

Conditioning Boltzmann generators for rare event sampling

Sebastian Falkner, Alessandro Coretti, Salvatore Romano, Phillip L. Geissler, Christoph Dellago

Angaben zur Veröffentlichung / Publication details:

Falkner, Sebastian, Alessandro Coretti, Salvatore Romano, Phillip L. Geissler, and Christoph Dellago. 2023. "Conditioning Boltzmann generators for rare event sampling." *Machine Learning: Science and Technology* 4 (3): 035050.
<https://doi.org/10.1088/2632-2153/acf55c>.



PAPER • OPEN ACCESS

Conditioning Boltzmann generators for rare event sampling

To cite this article: Sebastian Falkner *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 035050

View the [article online](#) for updates and enhancements.

You may also like

- [Probing reaction channels via reinforcement learning](#)
Senwei Liang, Aditya N Singh, Yuanran Zhu *et al.*
- [Normalizing flows as an enhanced sampling method for atomistic supercooled liquids](#)
Gerhard Jung, Giulio Biroli and Ludovic Berthier
- [Transition path sampling: throwing ropes over mountains in the dark](#)
Peter G Bolhuis, Christoph Dellago, Phillip L Geissler *et al.*



PAPER

Conditioning Boltzmann generators for rare event sampling

OPEN ACCESS

Sebastian Falkner¹, Alessandro Coretti², Salvatore Romano¹, Phillip L Geissler^{3,†} and Christoph Dellago^{2,*}

RECEIVED

19 May 2023

REVISED

3 August 2023

ACCEPTED FOR PUBLICATION

30 August 2023

PUBLISHED

22 September 2023

¹ Faculty of Physics & Vienna Doctoral School in Physics, University of Vienna, 1090 Vienna, Austria² Faculty of Physics, University of Vienna, 1090 Vienna, Austria³ Department of Chemistry, University of California, Berkeley, CA 94720, United States of America

† Deceased.

* Author to whom any correspondence should be addressed.

E-mail: christoph.dellago@univie.ac.at**Keywords:** rare events, transition path sampling, normalizing flows, statistical mechanics, enhanced samplingSupplementary material for this article is available [online](#)

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

Understanding the dynamics of complex molecular processes is often linked to the study of infrequent transitions between long-lived stable states. The standard approach to the sampling of such rare events is to generate an ensemble of transition paths using a random walk in trajectory space. This, however, comes with the drawback of strong correlations between subsequently sampled paths and with an intrinsic difficulty in parallelizing the sampling process. We propose a transition path sampling scheme based on neural-network generated configurations. These are obtained employing normalizing flows, a neural network class able to generate statistically independent samples from a given distribution. With this approach, not only are correlations between visited paths removed, but the sampling process becomes easily parallelizable. Moreover, by conditioning the normalizing flow, the sampling of configurations can be steered towards regions of interest. We show that this approach enables the resolution of both the thermodynamics and kinetics of the transition region for systems that can be sampled using exact-likelihood generative models.

1. Introduction

The exponential increase in computational power experienced by computers since the advent of molecular simulations has radically changed basically all aspects of the study of statistical mechanics via numerical experiments. Simulations of rare events have also benefited from such improvements, but progress in this area has relied even more on methodological developments rather than the exploitation of raw computing power. This is due to the intrinsic nature of rare events, which are phenomena that occur infrequently, but happen quickly if they occur. The resulting disparity between the long waiting times and the fastest time scales in the system is often so large that such processes cannot be simulated even on the fastest computers with straightforward methods. Examples are omnipresent in physics, chemistry and biology and include nucleation processes [1, 2], protein folding [3, 4], dynamics of ions in solution [5–8] and chemical reactions [9, 10]. All of these processes exhibit transitions between stable states separated by high energetic and/or entropic barriers. Resolving the thermodynamics and kinetics at the barrier top is the key challenge for understanding the rare event.

Over the years, many enhanced sampling methods were developed to focus the computational effort on regions of interest in phase space. For instance, when one aims to resolve the thermodynamic properties of a system, umbrella sampling [11] received widespread recognition. In this approach, a harmonic bias is added to the potential energy function, efficiently restricting the sampling to certain regions of configuration space. In contrast, when investigating the kinetics of a rare event, a properly weighted set of unbiased reactive trajectories is desired. Transition path sampling (TPS) [12] is an efficient strategy to achieve this goal by performing a Markov chain Monte Carlo simulation in trajectory space. A basic scheme for generating a new

path based on a previous one is the shooting move [13], where a point on the previous trajectory is randomly selected, possibly perturbed, and then integrated forward and backward in time until a stable state is reached. If the newly generated trajectory connects the stable states, it is accepted and used for the generation of the next path.

Despite improvements introduced by different shooting schemes [14–16], the foundation of these sampling approaches is the generation of a new path from the previous one. Therefore, these algorithms are inherently sequential and correlations between subsequently visited paths are inevitable. Even though a high acceptance rate may be achieved, a strong similarity between subsequent paths degrades the efficiency of sampling.

With recent developments in the field of generative neural networks [17–19], the sampling of independent equilibrium configurations from the Boltzmann distribution came into reach. In particular, normalizing flows [19, 20] have already been applied in physical sciences to solve a multitude of problems. These include free energy calculations [21], exploration of configuration space [22, 23], finding minimum energy paths [24], force field parametrization [25] and lattice field theory [26–29]. These models are of particular interest since they allow for unbiased estimation of physical observables [30]. More recently, conditioned flow-based models have been successfully applied in the context of physical problems to enhance sampling in systems studied using lattice field theory [31–33]. In this work, we propose a parallel sampling scheme to explore the reactive path space based on normalizing flows in the form of Boltzmann generators [22] for the generation of shooting points. The flow model is conditioned to steer the generation to regions of interest, which at the same time allows for the accurate reconstruction of free energy profiles.

The proposed path sampling scheme (see figure 1) starts by sampling points from a multivariate Gaussian distribution. These points are then transformed into shooting points using a conditioned Boltzmann generator. From these, trajectories are obtained by integrating forward and backward in time until a stable state is reached. The resulting paths are reweighted to obtain a properly weighted transition path ensemble. In the following, we will describe this algorithm in detail and demonstrate it using some illustrative models.

1.1. Flexible length TPS

The path ensemble targeted by our sampling procedure includes all reactive trajectories connecting stable states. Each trajectory is defined as a sequence of configurations $X(\tau) = \{x_0, x_{\Delta t}, x_{2\Delta t}, \dots, x_\tau\}$, where τ is the length of the path and is a multiple of the timestep Δt . Transition paths connect two given stable states, A and B, and they are required to have exactly one point in each of these states. Consequently, transition pathways have varying lengths τ .

Transition paths are sampled proportional to their statistical weight $P_{AB}[X(\tau)]$. Here we consider the probabilities within a small region dX^τ in path space. Accordingly, the probability of a reactive path $X(\tau)$ can be expressed as:

$$P_{AB}[X(\tau)] dX^\tau = \frac{1}{Z_{AB}} H_{AB}(x_0, x_\tau) \prod_{i=1}^{\tau/\Delta t - 1} \tilde{h}(x_{i\Delta t}) P[X(\tau)] dX^\tau, \quad (1)$$

where $H_{AB}(x_0, x_\tau)$ is unity if the trajectory connects states A and B in any order and is zero otherwise. More explicitly, this function is defined as

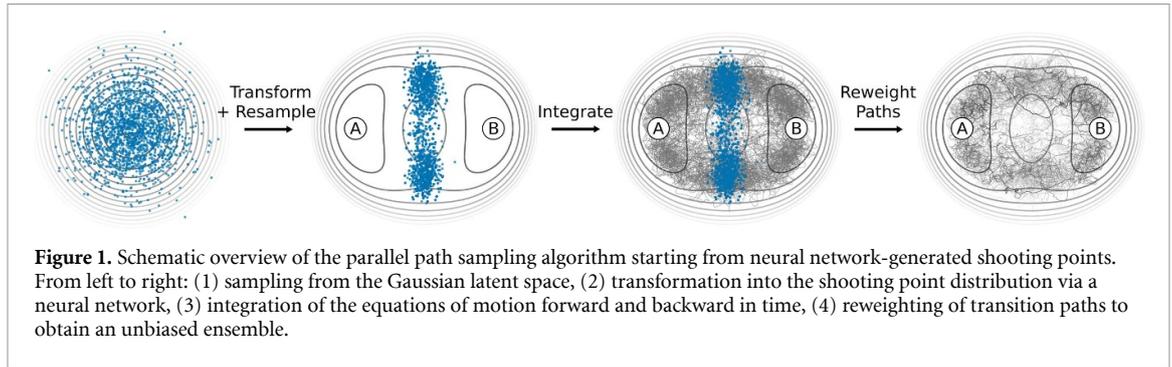
$$H_{AB}(x_0, x_\tau) = \begin{cases} 1 & \text{if } h_A(x_0)h_B(x_\tau) = 1 \\ & \text{or } h_A(x_\tau)h_B(x_0) = 1; \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here, h_A and h_B are population functions that return one if a point lies in state A and B, respectively, and vanish otherwise. The function $\tilde{h}(x)$ is defined as:

$$\tilde{h}(x) = \begin{cases} 0 & \text{if } h_A(x) = 1 \text{ or } h_B(x) = 1; \\ 1 & \text{otherwise,} \end{cases} \quad (3)$$

and acts as a constraint to focus attention only on values of τ that are comparable to a natural transition time. The normalizing factor Z_{AB} has the form of a partition function:

$$Z_{AB} = \sum_{\tau} \int dX^\tau H_{AB}(x_0, x_\tau) \prod_{i=1}^{\tau/\Delta t - 1} \tilde{h}(x_{i\Delta t}) P[X(\tau)]. \quad (4)$$



Assuming Markovian dynamics, the dynamical path probability $P[X(\tau)] dX^\tau$ is defined based on the equilibrium probability of the starting point $p_{\text{eq}}(x_0)$ and the short-time transition probabilities $p(x_{i\Delta t} \rightarrow x_{(i+1)\Delta t})$:

$$P[X(\tau)] dX^\tau = p_{\text{eq}}(x_0) \prod_{i=0}^{\tau/\Delta t - 1} p(x_{i\Delta t} \rightarrow x_{(i+1)\Delta t}) dX^\tau. \quad (5)$$

1.2. Parallel path sampling

Shooting moves are an integral part of most path sampling schemes. Their efficiency relies on the fact that shooting points in a region of high $p(\text{TP}|x)$, which is the probability of generating a transition path (TP) given a certain configuration x , lead to an efficient exploration of the path ensemble. Using Bayes' theorem, the transition path probability can be written as [34]:

$$p(\text{TP}|x) = \frac{p(x|\text{TP})p(\text{TP})}{p_{\text{eq}}(x)}, \quad (6)$$

where $p(x|\text{TP})$ is the density of x in the ensemble of points on transition paths, $p(\text{TP})$ is the fraction of time the system spends on transition paths and $p_{\text{eq}}(x)$ is the equilibrium probability density of x . Consequently, the transition path probability of x is defined up to a proportionality constant $p(\text{TP})$ by the ratio of its density in the ensemble of points on paths $p(x|\text{TP})$ and the equilibrium ensemble $p_{\text{eq}}(x)$. The scalability of TPS with shooting moves is limited by the inherently sequential nature of the sampling. The previous trajectory is indispensable for the generation of the new trajectory. To combine the efficiency of shooting moves with the possibility of parallel sampling, we propose an alternative algorithm to sample the transition path ensemble.

The basis of the scheme is a set of shooting points generated before the actual path sampling starts. These configurations can be sampled from an arbitrary distribution denoted as $p_{\text{SP}}(x)$, where SP stands for shooting point. From these shooting points, trajectories are obtained by integration forward and backward in time until a stable state is reached. As a result, the generation of paths becomes embarrassingly parallel because the trajectories are generated independently from each other. Fleeting trajectories initiated using configurations from $p_{\text{SP}}(x)$, however, do not correspond to a properly weighted transition path ensemble. It is important to note that this applies for any case where paths are generated from a set of in-advance-sampled points, no matter if the shooting point distribution is the result of a flow-based transformation or is obtained by well-known enhanced sampling methods. One difference are the missing population functions to distinguish paths that connect stable states from ones that end in the same state in both directions. Even more critically, paths that dwell a long time in high probability regions of $p_{\text{SP}}(x)$ are sampled preferentially. Accordingly, a reweighting factor $\Omega[X(\tau)]$ has to be included when calculating the expectation values of path observables:

$$\langle A[X(\tau)] \rangle \approx \frac{\sum_{i=1}^N \Omega[X(\tau)_i] A[X(\tau)_i]}{\sum_{i=1}^N \Omega[X(\tau)_i]}. \quad (7)$$

A similar path reweighting has already been successfully applied in other studies, e.g. in works by Daru and Stirling [35] and Menzl *et al* [36] for the calculation of rate constants.

To derive an expression for $\Omega[X(\tau)_i]$ in the context of our parallel sampling scheme, we first consider the generation probability of a trajectory obtained from an *a priori* sampled shooting point. By means of a

shooting move, a given trajectory can be generated from any of its points. Therefore, the total generation probability is the sum of the independent generation probabilities from each point on the trajectory:

$$P_{\text{gen}}[X(\tau)] dX^\tau = dX^\tau \sum_{k=0}^{\tau/\Delta t} \left[p_{\text{SP}}(x_{k\Delta t}) \frac{p_{\text{eq}}(x_0)}{p_{\text{eq}}(x_{k\Delta t})} \times [h_A(x_0) + h_B(x_0)] [h_A(x_\tau) + h_B(x_\tau)] \right. \\ \left. \times \prod_{j=1}^{\tau/\Delta t-1} \tilde{h}(x_{j\Delta t}) \prod_{i=0}^{\tau/\Delta t-1} p(x_{i\Delta t} \rightarrow x_{(i+1)\Delta t}) \right]. \tag{8}$$

The weight of reactive paths in equation (8) differs from the corresponding weight in the transition path ensemble (equation (1)) by the factor:

$$\Omega[X(\tau)] = \frac{1}{Z_{\text{AB}}} \left[\sum_{k=0}^{\tau/\Delta t} \frac{p_{\text{SP}}(x_{k\Delta t})}{p_{\text{eq}}(x_{k\Delta t})} \right]^{-1}. \tag{9}$$

A full derivation of the generation probability and of the reweighting factor is given in the Supplementary information (SI). Since we are solely interested in properly weighting a path relative to all others, the partition function Z_{AB} can be omitted. This leads to a tractable relative reweighting factor to recover a properly weighted transition path ensemble given a collection of trajectories generated from a distribution of shooting points. In the simplest case, one can choose the equilibrium distribution $p_{\text{eq}}(x)$ as the shooting point distribution so that $p_{\text{SP}}(x) \equiv p_{\text{eq}}(x)$ with the caveat that points already lying in a stable state must be sorted out. The reweighting factor then reduces to $(\tau/\Delta t + 1)^{-1}/Z_{\text{AB}}$. In this case, reactive paths are weighted by their inverse number of points. When all shooting points lie on a dividing surface and have weights according to the equilibrium distribution, the above reweighting factor reduces to the inverse number of crossings of the path with the surface, which agrees with the findings of Best and Hummer [34, 37]. Given an infinite number of samples from $p_{\text{SP}}(x)$ as a set of shooting points, ergodicity in transition path space is guaranteed if every configuration with a non-zero probability in $p_{\text{eq}}(x)$ also has a non-zero probability in $p_{\text{SP}}(x)$.

1.3. Targeted sampling using Boltzmann generators

Enhanced sampling revolves around the efficient exploration of low probability regions in configuration space. In standard equilibrium simulations, these regions are often not visited frequently enough to make accurate predictions about the thermodynamics of the system. Therefore, one common approach in enhanced sampling methods is to restrict the sampling to the region of interest and thereby focus computational resources. Often this is achieved by applying a bias towards the region of interest. For example, in umbrella sampling this bias is introduced in the form of a harmonic bias potential, which is added to the potential energy function of the system. This bias is defined using a collective variable, here denoted as $r(x)$, a bias center \bar{r} and a force constant k :

$$U^{\text{bias}}(x, \bar{r}) = U(x) + \frac{k}{2} [r(x) - \bar{r}]^2. \tag{10}$$

By the addition of the bias potential, configurations with a collective variable value close to \bar{r} will be sampled more often. The resulting configuration ensemble is referred to as an umbrella window. In the canonical ensemble, the probability of observing a configuration x given an applied bias potential centered at \bar{r} can be expressed as:

$$p^{\text{biased}}(x|\bar{r}) = \frac{1}{Z_x} \exp \left\{ -\beta \left[U(x) + \frac{k}{2} (r(x) - \bar{r})^2 \right] \right\}, \tag{11}$$

$$Z_x = \int dx \exp \left\{ -\beta \left[U(x) + \frac{k}{2} (r(x) - \bar{r})^2 \right] \right\}. \tag{12}$$

Boltzmann generators, as proposed by Noé et al [22], provide a way to obtain uncorrelated samples from such a distribution. They belong to the class of flow-based generative models and they allow to obtain unbiased samples from a given target distribution after reweighting. This unique feature makes them well-suited for our parallel path sampling scheme.

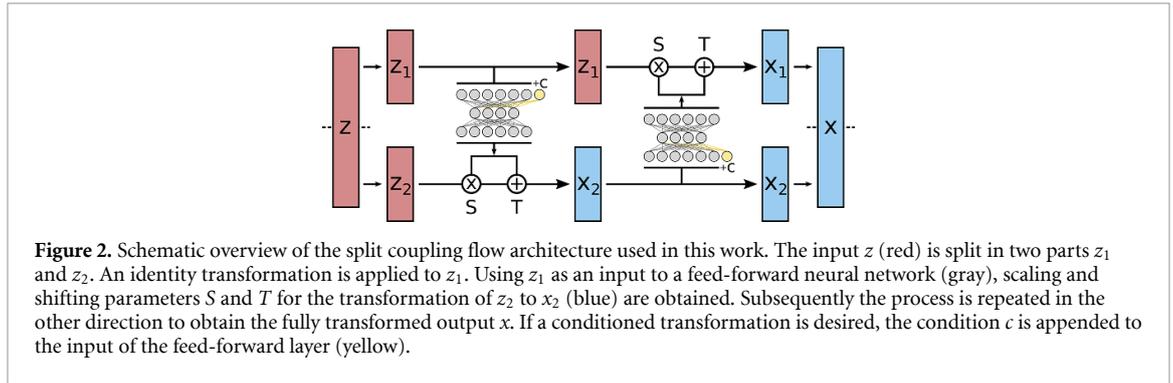


Figure 2. Schematic overview of the split coupling flow architecture used in this work. The input z (red) is split in two parts z_1 and z_2 . An identity transformation is applied to z_1 . Using z_1 as an input to a feed-forward neural network (gray), scaling and shifting parameters S and T for the transformation of z_2 to x_2 (blue) are obtained. Subsequently the process is repeated in the other direction to obtain the fully transformed output x . If a conditioned transformation is desired, the condition c is appended to the input of the feed-forward layer (yellow).

In flow-based models, a neural network learns an invertible coordinate transformation between an easy to sample latent distribution $p_z(z)$ and a complex data distribution $p_x(x)$:

$$x = f(z; \theta), \tag{13}$$

$$z = f^{-1}(x; \theta), \tag{14}$$

where x and z represent samples from the data space (denoted as a whole by x) and from the latent space (denoted as a whole by z), respectively, while θ represents the set of trainable network parameters that parameterize the transformation f .

The architecture of Boltzmann generators is based on a split-coupling flow using RealNVP blocks as proposed by Dinh *et al* [38]. Split-coupling flows allow to compute the determinant of the transformation's Jacobian efficiently [38]. Therefore, the distribution of neural-network generated samples can be expressed via the change of variable theorem:

$$q_x(x) = p_z [f^{-1}(x; \theta)] |\det J_{f^{-1}}(x; \theta)|, \tag{15}$$

$$q_z(z) = p_x [f(z; \theta)] |\det J_f(z; \theta)|, \tag{16}$$

where q represents the distributions generated by the network in the corresponding spaces, which will be, in general, different from p . The determinant of the Jacobian is tractable thanks to the particular construction of the network, as shown schematically in figure 2. The input is split into two parts, x_1 and x_2 . While an identity transformation is applied to one part, the other one is scaled and translated with parameters that are a function of the first. The generated distribution $q_x(x)$ is, in general, only an approximation to the Boltzmann distribution. However, since the probability of a generated sample can be obtained using equation (15), a statistical weight can be assigned to each generated configuration in order to correct for the bias and to recover the exact distribution. A simple choice for this reweighting factor is represented by the ratio between the reference and generated probability [22]:

$$\omega(x) = \frac{p_x(x)}{q_x(x)}. \tag{17}$$

For further discussions on the unbiased estimation of observables using flow-based models, we refer to the seminal work of Nicoli *et al* [30].

The generation of samples with probabilities according to equation (11) for a single bias center is possible in the framework of Boltzmann generators. However, it is rarely the case that a single window provides sufficient information on the rare event of interest. For this reason, the generation of samples should be possible at arbitrary bias centers.

Conditioning the transformation applied by the normalizing flow enables sampling at different bias centers using a single neural network. A simple scheme to condition a split-coupling flow architecture was proposed by Ardizzone *et al* [39] in relation to image generation. These ideas were then applied to physical problems in the context of lattice field theory by Gerdes *et al* [31] and Singha *et al* [32, 33]. In the model proposed by Ardizzone *et al* [39], the transformation is conditioned by concatenating the condition data c to the coupling layer network input, as indicated in figure 2. For the purpose of generating configurations at different bias centers with weights given by equation (11), the condition vector c corresponds to the bias center \bar{r} . This approach leaves the latent space distribution unconditioned and it imposes the condition

directly on the transformation which is then reflected on the generated distribution. The change of variable theorem then takes the form

$$q_x(x|\bar{r}) = p_z[f^{-1}(x|\bar{r};\theta)] |\det J_{f^{-1}}(x|\bar{r};\theta)|, \quad (18)$$

$$q_z(z|\bar{r}) = p_x[f(z|\bar{r};\theta)|\bar{r}] |\det J_f(z|\bar{r};\theta)|. \quad (19)$$

For all systems studied here, we use a multivariate Gaussian as a latent distribution $p_z(z)$ as in regular Boltzmann generators.

The training loss functions are formulated based on the forward and reverse Kullback–Leibler (KL) divergences between the generated and reference distributions. Conditioning of the transformation can then be incorporated in the definition of the loss functions for the training. In *training by example*, samples from different umbrella windows are transformed into Gaussian-distributed samples. Here, the training loss L_{fwd} is given by the conditional forward KL-divergence between the reference and generated data distribution $\text{KL}[p_x(x|r)||q_x(x|\bar{r};\theta)]$ (full derivation in SI):

$$\begin{aligned} L_{\text{fwd}} &= \mathbf{E}_{\bar{r}}\mathbf{E}_x[\log p_x(x|r) - \log q_x(x|\bar{r})] \\ &= \mathbf{E}_{\bar{r}}\mathbf{E}_x\left[\frac{1}{\sigma^2}\|f^{-1}(x|\bar{r};\theta)\|^2 - |\det J_{f^{-1}}(x|\bar{r};\theta)|\right] + \text{const.} \end{aligned} \quad (20)$$

During training, the expectation values $\mathbf{E}_{\bar{r}}\mathbf{E}_x$ are approximated using a set of configurations at discrete bias positions on the collective variable. These discrete positions should cover the regions of interest on the bias coordinate.

Training in the other direction, the *training by energy*, works by sampling from the latent Gaussian distribution and transforming to the desired umbrella windows. The conditional reverse KL-divergence between the reference and generated latent distribution $\text{KL}[p_z(z|r)||q_z(z|\bar{r};\theta)]$ is then minimized leading to the loss function L_{rev} (full derivation in SI):

$$\begin{aligned} L_{\text{rev}} &= \mathbf{E}_{\bar{r}}\mathbf{E}_z[\log p_z(z) - \log q_z(z|\bar{r};\theta)] \\ &= \mathbf{E}_{\bar{r}}\mathbf{E}_z\left[\beta U(f(z|\bar{r};\theta)) + \beta\frac{k}{2}[r(f(z|\bar{r};\theta)) - \bar{r}]^2 - |\det J_f(z|\bar{r};\theta)|\right] + \text{const.} \end{aligned} \quad (21)$$

Therefore, training on the reverse KL-divergence is performed by selecting bias centers of interest and sampling from the latent distribution. Latent points and corresponding bias centers are then transformed and parameters of the network are optimized with respect to equation (21). During training by energy, the network can be trained at different temperatures by adjusting the variance of the Gaussian latent space distribution [22]. The final loss function for the training can be computed as

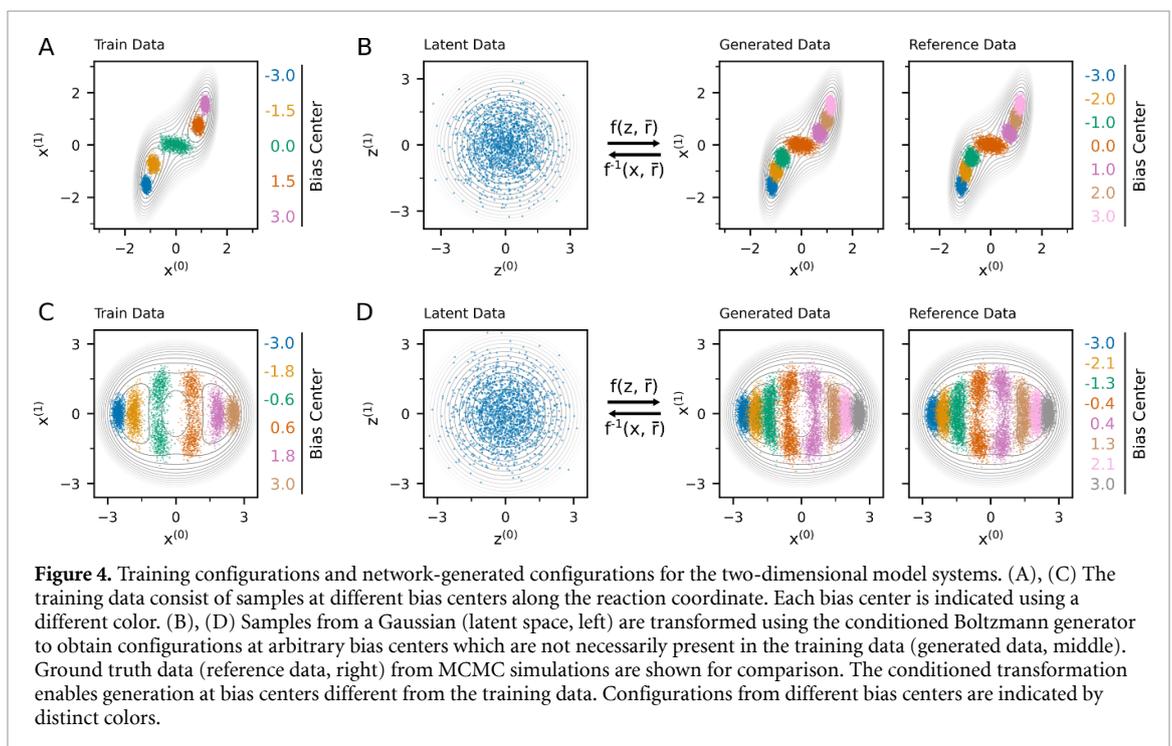
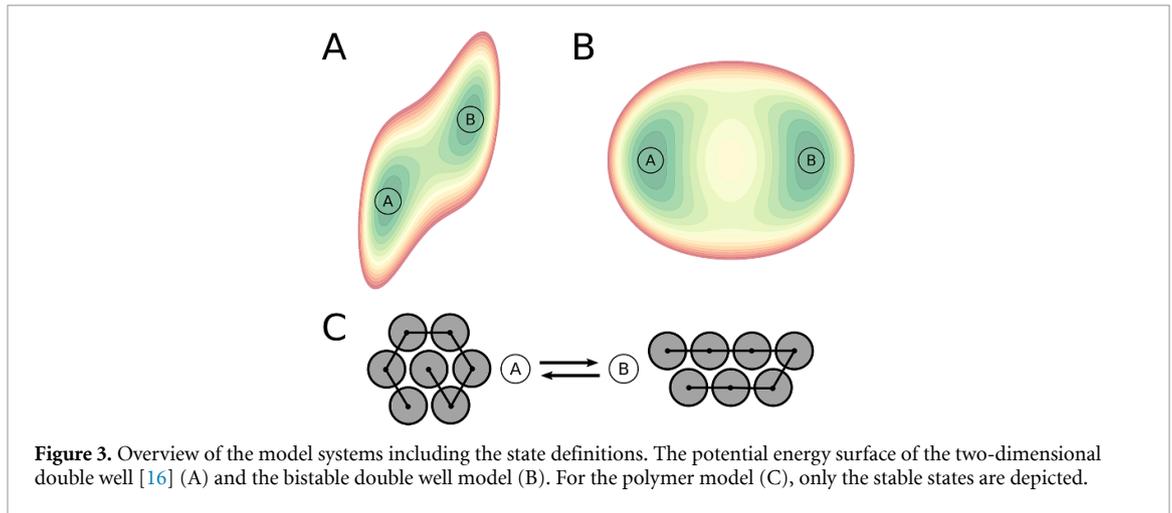
$$L = \lambda_{\text{fwd}}L_{\text{fwd}} + \lambda_{\text{rev}}L_{\text{rev}}, \quad (22)$$

where λ_{fwd} and λ_{rev} are weights used to tune the focus of the training. For a comprehensive discussion on the forward and reverse KL-divergence in flow-based sampling, we refer to [40].

2. Results

2.1. Resolving the barrier region

We first test the conditioned Boltzmann generators on a simple two-dimensional model [16] (figure 3(A), system parameters in SI). Here, the conditioning of the Boltzmann generator greatly improves the resolution of low probability regions in configuration space, as shown in figure 4 where the reaction coordinate $r(x) = x^{(0)} + x^{(1)}$ was used for the double well system. Analogous to umbrella sampling, a bias potential is applied to force the system to regions that are rarely seen in equilibrium at a given temperature. In the case of the original Boltzmann generator, low probability states can be included in the generated distribution by the introduction of a reaction coordinate loss [22]. Here, the entropy of samples projected on a reaction coordinate was maximized during training. While this is sufficient to encourage a broad sampling of the target distribution and to prevent a mode collapse, it does not allow a targeted sampling of low probability regions. For an accurate free energy estimate and especially for the generation of transition states, the sampling of specific low probability regions in configuration space must be enhanced. Due to the conditioning of the transformation, the generator can be steered to focus on certain regions in configurations space. In the following, we use the term network-generated configurations for samples generated by the *conditioned* Boltzmann generator if not stated otherwise.

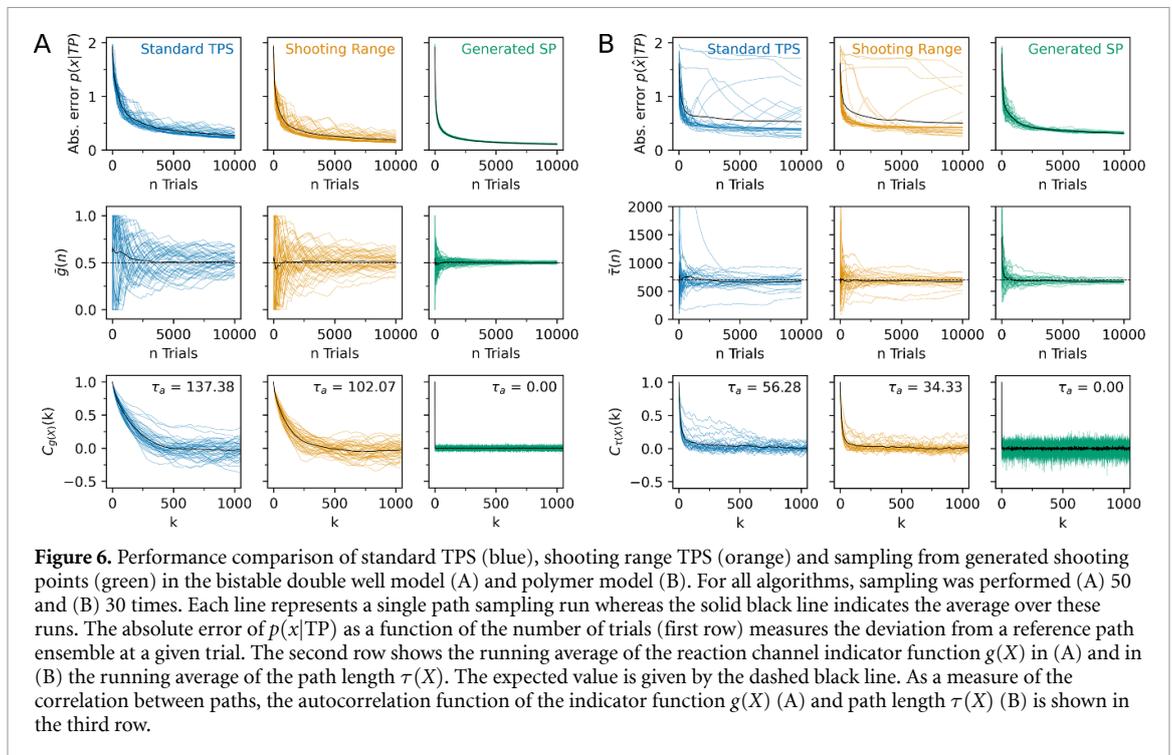
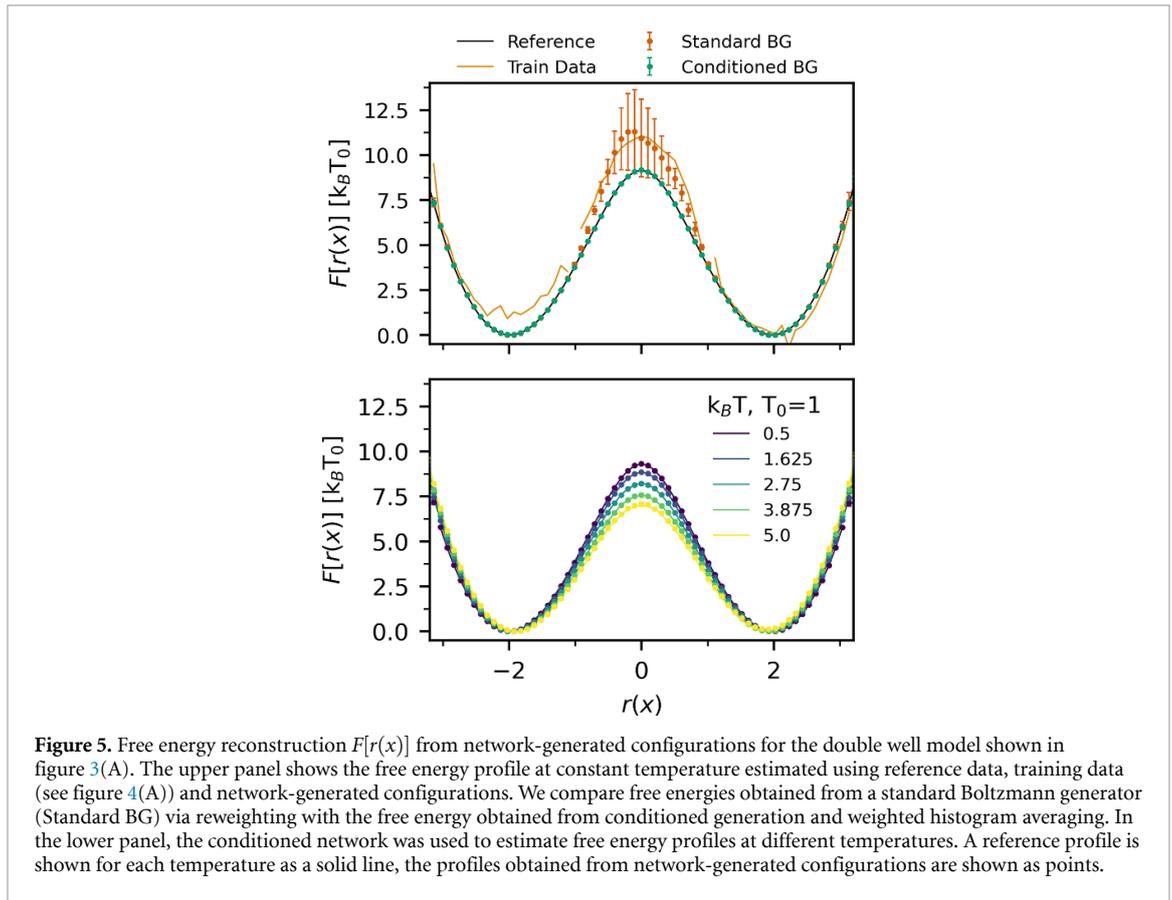


From network-generated configurations at different bias centers, accurate free energy profiles can be reconstructed. Using the weighted histogram analysis method [41], the free energy as a function of the reaction coordinate can be obtained. While the generator is trained using configurations from discrete windows, the network architecture and the process of training by energy allow the sampling at arbitrary bias centers (figure 4). For this reason, an accurate free energy estimate can be obtained by increasing samples or increasing the window count even if the training data alone is not sufficient for the free energy reconstruction, as shown in figure 5. In addition, the ability to train the conditioned Boltzmann generator at different temperatures allows for the reconstruction of free energies at different values of $k_B T$, see figure 5. A targeted sampling is not possible for the standard Boltzmann generator, for which the reconstructed free energy lacks sufficient samples at the barrier top given the same total number of generated configurations.

2.2. Exploring path space

The ability to sample independent configurations in targeted regions of phase space opens up new possibilities to investigate rare events, as generated points can serve as shooting points for trajectories. Initial tests that employ this path sampling scheme are performed on a bistable double well model using the reaction coordinate $r(x) = x^{(0)}$. The potential energy function is constructed in a way such that two reaction channels separated by an energy barrier connect the two stable basins (figure 3(B), system parameters in SI).

With the trained network at hand, we compare three different path sampling methods: TPS using two-way shooting with randomized velocities (standard TPS) [13], TPS with a bias on the shooting point



selection [16] and path sampling from network-generated shooting points. Shooting range TPS is included since it is the closest Markov chain-based scheme to the proposed path generation from presampled, biased shooting points. In shooting range TPS, the selection probability p_{sel} of a shooting point on the previous path is biased via an arbitrary, user-defined function of a reaction coordinate. For the comparison to the network-based scheme, we use a Gaussian centered at the top of the barrier.

The results of these initial tests show that in contrast to sampling from generated shooting points, both standard TPS and shooting range TPS struggle to estimate the ratio between paths in the upper and lower

reaction channel correctly (figure 6(A)). To obtain a quantitative measure of this ratio, we define an indicator function $g(X)$ that describes whether the path follows the upper or lower reaction channel. The function returns unity if the path follows the upper channel and vanishes otherwise. With reference to figure 4(D), the function $g(X)$ is given by

$$g(X) = \begin{cases} 1 & \text{if } \bar{X}'(\tau) > 0; \\ 0 & \text{otherwise,} \end{cases}$$

where \bar{X}' is the average of the y component of the trajectory over the single transition path.

The expectation value of $g(X)$ over all reactive paths is $1/2$ since the dynamics and the state definitions are symmetric. For standard TPS and shooting range TPS, the average value of the indicator function oscillates around the expected value. Since correlations between sampled paths are unavoidable with shooting moves, subsequently sampled paths are likely to remain in the same reaction channel. A switch to the other channel is only observed infrequently as indicated by the integrated autocorrelation times. This leads to the oscillating behavior of the average indicator function. In comparison, the path sampling from generated shooting points produces independent paths in the upper and lower reaction channel, leading to fast convergence of the average value of the indicator function, as shown in figure 6(A).

To further compare the performance of the different path sampling schemes, a reference path ensemble is sampled by means of 250 000 trials using two-way shooting TPS with randomized velocities. The difference between the reference path ensemble and a path ensemble at trial n can then be obtained by comparing the discretized density of configurations on transition paths $p(x|TP)$ as in [16]. The neural-network based sampling scheme outperforms standard TPS and shooting range TPS when looking at this difference between the path ensemble at trial n and the rigorously sampled reference ensemble. Moreover, due to the reweighting of generated configurations and paths, a proper distribution of paths can be obtained even if the generated shooting points are biased towards one reaction channel.

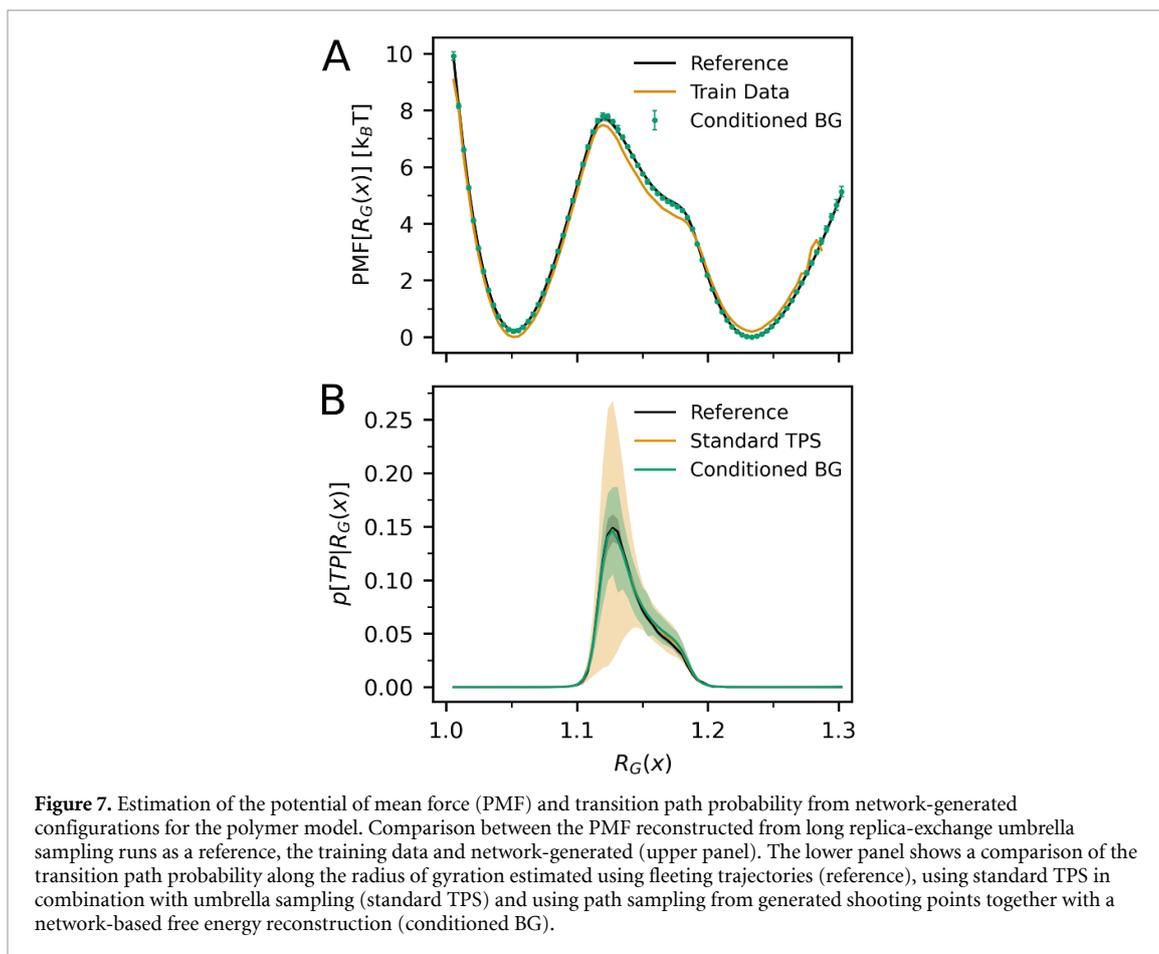
To test the scalability of the approach to higher dimensional systems, we consider a polymer model of $N = 7$ beads in two dimensions, as illustrated in figure 3(C). The interaction between beads includes a non-bonded Lennard–Jones interaction, a bond stretching term and an angular term (details of the potential and simulation parameters in SI). Two stable states can be identified in this system, an extended and a compact configuration (figure 3(C)). Since the states are solely identified by the radius of gyration and the Lennard–Jones interaction energy, all possible bonding permutations are included in the states. Not only does the transition between these states take place infrequently but it also occurs via different reaction channels, making it an ideal test system for rare event sampling.

Just like in the two-dimensional model case, we benchmark standard TPS, shooting range TPS and path sampling from network-generated shooting points on the transition from extended to circular states in the polymer model. As a coordinate for the shooting range bias and generation of initial points, we choose the radius of gyration R_G of the polymer. To compare the configuration density in the path ensemble at trial n with the reference ensemble for the different methods (figure 6(B)), we discretize the configuration space of the polymer (see SI for further details). In contrast to sampling from generated shooting points, some standard and shooting range TPS runs show slow or non-existent convergence towards the reference path ensemble. This effect may be explained by looking at the average transition path lengths or the autocorrelation function of the path length. Path sampling runs that do not converge to the reference ensemble come with an over- or underestimated average path length. Combined with the long correlation times of the path length, it can be concluded that standard TPS and shooting range TPS are prone to get stuck in a faster or slower reaction channel compared to the typical reaction channel. Since sampled paths from generated shooting points are uncorrelated, all reaction channels are visited independently in proportion to their statistical weight. This leads to a consistent convergence to the reference ensemble and an accurate estimate of the average path length.

2.3. Finding transition states

The conditioning of the Boltzmann generator also allows for a closer study of the bias coordinate. A free energy profile can be reconstructed from configurations at different bias centers as already shown for the double well. This also applies to the polymer model, where the Boltzmann generator learns the correct free energy profile along the reaction coordinate even with improperly weighted training data (figure 7(A)).

While the position of the barrier top can give initial information on the position of possible transition states, the central quantity of interest is the transition path probability $p(TP|r(x))$, i.e. the conditional probability for a point x to be on a transition path given the value of the reaction coordinate $r(x)$. Usually one obtains this measure by producing multiple fleeting trajectories from configurations with the specific reaction coordinate value [34]. An immediate drawback of this approach is that fleeting trajectories need to



be produced separately from the path sampling run, making the calculations expensive. Alternatively, one could estimate $p(\text{TP}|r(x))$ up to a proportionality constant using Bayes' theorem [34]:

$$p(\text{TP}|r(x)) \propto \frac{p(r(x)|\text{TP})}{p_{\text{eq}}(r(x))}. \quad (23)$$

However, since it is not trivial to extract information on the equilibrium distribution $p_{\text{eq}}(r(x))$ from the transition path ensemble, the calculation of $p(\text{TP}|r(x))$ using the Bayesian approach requires additional simulations and is therefore usually less efficient than estimation via fleeting trajectories. Moreover, both distributions $p(r(x)|\text{TP})$ and $p_{\text{eq}}(r(x))$ come with an uncertainty when estimated from simulation data and this uncertainty propagates to the estimated transition path probability.

The path sampling from network-generated shooting points proposed in this work allows for both the accurate reconstruction of the free energy and for the calculation of the path distribution along the reaction coordinate. To demonstrate this, we compare the calculation of $p(\text{TP}|R_G(x))$ for the polymer model using the Bayesian approach (figure 7(B)) with the results obtained using the standard approach employing fleeting trajectories. In the following comparison, we compute error estimates by performing simulations in replicas and by using Gaussian error propagation. We first estimate $p(\text{TP}|R_G(x))$ from 10 independent runs each with 25 000 fleeting trajectories as a reference. As a second baseline, we use the reference path ensemble (same as in figure 6(B)) and the reference free energy (same as in figure 7) to estimate the transition path probability. Here the resulting error does not allow for accurate determination of transition state regions on the reaction coordinate. In comparison, using the data from the Boltzmann generator (figure 7(A) and the right column in figure 6(B)) leads to a more accurate estimate of the transition path probability. The advantage of this approach is that the calculation is inexpensive since the trained network enables fast estimation of free energy profiles and efficient TPS.

3. Discussion and conclusion

In this work, we introduced the conditioning of Boltzmann generators for enhanced sampling of low probability regions in configuration space. The conditioned generators can be used, in the first place, to obtain more accurate free energy profiles. Secondly, we proposed a path sampling scheme based on a set of

presampled, network-generated shooting points. While Boltzmann generators were a natural choice for this proof of concept, the sampling scheme, including the path reweighting factor derived in equation (9), can be generalized to any exact-likelihood generative model that allows for computing the density of a generated sample. Recently proposed stochastic normalizing flows [42], smooth flows [43] or equivariant normalizing flows [21] can easily be adopted depending on the system to study.

From our experiments we see that the sampling acceleration and the computational cost of the workflow proposed here are both highly dependent on the physical system and on the observables of interest. The overall efficiency of the approach is negatively impacted by complex and multi-modal systems since they require larger network sizes and more configurations in the training set. On the other hand, long transition times or strong correlations between paths in TPS-based methods can quickly compensate the cost of additional energy evaluations required for training a flow. This is due to the numerical integration of the equations of motion requiring a full force evaluation per timestep. Finally, the accuracy required for the estimation of configuration or path observables is also a crucial factor, since training of a flow-based model is a one-time investment and only breaks even for longer sampling runs.

It is important to note that, for both use cases—free energy reconstruction and path sampling—we based our algorithm on a reaction coordinate $r(x)$. In the systems discussed, this coordinate is either trivial to find or could be obtained by educated guessing. Only after the training of the network and the whole path sampling process, it is possible to obtain a measure of the quality of the chosen reaction coordinate, e.g. by estimating the probability to generate a transition path. This approach is not straightaway transferable to more complex systems as reaction coordinates often turn into a less intuitive combination of order parameters [44, 45]. A direct approach to tackle this problem may be to use existing algorithms for reaction coordinate optimization such as the algorithm proposed by Peters and Trout [14, 46] or to use a reinforcement learning scheme as proposed by Jung *et al* [47]. Also, even though a reaction coordinate is often challenging to find, a reasonable order parameter may sometimes be more apparent. Here the difference is that an order parameter may distinguish between different states of the system but does not necessarily have a defined, compact region linked to a high probability to generate a transition path. Therefore, as an alternative to prior reaction coordinate analysis, the functional form of the bias potential could be adapted. Instead of a harmonic bias centered on a specific region on the coordinate, one could realize the biasing via a history-dependent bias potential as employed in metadynamics [48]. With this approach, low probability states along the whole reaction coordinate could be sampled eliminating the need for centering the shooting points on a specific region.

For the future, we see the approach of using generative neural networks for rare event sampling as especially useful if prior knowledge of a reaction coordinate exists and multiple orthogonal reaction channels complicate the sampling.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.5281/zenodo.7418398>.

Acknowledgments

We acknowledge the financial support of the Austrian Science Fund (FWF) through the SFB TACO, Grant Number F 81-N. The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC).

ORCID iDs

Sebastian Falkner  <https://orcid.org/0000-0003-2640-1761>

Alessandro Coretti  <https://orcid.org/0000-0002-7131-3210>

Salvatore Romano  <https://orcid.org/0000-0002-9762-9125>

Phillip L Geissler  <https://orcid.org/0000-0003-0268-6547>

Christoph Dellago  <https://orcid.org/0000-0001-9166-6235>

References

- [1] Menzl G and Dellago C 2016 Effect of entropy on the nucleation of cavitation bubbles in water under tension *J. Chem. Phys.* **145** 211918
- [2] Arjun, Berendsen T A and Bolhuis P G 2019 Unbiased atomistic insight in the competing nucleation mechanisms of methane hydrates *Proc. Natl Acad. Sci.* **116** 19305–10

- [3] Juraszek J and Bolhuis P G 2006 Sampling the multiple folding mechanisms of Trp-cage in explicit solvent *Proc. Natl Acad. Sci.* **103** 15859–64
- [4] Okazaki K-I, Wöhlert D, Warnau J, Jung H, Yildiz O, Kühlbrandt W and Hummer G 2019 Mechanism of the electroneutral sodium/proton antiporter PaNhaP from transition-path shooting *Nat. Commun.* **10** 1742
- [5] Geissler P L, Dellago C and Chandler D 1999 Kinetic pathways of ion pair dissociation in water *J. Phys. Chem. B* **103** 3706–10
- [6] Ballard A J and Dellago C 2012 Toward the mechanism of ionic dissociation in water *J. Phys. Chem. B* **116** 13490–7
- [7] Schwierz N 2020 Kinetic pathways of water exchange in the first hydration shell of magnesium *J. Chem. Phys.* **152** 224106
- [8] Falkner S and Schwierz N 2021 Kinetic pathways of water exchange in the first hydration shell of magnesium: influence of water model and ionic force field *J. Chem. Phys.* **155** 084503
- [9] Geissler P L, Dellago C, Chandler D, Hutter J and Parrinello M 2001 Autoionization in liquid water *Science* **291** 2121–4
- [10] Leitold C, Mundy C J, Baer M D, Schenter G K and Peters B 2020 Solvent reaction coordinate for an $s_n 2$ reaction *J. Chem. Phys.* **153** 024103
- [11] Torrie G M and Valleau J P 1977 Nonphysical sampling distributions in Monte Carlo free-energy estimation: umbrella sampling *J. Comput. Phys.* **23** 187–99
- [12] Bolhuis P G, Chandler D, Dellago C and Geissler P L 2002 Transition path sampling: throwing ropes over rough mountain passes, in the dark *Annu. Rev. Phys. Chem.* **53** 291–318
- [13] Dellago C, Bolhuis P G, Csajka F S and Chandler D 1998 Transition path sampling and the calculation of rate constants *J. Chem. Phys.* **108** 1964–77
- [14] Peters B and Trout B L 2006 Obtaining reaction coordinates by likelihood maximization *J. Chem. Phys.* **125** 054108
- [15] Faidon Brotzakis Z and Bolhuis P G 2016 A one-way shooting algorithm for transition path sampling of asymmetric barriers *J. Chem. Phys.* **145** 164112
- [16] Jung H, Okazaki K-I and Hummer G 2017 Transition path sampling of rare events by shooting from the top *J. Chem. Phys.* **147** 152716
- [17] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial nets *Advances in Neural Information Processing Systems* vol 27, ed Z Ghahramani, M Welling, C Cortes, N Lawrence and K Q Weinberger (Curran Associates, Inc.)
- [18] Kingma D P and Welling M 2014 Auto-encoding variational Bayes *2nd Int. Conf. on Learning Representations (ICLR 2014) (Banff, AB, Canada, 14–16 April 2014) (Conf. Track Proc.)*
- [19] Papamakarios G, Nalisnick E, Rezende D J, Mohamed S and Lakshminarayanan B 2019 Normalizing flows for probabilistic modeling and inference *J. Mach. Learn. Res.* **22** 1–64
- [20] Rezende D and Mohamed S 2015 Variational inference with normalizing flows *Int. Conf. on Machine Learning (PMLR)* pp 1530–8
- [21] Wirnsberger P, Ballard A J, Papamakarios G, Abercrombie S, Racanière S, Pritzel A, Rezende D J and Blundell C 2020 Targeted free energy estimation via learned mappings *J. Chem. Phys.* **153** 144112
- [22] Noé F, Olsson S, Köhler J and Wu H 2019 Boltzmann generators: sampling equilibrium states of many-body systems with deep learning *Science* **365** eaaw1147
- [23] Wirnsberger P, Papamakarios G, Ibarz B, Racanière S, Ballard A J, Pritzel A and Blundell C 2022 Normalizing flows for atomic solids *Mach. Learn.: Sci. Technol.* **3** 025009
- [24] Liu T, Gao W, Wang Z and Wang C 2022 Pathflow: a normalizing flow generator that finds transition paths *38th Conf. on Uncertainty in Artificial Intelligence*
- [25] Köhler J, Chen Y, Krämer A, Clementi C and Noé F 2023 Flow-matching: efficient coarse-graining of molecular dynamics without forces *J. Chem. Theory Comput.* **19** 942–52
- [26] Albergo M S, Kanwar G and Shanahan P E 2019 Flow-based generative models for Markov chain Monte Carlo in lattice field theory *Phys. Rev. D* **100** 034515
- [27] Nicoli K A, Anders C J, Funcke L, Hartung T, Jansen K, Kessel P, Nakajima S and Stornati P 2021 Estimation of thermodynamic observables in lattice field theories with deep generative models *Phys. Rev. Lett.* **126** 032001
- [28] Lawrence S and Yamauchi Y 2021 Normalizing flows and the real-time sign problem *Phys. Rev. D* **103** 114509
- [29] Abbott R et al 2022 Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions *Phys. Rev. D* **106** 074506
- [30] Nicoli K A, Nakajima S, Strodthoff N, Samek W, Müller K-R and Kessel P 2020 Asymptotically unbiased estimation of physical observables with neural samplers *Phys. Rev. E* **101** 023304
- [31] Gerdes M, de Haan P, Rainone C, Bondesan R and Cheng M C N 2022 Learning lattice quantum field theories with equivariant continuous flows (arXiv:2207.00283)
- [32] Singha A, Chakrabarti D and Arora V 2023 Conditional normalizing flow for Markov chain Monte Carlo sampling in the critical region of lattice field theory *Phys. Rev. D* **107** 014512
- [33] Singha A, Chakrabarti D and Arora V 2023 Sampling U(1) gauge theory using a re-trainable conditional flow-based model (arXiv:2306.00581)
- [34] Hummer G 2004 From transition paths to transition states and rate coefficients *J. Chem. Phys.* **120** 516–23
- [35] Daru J and Stirling A 2014 Divided saddle theory: a new idea for rate constant calculation *J. Chem. Theory Comput.* **10** 1121–7
- [36] Menzl G, Singraber A and Dellago C 2016 S-shooting: a Bennett-Chandler-like method for the computation of rate constants from committor trajectories *Faraday Discuss.* **195** 345–64
- [37] Best R B and Hummer G 2005 Reaction coordinates and rates from transition paths *Proc. Natl Acad. Sci.* **102** 6732–7
- [38] Dinh L, Sohl-Dickstein J and Bengio S 2016 Density estimation using real NVP *5th Int. Conf. on Learning Representations (ICLR 2017) (Conf. Track Proc.)*
- [39] Ardizzone L, Lüth C, Kruse J, Rother C and Köthe U 2019 Guided image generation with conditional invertible neural networks (arXiv:1907.02392)
- [40] Nicoli K A, Anders C J, Hartung T, Jansen K, Kessel P and Nakajima S 2023 Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories (arXiv:2302.14082)
- [41] Ferrenberg A M and Swendsen R H 1989 Optimized Monte Carlo data analysis *Phys. Rev. Lett.* **63** 1195–8
- [42] Wu H, Köhler J and Noe F 2020 Stochastic normalizing flows *Advances in Neural Information Processing Systems* vol 33, ed H Larochelle, M Ranzato, R Hadsell, M F Balcan and H Lin (Curran Associates, Inc.) pp 5933–44
- [43] Köhler J, Krämer A and Noe F 2021 Smooth normalizing flows *Advances in Neural Information Processing Systems* vol 34, ed M Ranzato, A Beygelzimer, Y Dauphin, P S Liang and J W Vaughan (Curran Associates, Inc.) pp 2796–809

- [44] Juraszek J and Bolhuis P G 2008 Rate constant and reaction coordinate of Trp-cage folding in explicit water *Biophys. J.* **95** 4246–57
- [45] Bolhuis P G, Dellago C and Chandler D 2000 Reaction coordinates of biomolecular isomerization *Proc. Natl Acad. Sci.* **97** 5877–82
- [46] Peters B 2016 Reaction coordinates and mechanistic hypothesis tests *Annu. Rev. Phys. Chem.* **67** 669–90
- [47] Jung H, Covino R, Arjun A, Leitold C, Dellago C, Bolhuis P G and Hummer G 2023 Machine-guided path sampling to discover mechanisms of molecular self-organization *Nat. Comput. Sci.* **3** 334–45
- [48] Laio A and Parrinello M 2002 Escaping free-energy minima *Proc. Natl Acad. Sci.* **99** 12562–6