t_0}^{-1} (t - t_0)}. \quad (\text{C13})$$

This uniform bound asymptotically reduces to

$$D_t \leq \frac{1}{2} E_{t_0} (\alpha t)^{-1} \quad \text{for } 2\alpha D_0 E_{t_0}^{-1} t \gg 1. \quad (\text{C14})$$

The constant E_{t_0} emphasizes the importance of having initial samples close to metastable basins of the target distribution to guarantee that $E_{t_0} < \infty$. This further highlights that the methods we explore here are not generically suitable for exploring distributions for which the metastable basins are not known (cf. Appendix D).

a. *Verification that $E_t = \int_{\Omega} \rho_*^2 / \rho_t dx$ is non-increasing.* We have

$$\begin{aligned} \frac{dE_t}{dt} &= -2 \int_{\Omega} V^{-2}(\rho_t) \partial_t \rho_t dx \\ &= -2 \int_{\Omega} V^{-2}(\rho_t) \nabla \cdot (\rho_* \nabla V(\rho_t)) dx + 2\alpha \int_{\Omega} V^{-2}(\rho_t) (V(\rho_t) - \bar{V}_t) \hat{\rho}_t dx \\ &= -4 \int_{\Omega} V^{-3}(\rho_t) |\nabla V(\rho_t)|^2 \rho_* dx - 2\alpha \int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx + 2\alpha \bar{V}_t \int_{\Omega} V^{-2}(\rho_t) \hat{\rho}_t dx \\ &\leq 2\alpha \int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx - 2\alpha \bar{V}_t \int_{\Omega} V^{-2}(\rho_t) \hat{\rho}_t dx. \end{aligned} \quad (\text{C15})$$

Using Jensen's inequality,

$$\int_{\Omega} V^{-2}(\rho_t) \hat{\rho}_t dx \geq \left(\int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \right)^2 \quad (\text{C16})$$

and so we can deduce from the last inequality in [C15] that

$$\begin{aligned} \frac{dE_t}{dt} &\leq 2\alpha \int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx - 2\alpha \bar{V}_t \left(\int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \right)^2 \\ &= 2\alpha \left(\int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \right) \left(1 - \bar{V}_t \int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \right). \end{aligned} \quad (\text{C17})$$

Using Jensen's again,

$$\int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \geq \left(\int_{\Omega} V(\rho_t) \hat{\rho}_t dx \right)^{-1} = \bar{V}_t^{-1} \quad (\text{C18})$$

which implies

$$1 - \bar{V}_t \int_{\Omega} V^{-1}(\rho_t) \hat{\rho}_t dx \leq 0. \quad (\text{C19})$$

Using this equation in [C17], we deduce

$$\frac{dE_t}{dt} \leq 0. \quad (\text{C20})$$

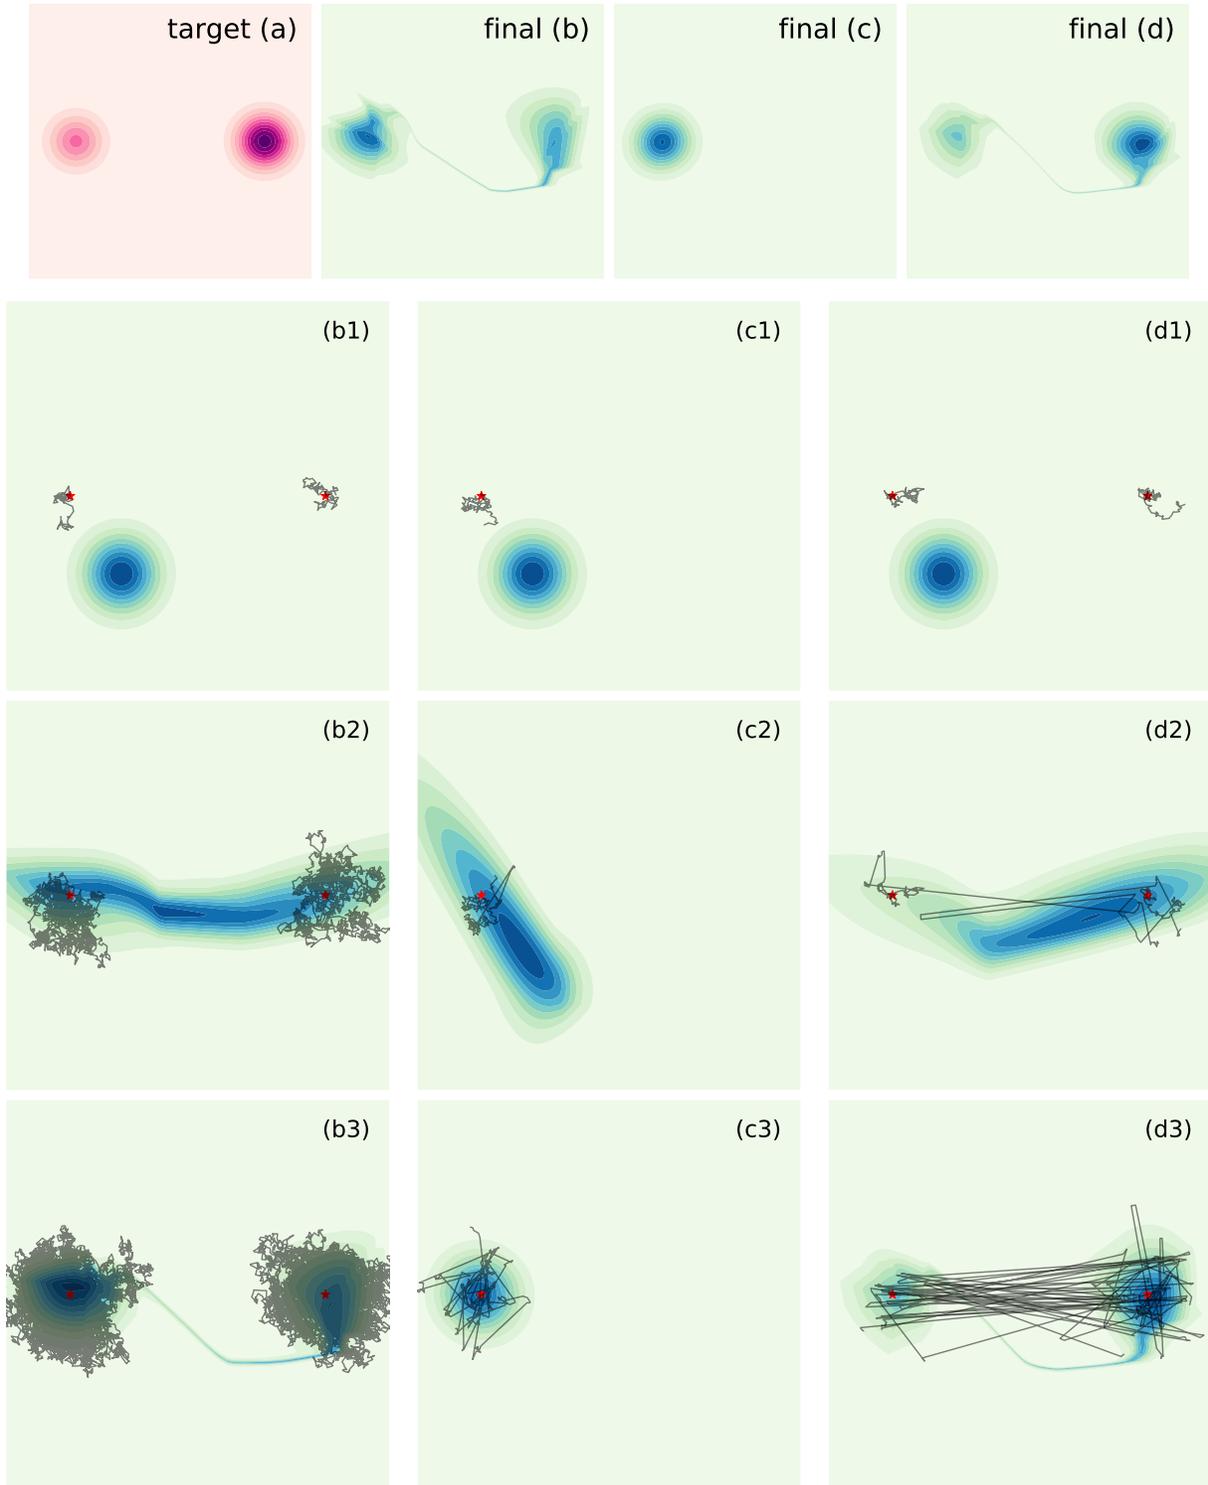


FIG. A1. **Advantage of resampling and necessity of good initialization.** (a) Target 2d mixture of two Gaussians. In all other panels, the push-forward density modelled by the normalizing flow is shown in blue. The trajectory of representative walkers of the concurrent sampling and training is plotted in black from the initialization marked by a red star. See Appendix E for details on the scenarios. (b) Relative weights of components are not recovered in final density learned using only local sampling with equal number of walkers starting in both components. (b1) through (b3) Early, middle, and later stage of training in this setting. The trajectory of walkers starting at the red star does not mix between modes. (c) Even if the combined sampling method is used, with nonlocal resampling, if walkers are initialized only in one mode, the second one is not learned by the flow. (c1) through (c3) As the flow learns to model one mode some non local moves are accepted by the walker but no proposition allows mixing with and modelling of the second mode. (d) When using the nonlocal resampling and proper initialization in both modes, the final learned density properly models the statistical weights of both modes. (d1) through (d3) The progressive learning of both modes by the flow allows more and more efficient mixing along the training.

Appendix D: Sampling unvisited metastable states

Once the normalizing flow T is trained to sample a given metastable basin, it is generally not suitable to discover unvisited metastable basins (i.e., those not in the data set). We illustrate this phenomenon by considering the probability of sampling a configuration near ϕ_- , the negative spin solution of the stochastic Allen-Cahn model with a map T_+ that has been training only on samples from the ϕ_+ metastable solution. First, we note that our analysis of the maps (Fig. 1) illustrates that the map is local; in this case, it means that when N is large (and hence we neglect the boundary terms, which will not change the overall scaling of the calculation that follows), the map will essentially be additive

$$T_+(\phi_i) = \phi_i + 1. \quad (\text{D1})$$

With this trained map T_+ , it is straightforward to see that the probability of sampling a configuration in the ϕ_- basin is actually worse than it would be with the untrained identity map. Denoting a configuration $\phi = \{\phi_i\}_{i=1}^N$, we have

$$P_B(\bar{T}_+(\phi)) = \exp\left(-\frac{a\beta}{2\Delta s} \sum_{i=1}^{N+1} (\phi_i - \phi_{i-1})^2 - \frac{\beta\Delta s}{4a} \sum_{i=1}^N (\phi_i - 1)^2\right), \quad (\text{D2})$$

where the inverse map in this case simply subtracts 1 and has unit Jacobian. When N is large, the probability of sampling in the ϕ_- metastable basin can be estimated by a Laplace approximation and it is easy to see that the integral over the possible states in the basin will be dominated by the configuration in which all the spins are aligned. Hence, we can estimate this probability by computing the probability of the configuration in which $u_i = -1$ for all $i = 1, \dots, N$, which we call u_- . This probability decays exponentially with the dimension,

$$P_B(\bar{T}_+(\phi_-)) = \exp(-\beta a^{-1} \Delta s N). \quad (\text{D3})$$

In the stochastic Allen-Cahn model, we take $\Delta s = 1/N$ so that the free energy barrier between the two metastable configurations is intensive (it should be noted that free energy barriers are typically extensive with system size for phase transitions). When βb is large ($\beta = 20$, $a = 1/10$ in our examples) it requires a rare event to spontaneously sample a configuration away from the metastable basin present in the original data. While this calculation is specific to the Allen-Cahn model, we have observed the same trends in other test systems, suggesting that, generically, learning about a given metastable basin reinforces bias for those samples and could hinder exploration of other basins of a target probability distribution. Taken together, these observations emphasize the importance of *a priori* identification of metastable states of interest.

Appendix E: Models, Computational Details and Additional Tests

a. Two-dimensional Gaussian mixture example We illustrate success and failures of concurrent training/sampling on the simple example of a Gaussian mixture in two dimensions with two modes (see Fig. A1(a)). Both components have covariance σI with $\sigma = 1$ and their means are separated by 10σ . The rightmost component is twice as likely as the leftmost. We choose the standard normal distribution as base measure ρ_B . At initialization the push-forward $\hat{\rho}$ is approximately equal to ρ_B and does not overlap with either of the target modes. We compare three scenarios: (b), (c) and (d) in Fig. A1. In scenario (b), we run Alg. 1 omitting lines 7-8, that is using exclusively the local sampling kernel, starting from an equal number of chains initialized in each mixture component. As the local chains fail to mix between modes, neither the learned density nor the samples reflect the relative statistical weights of the target. In scenario (c), Alg. 1 is run properly but with initial chains only in the left-most mode. The right-most mode is never sampled by the procedure. Finally in scenario (d) Alg. 1 is run with chains initialized in both modes and accurate and efficient sampling is achieved. For all cases above, batches of 400 samples from 40 independent walkers updated 10 times are used to compute each gradient step on the parameters of the normalizing flow. The optimization is done using Adam for 1500 iterations with a learning rate of 0.005. In scenario (d), acceptance rate of the push-forward proposals reaches 80 – 85%. The residual connection between modes in (d3) originates from the transformation of the uni-modal base measure into the bi-modal target. Additional iterations would make it thinner. These artefacts can be eliminated using stochastic normalizing flows [49], for which however the push-forward probability is not analytically tractable anymore.

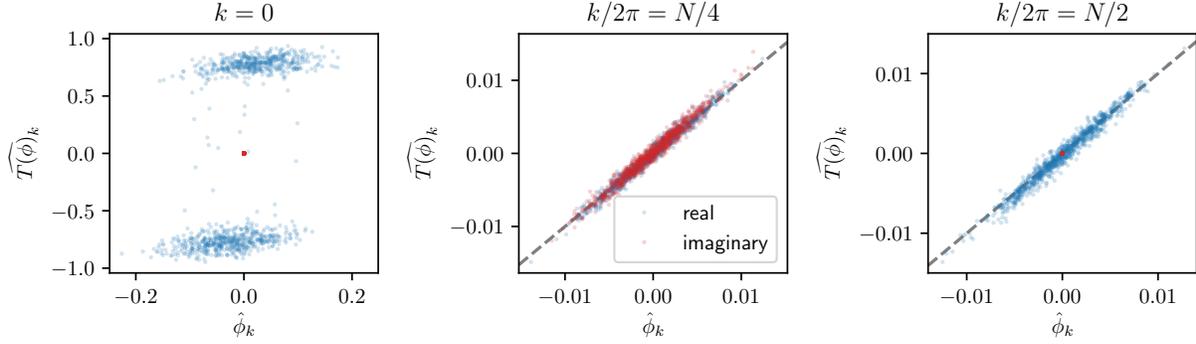


FIG. A2. **Diagnostic of the learned map T to sample the stochastic Allen-Cahn model on Fourier modes.** Correspondence between Fourier coefficients of base samples $\hat{\phi}_k$ with Fourier coefficients of the corresponding mapped fields $T(\hat{\phi})_k$, at low $k = 0$, intermediate $k = N\pi/2$ and high $k = N\pi$ frequency. Plots are done with 1000 independent samples from the base measure.

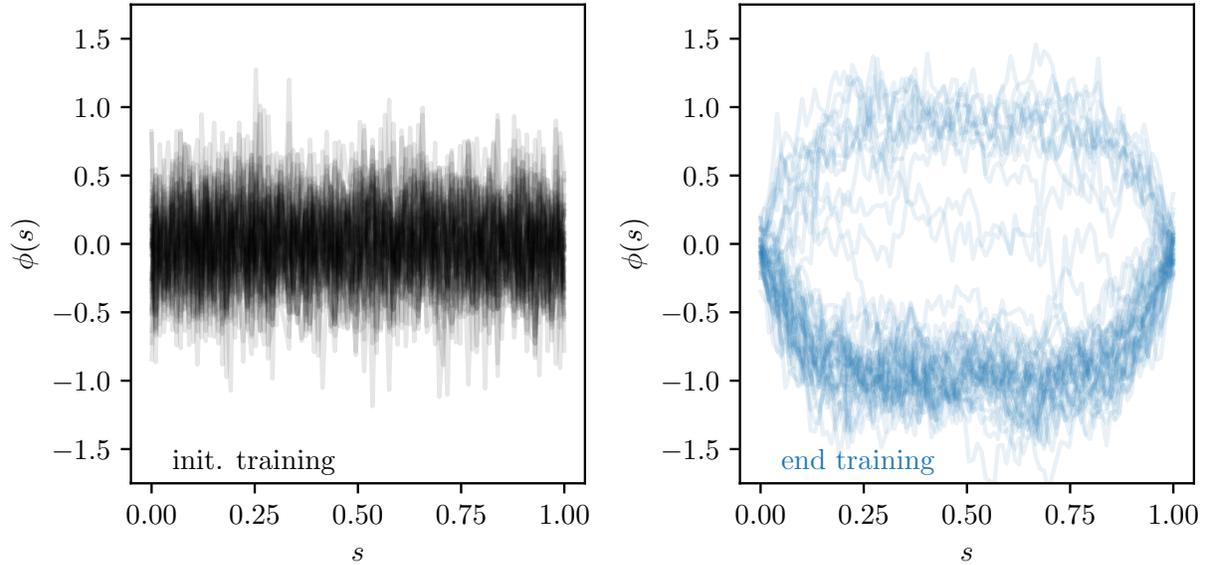


FIG. A3. **Failure of the uniformed prior.** Push-forward samples from the “uniformed” base distribution through the map T at the beginning and at the end of training. The initial walkers were initialized at 10% in ϕ_+ and 90% in ϕ_- . Final samples are roughly locating ϕ_+ and ϕ_- . Yet there is no rebalancing of the modes (50% in each mode) because of the Hamiltonian invariance to $\phi \rightarrow -\phi$ since almost no generated samples is accepted.

b. Stochastic Allen-Cahn experiments In our experiments the energy [18] is discretized on a grid with N points. The configuration is then a set of coupled continuous spins $\{\phi_i\}_{i=1}^N$ according to the discretized Hamiltonian

$$U_*(\phi) = \frac{a\beta}{2\Delta s} \sum_{i=1}^{N+1} (\phi_i - \phi_{i-1})^2 + \frac{\beta b \Delta s}{4} \sum_{i=1}^N (1 - \phi_i^2)^2 \quad (\text{E1})$$

with $\Delta s = 1/N$, and Dirichlet boundary conditions $\phi_0 = \phi_{N+1} = 0$. The discretized Hamiltonian of the base distribution similarly reads

$$U_B(\phi) = \frac{\beta a}{2\Delta s} \sum_{i=1}^{N+1} (\phi_i - \phi_{i-1})^2 + \frac{\beta b \Delta s}{2} \sum_{i=1}^N \phi_i^2, \quad (\text{E2})$$

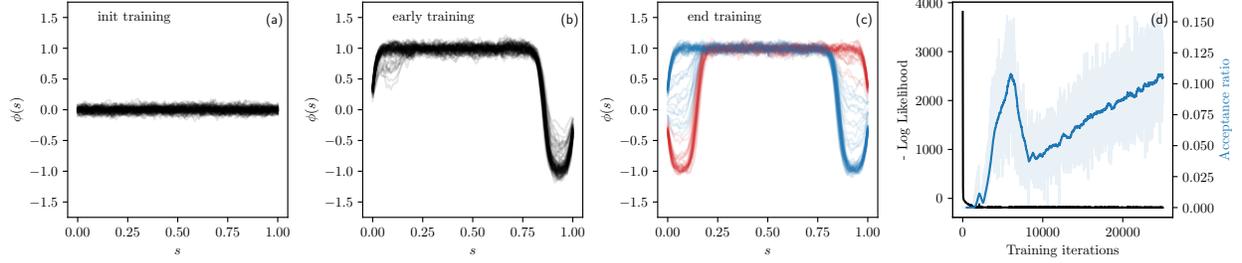


FIG. A4. **Concurrent training and sampling of the stochastic Allen-Cahn model with tilted Hamiltonian (E4).** (a,b,c) 100 independent samples from the push-forward $\hat{\rho}$ at different stage of training. The map gradually adjusts from generating samples close to the base measure samples to generating samples featuring a domain wall corresponding to the two modes with $\bar{\phi} = 0.7$. Colors are used to help distinguish between them. (d) In black, left axis, the negative log likelihood decreases rapidly along the gradient descent steps. In blue, right axis, the acceptance ratio features a drop in the middle of training when the map shifts from concentrating on a single mode to properly capture both states. The acceptance rate is of roughly 10% at the end of this run and could likely be further improved by pursuing the gradient descent.

with the same boundary conditions. Unless otherwise stated, we take $a = 0.1$, $b = 1/a = 10$, inverse temperature $\beta = 20$ and discretization $N = 100$.

During concurrent training/sampling, we gather batches of 1000 samples from 100 independent walkers updated 10 times to compute each gradient step on the parameters of the normalizing flow. The optimization is done using Adam for 10^5 iterations with a learning rate of 0.001. In Fig. 2 we report the evolution of the objective function, the acceptance ratio in the Metropolis steps and the autocorrelation time of the sampling chains during training, as well as the configurations visited by a single walkers in 10 steps at the end of training, demonstrating fast mixing. Fig. A2 examines Fourier modes of samples from the push-forward $\hat{\rho}$, at low and high frequencies, with respect to the Fourier decomposition of the base measure sample from which they originate. At high frequencies, the map implements the identity.

The most common implementations of normalizing flow use instead a standard Gaussian distribution for the base. In the continuous limit, the Hamiltonian of this “uniformed” base distribution is

$$U_{\text{BU}} = \frac{\beta}{2a} \int_0^1 \phi^2(s) ds, \quad (\text{E3})$$

with the coupling term of the target [18]. This defines a white-in-space process which gives back the standard multivariate Gaussian when discretized as $U_{\text{BU}}(\phi) = \frac{1}{2} \beta a^{-1} \Delta s \sum_{i=1}^N \phi_i^2$. Poor performance are obtained when using U_{BU} as shown on Fig. A3 for $N = 100$. As the discretization is refined, it gets harder and harder for the map to account for the fine scale structure when it is not encoded in the base distribution. Even after a long run of Alg. 1 generated samples are not probable. As a result almost none is accepted in the resampling step and there is no mixing between modes.

Finally, we also demonstrate the possibility to sample unlikely configurations with the presented algorithm using a biased measure of the stochastic Allen-Cahn model. As an example we consider the Allen-Cahn Hamiltonian (18) tilted towards configurations with a given spatial average $\bar{\phi}$,

$$U_{*,\lambda,\bar{\phi}}[\phi] = U_*[\phi] + \lambda\beta \left(\int_0^1 \phi(x) dx - \bar{\phi} \right)^2, \quad (\text{E4})$$

where λ is a positive Lagrange multiplier. By setting $\bar{\phi}$ away from ± 1 , the tilt favors configurations with domain walls. In the experiment displayed on Fig. A4, for which $a = 0.2$, $b = 1/a$, $\beta = 200$, $\bar{\phi} = 0.7$ and $\lambda = 10^4$, we can see that the normalizing flow learns to generate the two symmetric configurations with one domain wall thermodynamically favored by the tilted measure.

c. Nonequilibrium transition path sampling experiments To sample [21], as described in the main text, we discretize the path into N points with a time step Δt . We then carry out Langevin dynamics on the

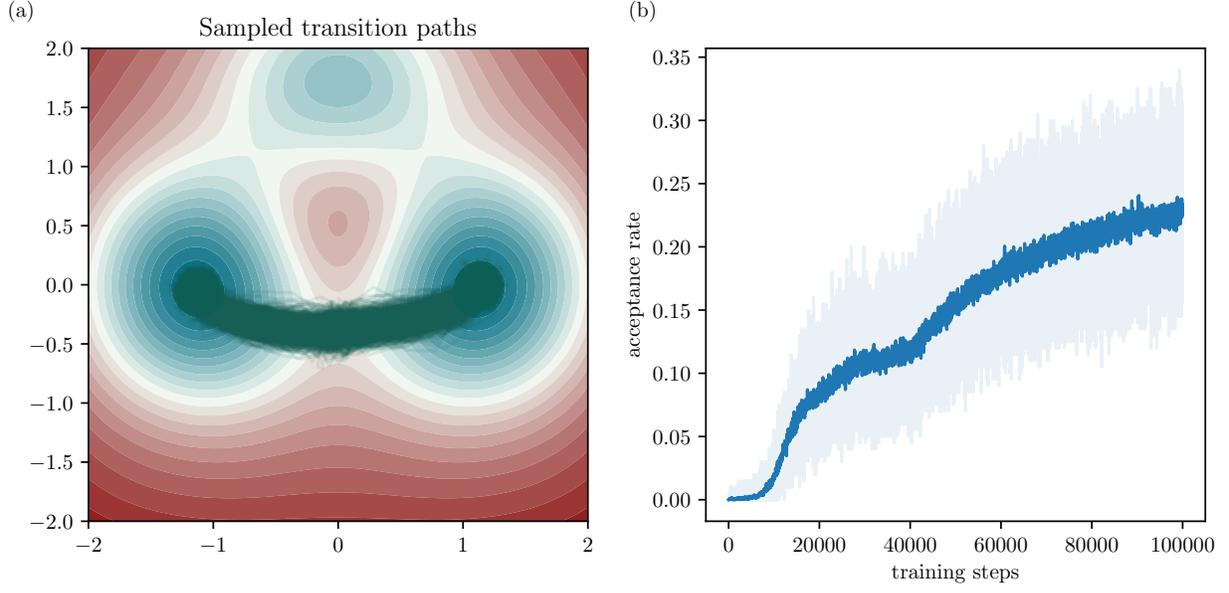


FIG. A5. **The map produces transition paths corresponding only to the lower channel when trained with resampling at each training step and a finite batch size** (as discussed in the main text). The resulting generator produces samples shown in (a). The acceptance rate is shown as a function of training time in (b). The rate of increase in acceptance plateaus when the initial samples from the upper channel have been eliminated.

target path action

$$S_* = \beta \Delta t \sum_{i=1}^{N-1} \left| \frac{x_{i+1} - x_i}{\Delta t} - b(x_i) \right|^2 \quad (\text{E5})$$

with an effective inverse temperature of $\beta \Delta t$, where the factor of Δt in the effective temperature ensures the proper scaling in the limit $\Delta t \rightarrow 0$ and the boundary conditions $x_0 = x_A$ and $x_N = x_B$ are imposed. The base measure has action

$$S_B = \beta \Delta t \sum_{i=1}^{N-1} \left| \frac{x_{i+1} - x_i}{\Delta t} \right|^2. \quad (\text{E6})$$

Langevin dynamics on the path action, written pointwise for $i = 1, \dots, N - 1$ evolves as

$$x_i(k+1) = x_i(k) - \frac{\partial S_*(x(k))}{\partial x_i} \tau + \sqrt{2\tau} \eta_i(k), \quad (\text{E7})$$

where $\tau > 0$ is the time step of the path-space Langevin dynamics. In all experiments $\beta = 4$, $N = 100$, $\Delta t = 6 \times 10^{-3}$, $\tau = 5 \times 10^{-5}$ and

$$b(x) = -\nabla V(x) + f(x) \quad (\text{E8})$$

with

$$V(x) = \sum_{i=1}^4 A_i e^{(x - \mu_i)^2} \quad (\text{E9})$$

with $A_1 = 30$, $A_2 = -30$, $A_3 = -50$, $A_4 = -50$ and $\mu_1 = (0, 1/3)^T$, $\mu_2 = (0, 5/3)^T$, $\mu_3 = (-1, 0)^T$, $\mu_4 = (1, 0)^T$. The nonconservative part is given by

$$f(x) = c(-x_2, x_1)^T \quad (\text{E10})$$

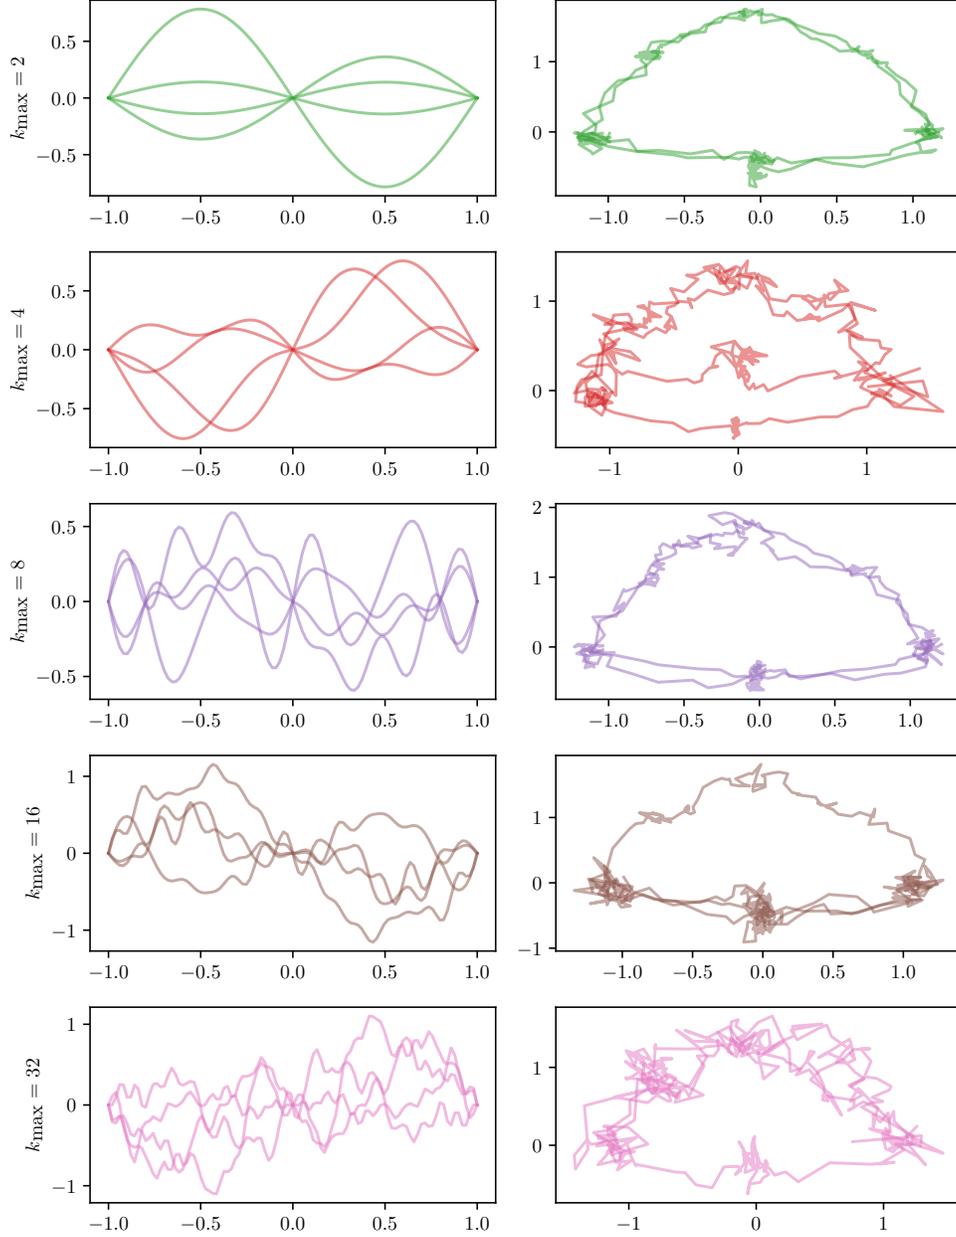


FIG. A6. The first four terms of the stochastic Karhunen-Loève expansion of the Brownian bridge are plotted over 5 realizations in (a). The map T , with parameters optimized from the training procedure with no resampling (cf. Fig 3), applied to the terms of this expansion are shown in (b).

with $c = 2.5$. We examine transition paths conditioned with $x_A = (-1, 0)$ and $x_B = (0, 1)$ and take the Brownian bridge process connecting these two points as the corresponding base measure.

To assess the smoothness and locality of the map, we plotted $T(B_i^{(k_{\max})})$ for a set of paths B_i $i = 1, \dots, 4$ which come from a truncated Karhunen-Loève expansion of the Brownian bridge process. Shown in Fig. A6, the left hand side plots

$$B_i^{(k_{\max})} = \sum_{k=1}^{k_{\max}} Z_k \frac{\sqrt{2} \sin(\pi k x)}{\pi k} \quad (\text{E11})$$

with $Z_k \sim \mathcal{N}(0, 1)$ for increasing values of k_{\max} . These plots illustrate that the high frequency fluctuations visible at larger values of k_{\max} lead to path fluctuations of higher frequency.

d. Normalizing flow architecture We parametrize the map T as a RealNVP [31], for which inverse and Jacobian determinant can be computed efficiently. Its building block is an invertible affine coupling layer updating half of the variables,

$$x_{1:d/2}^{(k+1)} = e^{s(x_{d/2:d}^{(k)})} * x_{1:d/2}^k + t(x_{d/2:d}^{(k)}) \quad (\text{E12})$$

where $s(\cdot)$ and $t(\cdot)$ are learnable neural networks from $\mathbb{R}^{d/2} \rightarrow \mathbb{R}^{d/2}$, $*$ is a component-wise multiplication, and k indexes the depth of the network.

In all the scenarios of the Gaussian mixture above, we used RealNVPs with 6 pairs of coupling layers. All the translation $t(\cdot)$ and scaling neural networks $s(\cdot)$ are multi-layer perceptrons with depth 3 and hidden layers of size 100. For the optimization, we used batches of 400 samples from 40 independent walkers updated 10 times to compute gradients and update parameters with Adam for 1500 iterations with a learning rate of 0.005.

Experiments reported on the Allen-Cahn model use RealNVPs with 10 pairs of coupling layers and again ReLU multi-layer perceptrons with depth 3 and width 100 for scalings and translations. During concurrent training/sampling, we gather batches of 1000 samples from 100 independent walkers updated 10 times to compute each gradient step on the parameters of the normalizing flow. The optimization is done using Adam for 10^5 iterations with a learning rate of 0.001. The parameters of the RealNVPs used to generate Fig. 1, Fig. 2, Fig. A2, Fig. A3 and Fig. A4 are given in the git repository listed below.

For the nonequilibrium transition path sampling experiments, we used a RealNVP architecture with 10 pairs of coupling layers and MLPs of depth 3 with hidden layers of 100 neurons. The parameters of the models used to generate configurations in Fig. 3, Fig. A5, and Fig. A5 are included in the git repository. We optimized the network using a batch size of 200 and 1 Langevin step per normalizing flow step. We optimized the parameters using the Adam optimizer with a learning rate of 5×10^{-4} with parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 1 \times 10^{-8}$. We carried out the optimization for a total of 10^5 steps.

e. Code availability All code and model parameters are available under an MIT license in a github repository for this project: <https://github.com/marylou-gabrie/flonaco>, archived at <https://doi.org/10.5281/zenodo.4783701>.

Appendix F: Computing free energy differences

Once the normalizing flow has been learned to assist sampling, there are two possible methods for estimating expectations. Ref. [22] trains the map with a different procedure, but compares empirical averages from Metropolized unbiased samples with an importance sampling reweighing procedure with direct samples from the push-forward $\hat{\rho}$.

Here we focus on the special case of computing free energy differences between metastable states. Denoting by A and B two sets of configurations in Ω , the difference between their free energies is given by

$$\begin{aligned} -\Delta F_{AB} &= \log \frac{\int_{\mathbb{R}^d} \mathbb{1}_A(x) e^{-U_*(x)} dx}{\int_{\mathbb{R}^d} \mathbb{1}_B(x) e^{-U_*(x)} dx} \\ &= \log \mathbb{E}_*(\mathbb{1}_A) - \log \mathbb{E}_*(\mathbb{1}_B). \end{aligned} \quad (\text{F1})$$

In Fig. 1 we used the set $A = \{\phi : \int_0^1 \phi(s) ds > 0\}$ and $B = \{\phi : \int_0^1 \phi(s) ds < 0\}$ and denote $F_+ - F_- = \Delta F_{AB}$

We can obtain unbiased samples from ρ_* by running a Metropolis-Hasting MCMC with the fast mixing kernel $\hat{\pi}_T$ [10]. From n samples $\{x_i\}_{i=1}^n$, an estimator of the partial partition function and its variance are

$$\begin{aligned} \mathbb{E}_*(\mathbb{1}_A) &\approx \hat{Z}_A^{\text{MC}} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_A(x_i), \\ \text{Var}(\hat{Z}_A^{\text{MC}}) &\approx \frac{\widehat{\text{Var}}_n(\mathbb{1}_A(x_i))}{n_{\text{eff}}} = \frac{\widehat{\text{Var}}_n(\mathbb{1}_A(x_i))}{n/\tau_{\text{eff}}} \end{aligned} \quad (\text{F2})$$

where n_{eff} is the effective sample size deduced from the autocorrelation time of the chains τ_{eff} (see e.g. [6] for details). While the map allows fast mixing between modes, it need not be perfect as the local MCMC kernel can compensate when some mass is not well represented by the map.

Alternatively, we can use an importance sampling estimator drawing $\{x_i\}_{i=1}^n$ from $\hat{\rho}$,

$$\mathbb{E}_*(\mathbb{1}_A) \approx \hat{Z}_A^{\text{IS}}/Z_* = \frac{1}{Z_*} \sum_{i=1}^n \mathbb{1}_A(x_i) \hat{w}(x_i), \quad (\text{F3})$$

where we use unnormalized weights $\hat{w}_i = e^{-U_*(x_i)}/\hat{\rho}(x_i)$, and the unknown Z_* cancels out in the free energy difference estimator:

$$-\Delta F_{AB} \approx \log \left(\sum_{i=1}^n \mathbb{1}_A(x_i) \hat{w}_i \right) - \log \left(\sum_{i=1}^n \mathbb{1}_B(x_i) \hat{w}_i \right). \quad (\text{F4})$$

Here again the quality of the estimator can be monitored using an estimate of the effective sample size

$$n_{\text{eff}} = \frac{(\sum_{i=1}^n \hat{w}_i)^2}{\sum_{i=1}^n \hat{w}_i^2}. \quad (\text{F5})$$

For both types of estimators, the map learned in certain thermodynamics conditions can be leveraged to sample in distinct conditions without relearning. This is possible thanks to the Metropolis-Hasting accept-reject in the MCMC and the importance sampling weights that correct mismatches between the target and push-forward distributions. As conditions are moved away from the training setting, quality of the Monte Carlo estimators can be assessed via the variance estimators above.