Deep Learning for Multi-Scale Molecular Modeling

Jun Zhang[1,2], Yao-Kun Lei[3], Yi Isaac Yang[1,†] and Yi Qin Gao[1,3,4,5,‡]

[1] *Institute of Systems Biology, Shenzhen Bay Laboratory, 518055 Shenzhen, China*

[2] *Department of Mathematics and Computer Science, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany*

[3] *Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, 100871 Beijing, China.*

[4] *Beijing Advanced Innovation Center for Genomics, Peking University, 100871 Beijing, China.*

[5] *Biomedical Pioneering Innovation Center, Peking University, 100871 Beijing, China.*

Corresponding authors: † (Y.I.Y) yangyi@szbl.ac.cn or ‡ (Y.Q.G) gaoyq@pku.edu.cn

Abstract

Molecular simulations are widely applied in the study of chemical and bio-physical systems. However, the accessible timescales of atomistic simulations are limited, and extracting equilibrium properties of systems containing rare events remains challenging. Two distinct strategies are usually adopted in this regard: either sticking to the atomistic level and performing enhanced sampling, or trading details for speed by leveraging coarse-grained models. Although both strategies are promising, either of them, if adopted individually, exhibits severe limitations. In this paper we propose a machine-learning approach to ally both strategies so that simulations on different scales can benefit mutually from their cross-talks: Accurate coarse-grained (CG) models can be inferred from the fine-grained (FG) simulations through deep generative learning; In turn, FG simulations can be boosted by the guidance of CG models via deep reinforcement learning. Our method defines a variational and adaptive training objective which allows end-to-end training of parametric molecular models using deep neural networks. Through multiple experiments, we show that our method is efficient and flexible, and performs well on challenging chemical and bio-molecular systems.

Keywords: Molecular dynamics, Artificial intelligence, Coarse graining, Enhanced sampling

I. Introduction

Molecular simulations, particularly all-atom and *ab initio* molecular dynamics (MD), have furthered our understanding of many chemical and bio-physical processes.[1-2] In molecular simulations, interactions between particles (e.g., atoms, residues or molecules) are described by an (potential) energy function $U(\mathbf{R})$ of the configuration $\mathbf{R}$. To investigate such systems, one is often not interested in the exact energy, but the free energy, or the equilibrium distribution, of some reduced descriptors, e.g., collective variables (CV)[3] or CG variables,[4] $\mathbf{s}(\mathbf{R})$, as a function of $\mathbf{R}$:

$$p(\mathbf{s}) = \frac{\int e^{-\beta U(\mathbf{R})}\delta\big(\mathbf{s}-\mathbf{s}(\mathbf{R})\big)d\mathbf{R}}{\int e^{-\beta U(\mathbf{R})}d\mathbf{R}} = \frac{e^{-\beta F(\mathbf{s})}}{Z} \tag{1}$$

$$F(\mathbf{s}) = -\frac{1}{\beta}[\log p(\mathbf{s}) + \log Z] \tag{2}$$

where $Z = \int e^{-\beta U(\mathbf{R})}d\mathbf{R}$ is the partition function, β is the inverse temperature and δ denotes the Dirac-delta function. Equation (2) holds up to an arbitrary additive constant. Usually $\mathbf{s}$ is selected to be slowly changing variables governing the process of interest, and the rest of degrees of freedom (DOF) can be treated in the mean-field fashion.[5-6] Under this setting, $F(\mathbf{s})$ becomes a CG description of the original thermodynamic system, and simulations performed under $F(\mathbf{s})$ are generally much faster than those run on the FG potential (i.e., $U(\mathbf{R})$) because $\mathrm{Dim}(\mathbf{s}) \ll \mathrm{Dim}(\mathbf{R})$ (where $\mathrm{Dim}(\cdot)$ denotes the dimensionality) despite the loss of finer details. Therefore, Eqs. (1-2) are also known as the principle of thermodynamic consistency for coarse graining.[5] However, practical implementation of Eqs. (1-2) is hindered by two major issues: (1) How can one approximate a reliable analytical form for the CG potential $F(\mathbf{s})$ given access to samples drawn from $U(\mathbf{R})$? (2) How can one draw equilibrium samples from $U(\mathbf{R})$, which are further used to infer $F(\mathbf{s})$? The former is known as the coarse graining problem, while the latter as the importance sampling problem, and both are of particular importance in physics, chemistry and biology.

Conventionally, if $\mathbf{s}$ is low-dimensional (say, $\mathrm{Dim}(\mathbf{s}) \leq 3$), non-parametric methods like kernel density estimation (KDE)[7] can be adopted to infer $F(\mathbf{s})$, but they become quickly infeasible as $\mathrm{Dim}(\mathbf{s})$ increases. Artificial neural networks (ANNs) and deep learning may offer extra flexibility and expressivity to this end.[8-9] For instance, several recent studies proposed supervised learning approaches to fit $F(\mathbf{s})$ by ANNs.[10-12] However, fitting $F(\mathbf{s})$ as a regression problem has several drawbacks: (1) it necessitates gridding the space of $\mathbf{s}$, which would be computationally prohibitive for large $\mathrm{Dim}(\mathbf{s})$; (2) it is rather data-inefficient because calculating $F(\mathbf{s})$ at one point needs a large amount of samples from $U(\mathbf{R})$ at the neighborhood of $\mathbf{s}$ (i.e., $\delta\big(\mathbf{s}-\mathbf{s}(\mathbf{R})\big)$); (3) Since regression learning lays equal importance over any point in the probability measure space regardless of the mass density distribution, it is not suitable for fitting an imbalanced distribution where the accuracy of areas with higher mass density should be more emphasized.

An alternative view of inferring $p(\mathbf{s})$ is provided by statistical modeling and machine learning community, where density estimation of high-dimensional data has been a long-standing goal.[9, 13] Particularly, a recent burst of work sparks new ideas to exploit deep learning to this end, giving rise to various deep generative models including variational auto-encoders,[14] auto-regressive[15] and normalizing-flow models.[16-17] These generative models enjoy the merits of fast sampling and scale well to data of large dimensionality. The price paid for such easy-sampling is that these models have to be footed on certain simple prior distributions, hence the complexity of the distributions resulted from these methods is bound by the manifold structure of the prior distribution. Therefore, they are known to suffer from issues like mode-dropping and often assigning probability mass to

areas unwarranted by the data.[18] On the other hand, generative learning with energy-based models (EBMs) can be dated back to even longer before,[19-21] which in principle can fit arbitrarily complex distributions due to the flexibility and plasticity of the energy landscapes.[22-23] Due to the flexibility, preserving the symmetry and invariance of many-particle systems which is difficult for other generative models remains tractable in EBMs. Moreover, EMBs can be easily conditioned on *a priori* restraints thanks to the additive compositionality of the energy functions. Despite of these advantages, implementation of EBMs is often hindered by the intractable sampling issue which are circumvented by above-mentioned deep generative models.

In this paper, combining the strengths of deep generative models and EBMs, we propose a variational approach to the CG problem, i.e., to infer an analytic form for the complex free energy surface (FES) without supervision in a tractable manner. Furthermore, our new approach allows simulations on different scales, which are launched simultaneously, cross talk and benefit from each other so that the inferred CG potential can in turn help enhance the sampling of the FG model in a reinforced and adaptive manner.

II. Methods

1. Variational adversarial density estimation (VADE).

We propose a deep learning approach to approximate the (possibly high-dimensional) FES, $F(\mathbf{s})$, by parametric models. Specifically, we denote the approximate free energy function and the associated probability distribution as $F_\theta(\mathbf{s})$ and $p_\theta(\mathbf{s}) \propto \exp\big(-\beta F_\theta(\mathbf{s})\big)$ respectively (where θ are optimizable parameters), and define a strict divergence $D(p||p_\theta)$ between p and p_θ. A strict divergence, including but not limited to Kullback-Leibler divergence $D_{\mathrm{KL}}(p||p_\theta)$[24] and Earth-Mover's distance $D_{\mathrm{W}}(p||p_\theta)$,[25] satisfies the condition that $D(p||p_\theta) \geq 0$ where the equality holds if and only if $p = p_\theta$, hence can be used as a variational objective.[26-28] Particularly, the gradient of $D_{\mathrm{KL}}(p||p_\theta)$ w.r.t. θ takes the following form in Eq. (3) (see SI or references[27] for the derivation),

$$\nabla_\theta D_{\mathrm{KL}}(p||p_\theta) = \langle \beta \nabla_\theta F_\theta(\mathbf{s}) \rangle_{p(\mathbf{s})} - \langle \beta \nabla_\theta F_\theta(\mathbf{s}) \rangle_{p_\theta(\mathbf{s})} \tag{3}$$

where $\langle f(x) \rangle_{p(x)}$ denotes the expectation value of f, a function of x, over a distribution $p(x)$. The gradient for $D_{\mathrm{W}}(p||p_\theta)$ has been separately derived in Targeted Adversarial Learning Optimized Sampling (TALOS)[29]. In this paper we perform experiments exclusively according to Eq. (3) and leave D_{W} for future research. Noteworthy, Eq. (3) bears intriguing similarity with Wasserstein Generative Adversarial Networks (W-GANs).[25] As in GANs, optimizing Eq. (3) requires both positive samples (drawn from the ground-true distribution $p(\mathbf{s})$) and negative samples (drawn from the model potential $p_\theta(\mathbf{s})$), thus effectively avoids overfitting in an adversarial manner. Given this observation, we refer Eq. (3) as variational adversarial density estimation (VADE), and F_θ optimized via Eq. (3) as VADE potential hereafter (Fig. 1A).

We note here that VADE can be viewed as a more general extension of the inversed Monte Carlo method.[30] The merit of VADE over GANs lies in the fact that F_θ actually yields a density estimation of $\mathbf{s}$ whereas GANs cannot. Compared to GANs, VADE is able to preserve the necessary *a priori* physics restraints of the data by choosing proper energy functional forms. Moreover, the mode-dropping issue suffered by GANs could be avoided through active mode-exploration (Eq. (S11); see SI for more details). VADE requires the evaluation of the two expectation values in Eq. (3). If we have access to, say, equilibrium MD samples drawn from $U(\mathbf{R})$, the distribution of which is denoted as $p_{\mathrm{FG}}(\mathbf{s})$, we can use p_{FG} to approximate p and optimize F_θ w.r.t. a surrogate objective $D_{\mathrm{KL}}(p_{\mathrm{FG}}||p_\theta)$, minimizing which is equivalent to maximum likelihood estimation. The remaining task is to calculate $\langle \nabla_\theta F_\theta(\mathbf{s}) \rangle_{p_\theta(\mathbf{s})}$. Since the analytical forms of $F_\theta(\mathbf{s})$ and $\nabla_{\mathbf{s}} F_\theta(\mathbf{s})$ are both available, we can perform CG simulations on $F_\theta(\mathbf{s})$ to generate negative samples and estimate $\langle \nabla_\theta F_\theta(\mathbf{s}) \rangle_{p_\theta(\mathbf{s})}$. If $\mathbf{s}$ is relatively low-dimensional, the CG simulations are usually computationally economic and converge relatively fast, hence brute-force simulations via Monte Carlo (MC) sampling or Langevin dynamics (LD) generally suffice.

2. Scale up VADE via Neural Samplers (VADE-NS).

In contrast to GANs where negative samples can be easily generated by drawing simple independent random noises,[31] VADE entails sampling over $F_\theta(\mathbf{s})$ through MC or LD simulations which could be intractable as $\mathrm{Dim}(\mathbf{s})$ grows. Worse still, after each update of θ, Eq. (3) requires sampling over the newly updated F_θ, thus optimizing F_θ through MC or LD sampling could be extremely time-consuming. To conquer this sampling issue, we propose to equip VADE with a neural sampler, i.e., a deep generative model, in replacement of the less efficient MC or LD simulations. Specifically, given a fixed-form energy function

F_θ, one can approximate an upper bound for the free energy (or equivalently, the normalized equilibrium distribution) via a surrogate distribution q_ψ by minimizing the KL-divergence between q_ψ and p_θ ,

$$D_{\mathrm{KL}}\left(q_\psi||p_\theta\right) = \langle \log q_\psi(\mathbf{s}) + \beta F_\theta(\mathbf{s})\rangle_{q_\psi(\mathbf{s})} + \log Z_\theta \tag{4}$$

The first term on the R.H.S. of Eq. (4) is known as the variational free energy of F_θ.[32-33] Note that calculating the constant Z_θ is unnecessary for solving this optimization (hence circumventing sampling over F_θ). The optimal $q_\psi(\mathbf{s})$ can be determined variationally according to Eq. (4) without the concern of overfitting, but the optimization requires that sampling from q_ψ is efficient and the log-likelihood $\log q_\psi(\mathbf{s})$ can be computed with ease.[33-34] Considering that a function f_ψ (parametrized by ψ) can map a random variable $\mathbf{z}$ *bijectively* into $\mathbf{s}$, i.e., $\mathbf{s} = f_\psi(\mathbf{z})$, and that $\mathbf{z}$ comes from a known distribution $q(\mathbf{z})$, we can thus compute $q_\psi(\mathbf{s})$ according to the change-of-variable formula (Eq. (5)),[17]

$$\log q_\psi(\mathbf{s}) = \log q(\mathbf{z}) - \log\left|\det\left(\frac{\partial f_\psi}{\partial \mathbf{z}}\right)\right| \tag{5}$$

where $\det(\partial f/\partial \mathbf{z})$ is the determinant of the Jacobian matrix $\partial f/\partial \mathbf{z}$. In practice, $q(\mathbf{z})$ can be chosen as a tractable distribution (e.g., independent multi-variate Gaussians) so that sampling over $\mathbf{z}$ as well as $q_\psi(\mathbf{s})$ can be convenient. Many deep bijective models[16, 35-38] can be adopted as f_ψ to approximate p_θ, and we term them as the neural sampler (NS) hereafter (see SI for more details about NS).

Equipping VADE with NS (VADE-NS) (Fig. 1B), we arrive at a nested optimization problem similar to GANs: Given a fixed F_θ, a NS is optimized according to Eq. (4); while given an optimized NS, F_θ can be updated according to Eq. (3). In this sense, the NS here takes a flavor of the generator in GANs whereas F_θ behaves like a discriminator. This nested optimization problem even allows f_ψ to be non-bijective mappings as shown in recent studies.[39] Noteworthy, in some applications of VADE where no *a priori* restraints are imposed over the functional form of F_θ, the nested optimization problem can be further simplified by modeling F_θ with a black-box deep bijective model (Eq. (6)),

$$\beta F_\theta(\mathbf{s}) \equiv -\log q_\psi(\mathbf{s}) \tag{6}$$

and optimizing F_θ directly through Eq. (3). More details about the training and model choices for VADE-NS can be found in SI.

3. Reinforced VADE (RE-VADE)

In a more common setting where p_{FG} is not available beforehand, we have to perform sampling over $U(\mathbf{R})$ from scratch. Since $\mathrm{Dim}(\mathbf{R})$ is usually very large, i.e., $\mathrm{Dim}(\mathbf{R}) \gg \mathrm{Dim}(\mathbf{s})$, estimating ensemble averages over $p(\mathbf{s})$ is often infeasible for brute-force FG (e.g., all-atom or *ab initio*) simulations or variational inference. A plethora of enhanced sampling methods have been developed trying to solve this issue, and there exist several excellent reviews on this topic.[3, 40-42] Here we will show that VADE, combined with reinforced imitation learning, provides a new solution to this problem. We name our new approach RE-VADE, which exhibits several compelling merits: (1) RE-VADE is able to handle high-dimensional CG variables while most of enhanced sampling methods become ineffective as $\mathrm{Dim}(\mathbf{s})$ grows; (2) RE-VADE casts a well-defined optimization problem which allows cross-fertilization with deep learning; (3) RE-VADE formulates an adaptive training objective which can be optimized variationally, hence ensuring the efficiency and convergence of the method.

Mathematically, if we have an optimal $F_\theta(\mathbf{s})$ which is equal to the ground-true $F(\mathbf{s})$, and perform sampling under $U - F_\theta$, we will arrive at a uniform distribution over $\mathbf{s}$. Therefore, one is motivated to employ $F_\theta(\mathbf{s})$ as a bias potential in order to achieve a flattened distribution over $\mathbf{s}$ which may originally involve high free energy barriers.[43] However, the situation is complicated by the errors in F_θ: Since F_θ is optimized w.r.t. available FG samples (usually corresponding to metastable states in MD simulations), its value can be very inaccurate in those under-sampled regions (e.g., transition regions). Therefore, directly inserting F_θ as the bias potential could be non- or even counter-productive. We note here that similar problems where one has to deal with moving distributions and partial sampling are commonly encountered and addressed in Reinforcement Learning (RL).[44] Inspired by RL, we introduce a two-timescale learning scheme, where a bias potential $V_\phi(\mathbf{s})$ with ϕ denoting optimizable parameters (equivalent to a policy function in RL) is separately trained in addition to $F_\theta(\mathbf{s})$ which can now be viewed as a value function in the spirit of actor-critic RL.[45] As in variationally enhanced sampling (VES)[28] or TALOS,[29] we can define a target distribution $p_\mathrm{T}(\mathbf{s};\theta)$ where the free energy barrier is lowered (in other words, less-visited regions are more encouraged) according to $F_\theta(\mathbf{s})$. Following the well-tempered (WT) metadynamics,[46] one reasonable choice of $p_\mathrm{T}(\mathbf{s};\theta)$ is Eq. (S23) (see SI for more information about p_T), and one can then optimize the bias potential $V_\phi(\mathbf{s})$ by minimizing a strict divergence, for instance, $D_\mathrm{KL}(p_\mathrm{T}||p_\phi)$,

$$\nabla_\phi D_\mathrm{KL}(p_\mathrm{T}||p_\phi) = \langle\beta\nabla_\phi V_\phi(\mathbf{s})\rangle_{p_\mathrm{T}} - \langle\beta\nabla_\phi V_\phi(\mathbf{s})\rangle_{p_\phi} \tag{7}$$

where p_ϕ denotes the Boltzmann distribution induced by $(U + V_\phi)$ which can be approximated through MD sampling. The separate parametrization allows us to use an imbalanced learning schedule for the free energy function F_θ and the bias potential V_ϕ. Particularly, we can train $F_\theta(\mathbf{s})$ based on the latest $p_\mathrm{FG}(\mathbf{s})$ (reweighted from p_ϕ) with a higher rate, and update $V_\phi(\mathbf{s})$ in a more conservative manner. In terms of imitation learning, $F_\theta(\mathbf{s})$ plays the role of a leader that coins a moving target based on the current density estimation, while $V_\phi(\mathbf{s})$ learns to tune the policy in order to catch up (Fig. 1C). A dramatic advantage of such leader-chaser scheme lies in the fact that F_θ along with p_T is constructed based on the samples drawn from simulations under V_ϕ, so p_T and p_ϕ always share substantial overlap; otherwise D_KL would fall victim to the notorious vanishing gradient issue.[25] Remarkably, such sort of separate parametrization and two-timescale updated rule were also adopted in some unsupervised and reinforcement learning settings like generative adversarial networks,[31, 47] double Q-learning[48] and parallel WaveNets,[49] *etc*.

Another major advantage of this separate parametrization scheme lies in the fact that density estimation of high-dimensional data inevitably involves high variance and large uncertainties, while RE-VADE still enables enhanced sampling with higher tolerance of these errors and even negates the necessity of knowing them. In RE-VADE, albeit $F_\theta(\mathbf{s})$ may be imperfect globally (especially at the beginning of training), it still gives relatively accurate density estimation of the recently visited regions, according to which $V_\phi(\mathbf{s})$ can be improved locally via Eq. (7) and gradually enhance the sampling efficiency. On the other hand, as the FG sampling is enhanced by $V_\phi(\mathbf{s})$, we can get a better approximation of $p(\mathbf{s})$, thus gradually push $F_\theta(\mathbf{s})$ to the optimum according to Eq. (3). Specifically, if one chooses a well-tempered target distribution as in Eq. (S13), the optimal V_ϕ will converge to the well-tempered free energy (Eq. (S14)). From this perspective, metadynamics [43] can be viewed as a special case of RE-VADE if $F_\theta(\mathbf{s})$ is replaced by KDE and accumulative Gaussians serve as $V_\phi(\mathbf{s})$. In other words, RE-VADE is a generalization of metadynamics to large $\mathrm{Dim}(\mathbf{s})$ and parametric bias potential functions. The assembled training algorithm of

RE-VADE is summarized in Algorithm S1. Compared to existing methods,[43, 50] VADE-NS allows us to implement Eq. (7) over high-dimensional CG space because density estimation and samples from the resulting target distribution (p_{T}) can be drawn conveniently. Particularly, in many enhanced sampling applications, *a priori* knowledge of the underlying $F(\mathbf{s})$ is scarcely available, hence we can model $F_\theta(\mathbf{s})$ with deep bijective models for simplicity (Eq. (6)). Besides, to attack the issue of irregular gradients while maximally harnessing the expressivity of ANNs, the architecture of ANNs should be carefully designed, and special regularization techniques may be needed to smooth the gradients (see SI for more details).

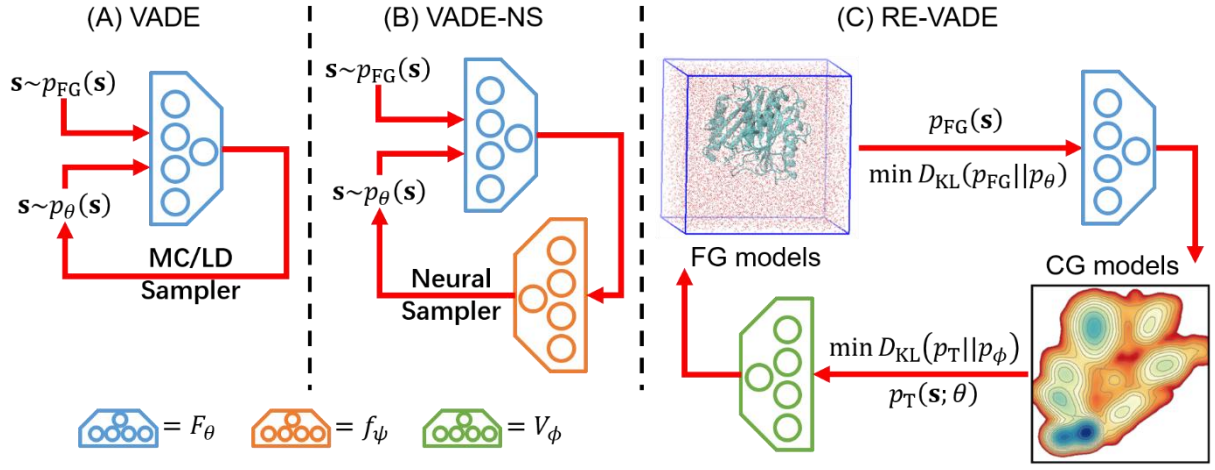

Figure 1. Workflow of variational adversarial density estimation (VADE). (A) VADE: the free energy approximator F_θ simultaneously takes in positive samples ($\mathbf{s}\sim p_{\mathrm{FG}}(\mathbf{s})$) and negative samples ($\mathbf{s}\sim p_\theta(\mathbf{s})$). The negative samples are generated by Monte Carlo or Langevin Dynamics simulations over F_θ. Red arrows indicate the direction of data flow or the computational graph. (B) VADE with a neural sampler (VADE-NS): identical to VADE except that the negative samples are drawn by a neural sampler f_ψ. (C) Reinforced VADE (RE-VADE): Given simulation samples from fine-grained (FG) models (p_{FG}), a coarse-grained (CG) potential F_θ can be approximated through VADE. In turn, a target distribution $p_{\mathrm{T}}(\mathbf{s};\theta)$ can be defined based on F_θ, according to which a bias potential V_ϕ can be variationally optimized to boost the FG simulation.

III. Results

I. Benchmark VADE as density estimator

We benchmarked VADE on the 2-dimensional Tiwary-Berne model[51] which consists of three potential wells. The contour map of the potential energy surface (PES), $U(x, y)$, is shown in Fig. 2A. A Langevin dynamics simulation was conducted on $U(x, y)$ as p_{FG} in Eq. (3) (see SI for more simulation details). We chose $\mathbf{s} = (x, y)$ in order to approximate the equilibrium distribution under $U(x, y)$ via VADE. We first performed kernel density estimation over the training data as a baseline (Fig. 2B). the FES obtained by KDE (Fig. 2B) agrees well with the ground-true PES (Fig. 2A) in all the metastable regions. However, the difference is relatively large in the transition regions due to the lack of training data in those regions which resulted from by insufficient sampling of LD simulations.

We then performed VADE and implemented importance MC sampling to optimize an ANN which serves as the density estimator, $F_\theta(x, y)$ (see SI for more details about model setup and training details). We trained $F_\theta(x, y)$ through VADE (Eq. (3) for 100 epochs, and plotted the estimated D_{KL} between p_θ and p_{FG} along the training progress (Fig. 2F). It can be seen that D_{KL} between the two distributions quickly diminished during training, and the optimization of F_θ around 100 epochs, demonstrating the efficiency and efficacy of VADE. Besides D_{KL}, another important indicator of the quality of models for VADE is whether or not F_θ of positive samples (i.e. training data) and negative samples (generated by VADE sampler) are identically (or similarly) distributed. Therefore, we examined the two distributions accordingly (Fig. 2E, panel 1), and found that they overlapped perfectly, proving the resulting F_θ is of high quality. The optimized $F_{\theta^*}(x, y)$ obtained by VADE was shown in Fig. 2C, and we found that F_{θ^*} not only agreeing well with the baseline KDE (Fig. 2B), but also appears quite smooth as the original PES thanks to the gradient regularizations we imposed over F_θ (see SI for more details). We also showcased the negative samples generated by MC sampling over F_{θ^*} (Fig. 2C), which indeed cannot be visually distinguished with those coming from the training set (Fig. 2B).

Furthermore, we chose a continuous normalizing-flow model with a free-form Jacobian[38] as the neural sampler, and performed VADE-NS over the same training set (see SI for more details about model setup and training details). For a fair comparison, VADE-NS was trained under the same settings with VADE. Tracking the optimization process, we found that D_{KL} between p_θ (which was yielded by VADE-NS) and p_{FG} also quickly diminished, and even dropped more rapidly than VADE in the early epochs showing higher initial training efficiency (Fig. 2F). However, as the training proceeded, VADE-NS did not reach the same minimum level in D_{KL} achieved by VADE, possibly due to insufficient training or the lack of expressivity commonly suffered by deep bijective models. Intriguingly, the distributions of F_θ (yielded by VADE-NS) over positive and negative samples also overlapped well (Fig. 2E, panel2), indicating that VADE-NS also found a minimum to the optimization problem in Eq. (3), but the solution is sub-optimal due to the insufficient variational flexibly of the bijective model. The optimized $F_{\theta^*}(x, y)$ obtained by VADE-NS was plotted in Fig. 2D. The F_{θ^*} obtained by VADE-NS was not so smooth as the model obtained by VADE, because it is non-trivial to impose the gradient regularizations over the bijective model in VADE-NS. Nonetheless, the overall contour of F_{θ^*} is still in line with the baseline KDE, capturing all the three metastable states while leaving blank for the transition regions, and the negative samples produced by F_{θ^*} also resemble those from the training set.

In summary, the surrogate energy function F_{θ^*} yielded by VADE and VADE-NS successfully preserves all the three local minima correctly, showing no signs of mode-dropping. This is a particular advantage of VADE provided that metastable states

usually play functionally important roles for bio-molecules, and dropping any of them during coarse graining may lead to failure of the resulting model. Now that we obtained the analytical form of F_θ based merely on samples from $U(x, y)$ (rather than knowing the mathematical form of $U(x, y)$), we can locate free energy minima and cluster samples. We minimized the simulation samples over F_{θ^*} obtained by VADE, and observed that all the samples finally fall into three distinct local minima (Fig. S1A). We also colored the simulation samples according to their final minimizers, as shown in Fig. S1B, and found that the noisily distributed samples were indeed assigned to different metastable states quite reasonably. This result implies important potential application of VADE in identifying free energy minima (or metastable states) and clustering noisy high-dimensional samples, which is demanded by many mechanism analysis and kinetic modeling methods. Besides, since F_θ deviates from the ground-true $U(x, y)$ due to the lack of samples over the transition region, F_θ would be further improved if more training samples over the transition regions can be obtained, and such consideration is a strong motivation behind RE-VADE.

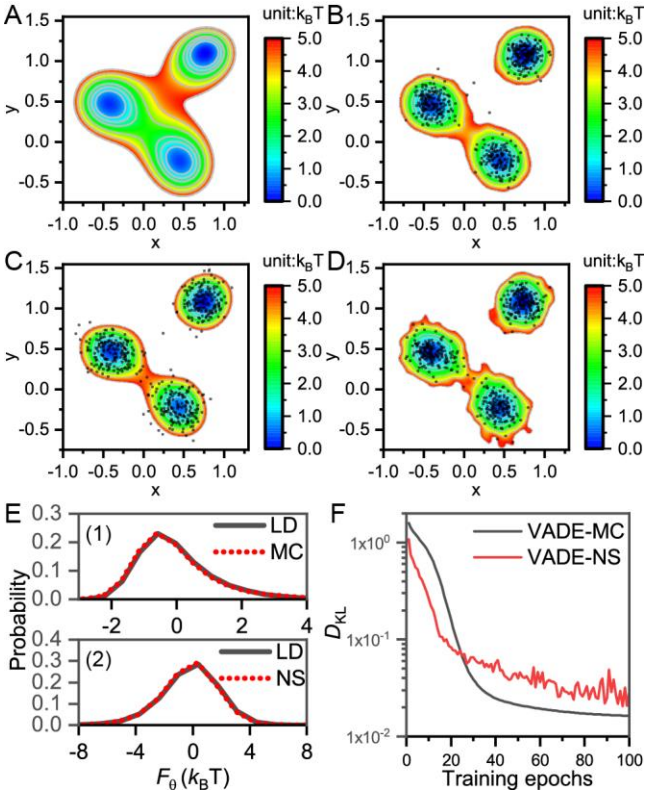

Figure 2. Benchmark VADE on a 2D numerical model. (A) 2D potential energy surface (PES), $U(x, y)$, of the model. The PES was shifted w.r.t. the global minimum so that $U(x, y) \geq 0$ everywhere. (B) Kernel density estimation of the training samples. The estimated free energy surface (FES) is shown in colored contour plot, superimposed with randomly selected training samples (black dots). The FES is shifted (and hereafter) as in (A) for a fair comparison. (C) $F_{\theta^*}(x, y)$ obtained by VADE which was optimized via Monte Carlo sampling (VADE-MC) is shown in colored contours. Black dots are random samples generated by MC over the optimized F_{θ^*}. (D) $F_{\theta^*}(x, y)$ obtained by VADE-NS is shown in colored contours, superimposed with random samples generated by the neural sampler. (E) Panel 1: The distribution of the final $F_{\theta^*}(x, y)$ obtained by VADE-MC over the training samples (black solid line) and those generated by MC (red dotted line); Panel 2: The distribution of the final $F_{\theta^*}(x, y)$ obtained by VADE-NS over training samples (black solid line) and those generated by NS (red dotted line). (F) Evolution of KL-divergence $D_{\mathrm{KL}}(p_{\mathrm{FG}}||p_\theta)$ against training epochs of VADE-MC (black) and VADE-NS (red).

2. Coarse grain a mini-protein by VADE

Particle-based CG models are widely adopted in bio-molecular modeling.[52] Conventionally, such models are premised on some empirical forms of force fields and fitted w.r.t. experimental and/or high-level calculation data. Since the force fields are generally designed to reflect certain physics restraints, the number of parameters are relatively limited and bottlenecks the expressivity and flexibility of the resulting models. In contrast, it is also possible to construct a CG potential entirely through an ANN in order to improve the expressivity, however, the model would require excessive amount of samples for training and might not generalize well due to the large parameter space. VADE can help combine the best of the two worlds if the CG potential takes a hybrid form (Eq. (8)):

$$F_\theta(\mathbf{s}) = F_\Theta(\mathbf{s}) + F_{\text{prior}}(\mathbf{s}) \tag{8}$$

where $F_\Theta(\mathbf{s})$ is a trainable parametric model while $F_{\text{prior}}(\mathbf{s})$ is a force-field-like term accounting for some *a priori* knowledge or physics restraints. Equation (8) can be interpreted from two perspectives: On the one hand, $F_\Theta(\mathbf{s})$ can be regarded as a correction term for the traditional force fields. On the other hand, $F_{\text{prior}}(\mathbf{s})$ serves as a reasonable prior distribution which effectively restricts the hypothesis space of F_θ thus expedites the training.

We tested this idea on a mini-protein Chignolin using long all-atom MD simulation trajectories contributed by Lindorff-Larsen *et al.*[53] The positions of the ten C_α atoms were selected as the CG variables **s** (Fig. 3A). $F_{\text{prior}}(\mathbf{s})$ contains repulsive restraints to penalize the overlapping between CG particles. We adopted a deep bijective model[36] as $F_\theta(\mathbf{s})$ and performed VADE-NS according to Eq. (S19), a regularized version of Eqs. (3-4), to incorporate the restraint F_{prior} (see SI for more details about model setups and training details). F_θ was trained over c.a. 500,000 FG samples for 10 epochs (one epoch takes less than 5 minutes in wall-clock time on a single NVIDIA GeForce GTX 1650 GPU card). After training finished, we plotted the distributions of F_θ over the fine-grained MD samples and the negative CG samples produced by VADE-NS (Fig. 3B). The two distributions agreed qualitatively well, indicating good quality of the resulting models.

In order to examine whether F_θ captures the important protein conformations, we characterized the generated CG samples from F_θ with two widely-used metrics: the root-mean-squared deviation (RMSD) w.r.t. the folded structure and the radius of gyration (Rg). We compared the distribution of the generated samples against the real ones in terms of these two metrics (Fig. 3C), and found that the CG potential faithfully reproduced the overall conformational features. Specifically, a sharp local minimum corresponds to the compact native conformation (small RMSD and small Rg) is well defined by F_θ, meanwhile the distribution of the extended unfolded structures is preserved as well.

In addition to the overall conformational patterns, we are also concerned with the fine details produced by the optimized CG model. Particularly, we examined the distributions of the C_α4-C_α5 bond length, C_α4-C_α5-C_α6 angle and C_α4-C_α5-C_α6-C_α7 torsion which locate at the turning point of the β-hairpin structure (Fig. 3A). We compared the distributions of these local structural features produced by F_θ against the reference FG MD samples (Fig. 3E), and confirmed that the optimized model $F_\theta(\mathbf{s})$ also faithfully reproduced these local structural patterns. Finally, we showcased some randomly sampled CG structures from F_θ, superimposed with their best-aligned FG counterparts (Fig. 3F). These CG conformations are diverse and realistic comparable to FG ones, including the folded structure, partially folded intermediates and extended unfolded structures,

demonstrating that the CG model obtained by VADE is free of mode-dropping and is able to capture the conformational features of the protein.

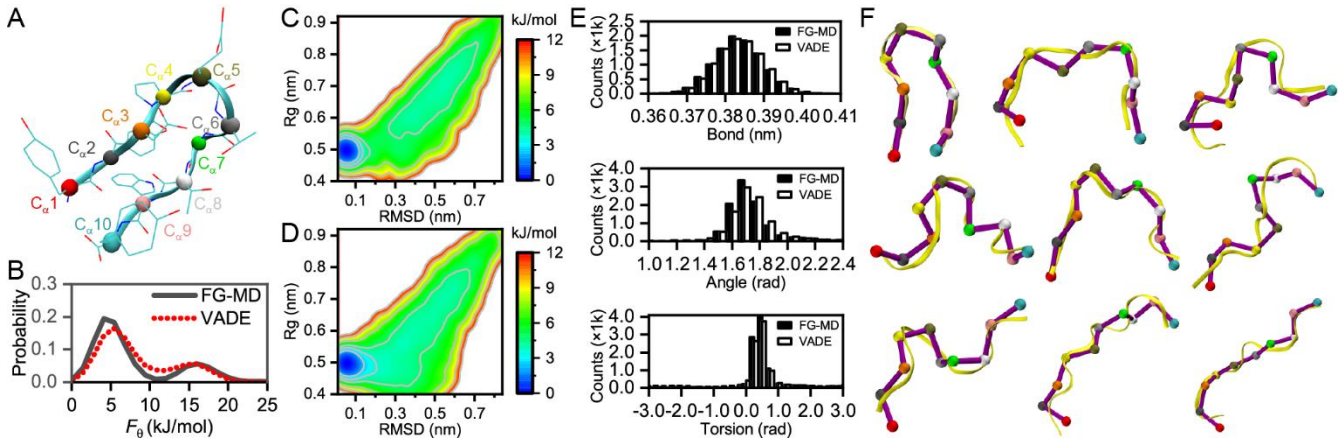

Figure 3. Coarse grain Chignolin by VADE. (A) FG protein structure of Chignolin. The CG particles are the 10 C_α atoms, which are highlighted by colored beads and indexed accordingly. (B) Distributions of VADE potential F_θ for FG MD simulation samples (black solid line) and for CG structures sampled from F_θ (red dotted line). (C) 2D FES spanned by RMSD and Rg of FG MD simulation samples (D) 2D FES spanned by RMSD and Rg of VADE CG structures. (E) Distributions of the bond length between C_α4-C_α5 (upper panel), the angle between C_α4-C_α5-C_α6 (middle) and the torsion between C_α4-C_α5-C_α6-C_α7 (bottom). Black histograms correspond to FG samples and white histograms to VADE samples. Statistics were performed over 10,000 FG MD samples and VADE CG samples, respectively. (F) VADE CG structures are shown beads (colored according to Fig. 3A) and purple bonds, superimposed with their best-aligned FG counterparts (shown in yellow ribbons) selected from the MD trajectories.

3. RE-VADE of alanine dipeptide in explicit water

Next, we proposed to construct a CG model for a prototypical bio-molecular system, alanine dipeptide (Ala2) in explicit water without available FG samples. Specifically, the backbone torsional angles were chosen to be the CG variables (Fig. 4A), that is, $\mathbf{s} = (\phi, \varphi)$, and our goal is to infer a reasonable FES, $F_\theta(\mathbf{s})$. However, different from the previous toy model, isomerization of (ϕ, φ) involves relatively high barrier, thus brute-force FG simulations of Ala2 converge too slowly to obtain an accurate estimate of $\langle \nabla_\theta F_\theta(\mathbf{s}) \rangle_{p_{\mathrm{FG}}}$. Therefore, we adopted RE-VADE to enhance the sampling over $\mathbf{s}$. Technically, we simultaneously launched two simulations: one FG (all-atom) MD simulation under a bias potential $V_\phi(\mathbf{s})$, and a CG MC simulation over $F_\theta(\mathbf{s})$. Both F_θ and V_ϕ are initialized to be zero everywhere. After a period of FG sampling (40 ps in length) biased by V_ϕ, we reweight the yielded FG samples to represent p_{FG}, and optimize F_θ according to Eq. (3). Based on the newly-trained F_θ, a well-tempered target distribution p_{T} according to Eq. (S23) (see SI for more details) is established, w.r.t. which the bias potential V_ϕ is optimized. V_ϕ is then fed into the FG simulations of Ala2 and yields an updated collection of samples representing p_{FG}. This procedure constitutes one iteration of RE-VADE, and the entire process continues till the convergence criteria are met (Algorithm S1).

We tracked how F_θ of FG simulation samples and CG ones distribute (Fig. 4B), provided that the similarity (or overlap) between these two distributions is a good indicator of convergence. We trained the FG and CG models by RE-VADE for 8 ns (or equivalently, 200 iterations), and found that the two distributions overlap well and that both spread for a relatively wide

range (implying no mode-dropping). Noteworthy, in such a short simulation length, it is impossible for brute-force MD to produce equilibrium samples covering all important metastable states. To illustrate how RE-VADE helps enhance the sampling of FG models, we presented the 1-ns trajectories for torsions ϕ and φ produced by vanilla MD in contrast to those produced by simulations biased by V_ϕ in Fig. 4C. It can be seen that isomerization of torsion ϕ is fairly frequent in biased MD but hardly found in vanilla MD. Similarly, rotation of torsion φ is also boosted significantly by V_ϕ. One may wonder how V_ϕ looks like and why it is able to boost the FG sampling so efficiently. By drawing the contour map of the bias potential (Fig. 4D), we observed that the optimized V_ϕ appears complementary to the ground-true FES of $\mathbf{s} = (\phi, \varphi)$ (a reference FES of $\mathbf{s}$ was presented in Fig. S2). We also superimposed some randomly selected FG samples produced under V_ϕ over the contour map (Fig. 4D), demonstrating an excellent coverage over both the free energy minima and the transition regions. Therefore, samples from V_ϕ are better representatives of p_{FG} and can be reliably used to optimize the CG models. The final CG model (F_θ) optimized via RE-VADE, which can be regarded as a variationally approximated FES for Ala2 in explicit water, is shown in Fig. 4E. We found that F_θ not only captures all known metastable states of Ala2 w.r.t. (ϕ, φ) (i.e., no mode-dropping), but also quantitatively agrees well with the reference FES (Fig S2). This example demonstrates that simulations on multiple scales can be bridged by RE-VADE and that CG models can be reliably inferred even without access to FG samples *a priori*.

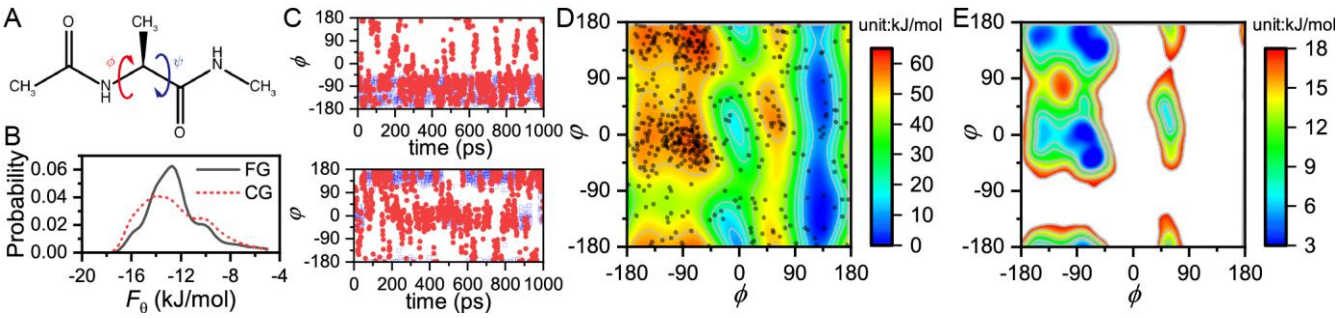

Figure 4. RE-VADE sampling of Ala2 in explicit water. (A) Chemical structure of Ala2 and two coarse-graining variables: the torsions ϕ and φ. (B) Distributions of F_θ for all-atom MD simulation samples (black solid line) and for CG simulation samples (red dotted line). (C) 1-ns simulation trajectories projected on torsion ϕ (upper panel) and torsion φ (lower panel). Blue squares correspond to vanilla MD, red dots to MD biased by V_ϕ. (D) The contour map of V_ϕ optimized via RE-VADE. Superimposed black dots are representative samples produced by the enhanced MD simulation under V_ϕ. (E) Contour map of F_θ optimized via RE-VADE.

4. RE-VADE of chemical reactions in explicit solvent

Chemical reactions in condensed phase are notoriously known to be difficult to simulate, and enhanced sampling approaches are often needed to this end. We thus employed RE-VADE to boost the simulation of a Claisen rearrangement reaction (Fig. 5A), which involves relatively high energy barrier, in the media of ionic liquid (see more details about simulation setup in SI). According to a previous study,[54] a linear combination of the breaking/forming bonds (i.e., $d1$ and $d2$ in Fig. 5A) is selected as the CV s, and the target distribution over s takes a Lorentzian form[55] (Eq. (S25)) which is developed to help enhance the sampling of the transition state regions (see more details about the CV, target distribution and training details in SI). The reactant is treated quantum mechanically (QM) while the solvent treated molecular mechanically (MM). One iteration of RE-VADE consists of 60-ps QM/MM simulation. As shown in Fig. 5B, after 100 iterations of RE-VADE, the chemical transitions

can take place back-and-forth within several nano-seconds by virtue of RE-VADE (i.e., the enhanced reaction rate is about 0.1 ns^{-1}). Noteworthy, as can be inferred from F_θ (Fig. 5C) which approximates the equilibrium potential of mean force over s (the putative reaction coordinate), the reaction rate is indeed very slow (less than 0.1 s^{-1} according to transition state theory). Therefore, within less than 10 ns simulation and training time, RE-VADE achieves an acceleration of the reaction rate by nearly 9 orders of magnitude.

To understand how RE-VADE achieves such impressive performance, we examined the target distribution p_T and the optimized bias potential $V_\phi(s)$. It can be seen from Fig. 5C that, based on the well-trained F_θ, the target distribution automatically emphasizes on the TS region while understates the reactant and product regions. In order to arrive at the target distribution, the bias potential V_ϕ is optimized in the way that the energy barrier is lowered and the metastable wells are lifted (Fig. 5C). By doing so, RE-VADE enables one to sample various molecular configurations efficiently (Fig. 5E), based on which we can investigate interesting thermodynamic properties of the reaction. For instance, the products of this reaction consists of a pair of enantiomers (Figs. 5A and 5E), the sampling of which is challenging for many sampling methods like umbrella sampling [6]. We constructed the 2-dimensional FES (Fig. 5D) spanned by the chirality order parameter q_C (the definition of q_C can be found in SI and references[56]) and s via KDE over samples generated by RE-VADE. We found that both enantiomers are adequately sampled. More importantly, the final FES appears fairly symmetric w.r.t. q_C, indicating that the distributions of the two enantiomers are almost identical, agreeing with the fact that the enantiomeric excess of this reaction should be zero. This result also demonstrates that the sampling of such a complex chemical reaction has converged.

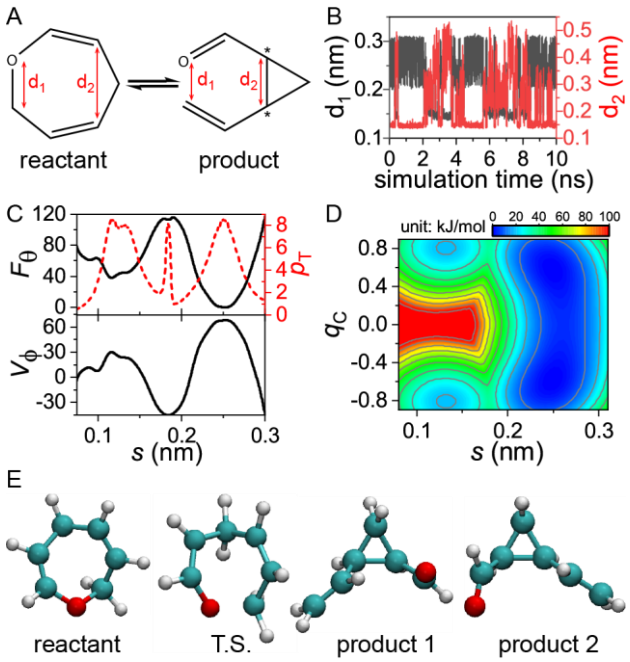

Figure 5. RE-VADE sampling of chemical reactions in ionic liquid. (A) The retro-Claisen rearrangement under study. The reactant is a 7-member-ring ether while the product is an aldehyde containing a 3-member ring. Two potential chiral carbons in the product are indicated by star symbols. The breaking and forming bonds are denoted as d_1 and d_2, respectively. (B) Trajectories of d_1 (black) and d_2 (red) during RE-VADE simulations. (C) Upper panel: the optimized CG potential $F_\theta(s)$ (black solid line) in unit of kJ/mol and the associated (unnormalized) target distribution $p_T(s)$ (red dashed line). Lower panel: the optimized bias potential V_θ in unit of kJ/mol. (D) 2D FES over

the CG variable s and chirality order parameter q_C. (E) Snapshots of various molecular configurations during RE-VADE. From left to right: reactant, transition state (T.S) and two enantiomer products.

IV. Concluding remarks

Despite that all-atom and *ab initio* MD simulations have assisted scientists gain insights over many important physical, chemical and biological processes, their applications to complex systems containing rare events are limited, because experimentally related timescales of such systems (like protein folding and chemical reactions) are well beyond the reachable scope for even the most powerful supercomputers. Two distinct strategies are separately developed to combat this issue: either to perform enhanced sampling over the atomistic model, or to leverage CG models at the cost of losing atomic details. However, traditional enhanced sampling methods can neither scale well to large system sizes, nor transfer well to different system types due to the requirement of system-specific expert knowledge. On the other end, although different attempts exist to build CG models incorporating atom-level knowledge, transitioning from atomistic models to CG models still remains challenging. In this paper we developed a machine-learning approach, (RE-)VADE, to connect FG and CG models. In (RE-)VADE, simulations on different scales can benefit from each other: CG models are optimized w.r.t. the FG simulations hence incorporating information on finer scales; In turn, FG simulations are enhanced under the guidance of CG models. Mathematically, (RE-)VADE belongs to the realm of unsupervised and reinforcement learning. The variational and self-adaptive training objective allows end-to-end and online training of parametric models like ANNs. Through several experiments we show that (RE-)VADE is able to yield flexible CG models more rapidly than traditional CG methods, moreover, it can also boost the sampling efficiency of chemical reactions and conformational transitions of prototypical biomolecules by several orders of magnitude. In (RE-)VADE, CG models can be variationally inferred based merely on equilibrium FG samples, thus involving less artifacts and computational cost than existing methods. This feature allows researchers to fully exploit the available atomistic simulations in order to construct transferrable CG models. More importantly, (RE-)VADE also allows one to construct CG models even without access to FG samples *a priori*. We remark here that although the CG models obtained this way are able to reproduce the thermodynamic properties of finer-scale simulations, the dynamics is not necessarily correct. In terms of Langevin dynamics, our CG models only provide a reliable description of the drifting field, yet the diffusion field remains unknown. Inferring the diffusion field based on the available drifting field can be an interesting direction for further studies. Besides, in the presented examples, RE-VADE is composed of one CG-FG cycle over merely two scales under the assumption that sampling over CG models can be readily achieved with simple simulation techniques and converges much faster than the FG simulations. While this is true for most cases because CG models are designed for computational tractability, some CG models for very large or complex systems may entail heavy computation and may also suffer from the sampling issue like FG models. If coarser-grained models were built on top of the CG model, they could possibly help enhance the CG sampling in a way similar to how the CG model boosts the FG sampling. Following this line, it is appealing to develop models over a cascade of scales via (RE-)VADE.

Acknowledgements

The authors thank Xing Che, Frank Noé, Lijiang Yang and Weinan E for useful discussion. This research was supported by National Natural Science Foundation of China [21927901, 21821004, 21873007 to Y.Q.G], the National Key Research and Development Program of China [2017YFA0204702 to Y.Q.G.] and Guangdong Basic and Applied Basic Research Foundation [2019A1515110278 to Y.I.Y.]. The author J.Z. thanks the Alexander von Humboldt Foundation for supporting part of this research.

Associated Content

Supplementary texts including introduction to derivation of (RE-)VADE training objectives, simulation setups and supplemental figures are provided in Supplemental Information.

Author contributions: J.Z., Y.K.L., Y.I.Y and Y.Q.G. designed the research; J.Z., Y.K.L. and Y.I.Y performed the research; J.Z., Y.K.L., Y.I.Y and Y.Q.G. analyzed the data; J.Z., Y.K.L., Y.I.Y and Y.Q.G. wrote the paper.

Notes: The authors declare no conflict of interest.

REFERENCES

1. Frenkel, D.; Smit, B., *Understanding molecular simulation: from algorithms to applications*. Elsevier: 2001; Vol. 1.
2. Tuckerman, M. E., Ab initio molecular dynamics: basic concepts, current trends and novel applications. *Journal of Physics: Condensed Matter* **2002,** *14* (50), R1297.
3. Fiorin, G.; Klein, M. L.; Hénin, J., Using collective variables to drive molecular dynamics simulations. *Molecular Physics* **2013,** *111* (22-23), 3345-3362.
4. Voth, G. A., *Coarse-graining of condensed phase and biomolecular systems*. CRC press: 2008.
5. Noid, W.; Chu, J.-W.; Ayton, G. S.; Krishna, V.; Izvekov, S.; Voth, G. A.; Das, A.; Andersen, H. C., The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models. *The Journal of chemical physics* **2008,** *128* (24), 244114.
6. Torrie, G. M.; Valleau, J. P., Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics* **1977,** *23* (2), 187-199.
7. Silverman, B. W., *Density estimation for statistics and data analysis*. Routledge: 2018.
8. LeCun, Y.; Bengio, Y.; Hinton, G., Deep learning. *nature* **2015,** *521* (7553), 436.
9. Goodfellow, I.; Bengio, Y.; Courville, A., *Deep learning*. MIT press: 2016.
10. Schneider, E.; Dai, L.; Topper, R. Q.; Drechsel-Grau, C.; Tuckerman, M. E., Stochastic neural network approach for learning high-dimensional free energy surfaces. *Physical review letters* **2017,** *119* (15), 150601.
11. Zhang, L.; Wang, H.; E, W., Reinforced dynamics for enhanced sampling in large atomic and molecular systems. *The Journal of chemical physics* **2018,** *148* (12), 124113.
12. Wang, J.; Olsson, S.; Wehmeyer, C.; Pérez, A.; Charron, N. E.; De Fabritiis, G.; Noé, F.; Clementi, C., Machine learning of coarse-grained molecular dynamics force fields. *ACS central science* **2019**.
13. Vapnik, V. N., An overview of statistical learning theory. *IEEE transactions on neural networks* **1999,** *10* (5), 988-999.
14. Kingma, D. P.; Welling, M., Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* **2013**.
15. Van den Oord, A.; Kalchbrenner, N.; Espeholt, L.; Vinyals, O.; Graves, A. In *Conditional image generation with pixelcnn decoders*, Advances in neural information processing systems, 2016; pp 4790-4798.
16. Dinh, L.; Krueger, D.; Bengio, Y., Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516* **2014**.
17. Rezende, D. J.; Mohamed, S., Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770* **2015**.
18. Abbasnejad, M. E.; Shi, Q.; Hengel, A. v. d.; Liu, L. In *A Generative Adversarial Density Estimator*, The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019/june.

19. Salakhutdinov, R.; Mnih, A.; Hinton, G. In *Restricted Boltzmann machines for collaborative filtering*, Proceedings of the 24th international conference on Machine learning, ACM: 2007; pp 791-798.
20. Salakhutdinov, R.; Hinton, G., Deep Boltzmann Machines. In *Proceedings of the Twelth International Conference on Artificial Intelligence and Statistics*, David van, D.; Max, W., Eds. PMLR: Proceedings of Machine Learning Research, 2009; Vol. 5, pp 448--455.
21. Hopfield, J. J., Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences* **1982,** *79* (8), 2554-2558.
22. Wales, D., *Energy landscapes: Applications to clusters, biomolecules and glasses*. Cambridge University Press: 2003.
23. Du, Y.; Mordatch, I., Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689* **2019**.
24. Kullback, S.; Leibler, R. A., On information and sufficiency. *The annals of mathematical statistics* **1951,** *22* (1), 79-86.
25. Arjovsky, M.; Chintala, S.; Bottou, L. In *Wasserstein generative adversarial networks*, International conference on machine learning, 2017; pp 214-223.
26. Shell, M. S., The relative entropy is fundamental to multiscale and inverse thermodynamic problems. *The Journal of chemical physics* **2008,** *129* (14), 144108.
27. Chaimovich, A.; Shell, M. S., Coarse-graining errors and numerical optimization using a relative entropy framework. *The Journal of chemical physics* **2011,** *134* (9), 094112.
28. Valsson, O.; Parrinello, M., Variational approach to enhanced sampling and free energy calculations. *Physical review letters* **2014,** *113* (9), 090601.
29. Zhang, J.; Yang, Y. I.; Noé, F., Targeted Adversarial Learning Optimized Sampling. **2019**.
30. Lyubartsev, A. P., Multiscale modeling of lipids and lipid bilayers. *European Biophysics Journal* **2005,** *35* (1), 53.
31. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. In *Generative adversarial nets*, Advances in neural information processing systems, 2014; pp 2672-2680.
32. Blei, D. M.; Kucukelbir, A.; McAuliffe, J. D., Variational inference: A review for statisticians. *Journal of the American Statistical Association* **2017,** *112* (518), 859-877.
33. Wu, D.; Wang, L.; Zhang, P., Solving statistical mechanics using variational autoregressive networks. *Physical Review Letters* **2019,** *122* (8), 080602.
34. Noé, F.; Olsson, S.; Köhler, J.; Wu, H., Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *science* **2019,** *365* (6457).
35. Dinh, L.; Sohl-Dickstein, J.; Bengio, S., Density estimation using Real NVP. In *International Conference on Learning Representations*, 2017.
36. Papamakarios, G.; Pavlakou, T.; Murray, I., Masked Autoregressive Flow for Density Estimation. In *Neural Information Processing Systems*, 2017; pp 2338-2347.
37. Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D., Neural Ordinary Differential Equations. In *Neural Information Processing Systems*, 2018; pp 6572-6583.
38. Grathwohl, W.; Chen, R. T. Q.; Bettencourt, J.; Sutskever, I.; Duvenaud, D., FFJORD: Free-Form Continuous Dynamics for Scalable Reversible Generative Models. In *International Conference on Learning Representations*, 2019.
39. Abbasnejad, M. E.; Shi, Q.; Hengel, A. v. d.; Liu, L., A Generative Adversarial Density Estimator. In *Computer Vision and Pattern Recognition*, 2019; pp 10782-10791.
40. Abrams, C.; Bussi, G., Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. *Entropy* **2014,** *16* (1), 163-199.
41. Okamoto, Y., Generalized-ensemble algorithms: enhanced sampling techniques for Monte Carlo and molecular dynamics simulations. *Journal of Molecular Graphics and Modelling* **2004,** *22* (5), 425-439.
42. Yang, Y. I.; Shao, Q.; Zhang, J.; Yang, L.; Gao, Y. Q., Enhanced sampling in molecular dynamics. *The Journal of Chemical Physics* **2019,** *151* (7), 070902.
43. Laio, A.; Parrinello, M., Escaping free-energy minima. *Proceedings of the National Academy of Sciences* **2002,** *99* (20), 12562-12566.
44. Sutton, R. S.; Barto, A. G., *Reinforcement learning: An introduction*. MIT press: 2018.
45. Grondman, I.; Busoniu, L.; Lopes, G. A.; Babuska, R., A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **2012,** *42* (6), 1291-1307.
46. Barducci, A.; Bussi, G.; Parrinello, M., Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Physical review letters* **2008,** *100* (2), 020603.
47. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. In *Gans trained by a two time-scale update rule converge to a local nash equilibrium*, Advances in Neural Information Processing Systems, 2017; pp 6626-6637.
48. Van Hasselt, H.; Guez, A.; Silver, D. In *Deep reinforcement learning with double q-learning*, Thirtieth AAAI conference on artificial intelligence, 2016.
49. Oord, A. v. d.; Li, Y.; Babuschkin, I.; Simonyan, K.; Vinyals, O.; kavukcuoglu, k.; Driessche, G. v. d.; Lockhart, E.; Cobo, L.; Stimberg, F.; Casagrande, N.;

Grewe, D.; Noury, S.; Dieleman, S.; Elsen, E.; Kalchbrenner, N.; Zen, H.; Graves, A.; King, H.; Walters, T.; Belov, D.; Hassabis, D., Parallel WaveNet: Fast High-Fidelity Speech Synthesis. In *International Conference on Machine Learning*, 2018; pp 3915-3923.

50. Bonati, L.; Zhang, Y.-Y.; Parrinello, M., Neural networks based variationally enhanced sampling. *arXiv preprint arXiv:1904.01305* **2019**.
51. Tiwary, P.; Berne, B., Predicting reaction coordinates in energy landscapes with diffusion anisotropy. *The Journal of chemical physics* **2017,** *147* (15), 152701.
52. Kmiecik, S.; Gront, D.; Kolinski, M.; Wieteska, L.; Dawid, A. E.; Kolinski, A., Coarse-Grained Protein Models and Their Applications. *Chemical Reviews* **2016,** *116* (14), 7898-7936.
53. Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E., How fast-folding proteins fold. *Science* **2011,** *334* (6055), 517-520.
54. Zhang, J.; Zhang, Z.; Yang, Y. I.; Liu, S.; Yang, L.; Gao, Y. Q., Rich Dynamics Underlying Solution Reactions Revealed by Sampling and Data Mining of Reactive Trajectories. *ACS Central Science* **2017,** *3* (5), 407-414.
55. Debnath, J.; Invernizzi, M.; Parrinello, M., Enhanced Sampling of Transition States. *Journal of Chemical Theory and Computation* **2019,** *15* (4), 2454-2459.
56. Zhang, J.; Yang, Y. I.; Yang, L.; Gao, Y. Q., Dynamics and Kinetics Study of “In-Water” Chemical Reactions by Enhanced Sampling of Reactive Trajectories. *The Journal of Physical Chemistry B* **2015,** *119* (45), 14505-14514.

Supplemental Information

Deep Learning for Multi-Scale Molecular Modeling

Jun Zhang[1,2], Yao-Kun Lei[3], Yi Isaac Yang[1,†] and Yi Qin Gao[1,3,4,5,‡]

[1] *Institute of Systems Biology, Shenzhen Bay Laboratory, 518055 Shenzhen, China*

[2] *Department of Mathematics and Computer Science, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany*

[3] *Beijing National Laboratory for Molecular Sciences, College of Chemistry and Molecular Engineering, Peking University, 100871 Beijing, China.*

[4] *Beijing Advanced Innovation Center for Genomics, Peking University, 100871 Beijing, China.*

[5] *Biomedical Pioneering Innovation Center, Peking University, 100871 Beijing, China.*

Corresponding authors: † (Y.I.Y) yangyi@szbl.ac.cn or ‡ (Y.Q.G) gaoyq@pku.edu.cn

TABLE OF CONTENTS

PART I. SUPPUPLEMTAL TEXTS

I. Training objectives of (RE-)VADE

Given the target distribution $p(\mathbf{s})$ and an approximate distribution $p_\theta(\mathbf{s})$, there are some practical ways to define the divergence between the two. The most commonly used one is the Kullback-Leibler divergence or relative entropy.

A. Kullback-Leibler divergence

$$D_{\mathrm{KL}}(p \,\|\, p_\theta) = \int p(\mathbf{s}) \log\left(\frac{p(\mathbf{s})}{p_\theta(\mathbf{s})}\right) d\mathbf{s} \tag{S1}$$

where both distributions are assumed to admit densities (i.e. absolutely continuous) with respect to a same measure. KL-divergence is known to be asymmetric and possibly infinite when there is non-overlapping area between the two distributions. Therefore, effective utilization of KL divergence as a optimization objective requires substantial overlap between $p(\mathbf{s})$ and $p_\theta(\mathbf{s})$[1].

B. Training objective of VADE

Now we consider the derivatives of Eq. (S1) w.r.t. the parameters θ, namely, $\nabla_\theta D_{\mathrm{KL}}(p \,\|\, p_\theta)$. First, we can re-write p and p_θ in terms of the free energy,

$$p(\mathbf{s}) = \frac{\exp\left(-\beta F(\mathbf{s})\right)}{Z} \tag{S2}$$

$$p_\theta(\mathbf{s}) = \frac{\exp\left(-\beta F_\theta(\mathbf{s})\right)}{Z_\theta} \tag{S3}$$

following which Eq. (S1) becomes

$$\begin{aligned} D_{\mathrm{KL}}(p \,\|\, p_\theta) = &\int \left[\beta F_\theta(\mathbf{s}) - \beta F(\mathbf{s})\right] p(\mathbf{s})\, d\mathbf{s} \\ &+ \int \left(\log Z_\theta - \log Z\right) d\mathbf{s} \end{aligned} \tag{S4}$$

It can be found that $D_{\mathrm{KL}}(p \,\|\, p_\theta)$ is now a functional of F_θ, hence we can obtain the functional derivatives,

$$\frac{\delta Z_\theta}{\delta F_\theta} = -\beta \exp\left(-\beta F_\theta\right) \tag{S5}$$

$$\frac{\delta D_{\mathrm{KL}}(p \,\|\, p_\theta)}{\delta F_\theta} = \frac{\beta \exp\left(-\beta F\right)}{Z} - \frac{\beta \exp\left(-\beta F_\theta\right)}{Z_\theta} \tag{S6}$$

Notice that Eq. (S6) is equivalent to

$$\frac{\delta D_{\mathrm{KL}}}{\delta F_\theta} = \beta p(\mathbf{s}) - \beta p_\theta(\mathbf{s}) \tag{S7}$$

so we finally arrive at the derivatives of D_{KL} w.r.t. θ,

$$\begin{aligned} \frac{\partial D_{\mathrm{KL}}}{\partial \theta} &= \int \frac{\delta D_{\mathrm{KL}}}{\delta F_\theta} \frac{\partial F_\theta(\mathbf{s})}{\partial \theta} d\mathbf{s} \\ &= \int \beta \frac{\partial F_\theta(\mathbf{s})}{\partial \theta} \left(p(\mathbf{s}) - p_\theta(\mathbf{s})\right) d\mathbf{s} \\ &= \left\langle \beta \frac{\partial F_\theta(\mathbf{s})}{\partial \theta} \right\rangle_{p(\mathbf{s})} - \left\langle \beta \frac{\partial F_\theta(\mathbf{s})}{\partial \theta} \right\rangle_{p_\theta(\mathbf{s})} \end{aligned} \tag{S8}$$

We remark here that Eq. (S8) holds for arbitrary fixed $p(\mathbf{s})$ when $p_\theta(\mathbf{s})$ is a learnable distribution used to approximate $p(\mathbf{s})$. So it can be easily generalized to the training objective of RE-VADE (Eq. (4) in the main text). Note that Eq. (S8) can be considered as the gradient of the following loss function,

$$L(\theta) = \left\langle \beta F_\theta(\mathbf{s}) \right\rangle_{p(\mathbf{s})} - \left\langle \beta F_\theta(\mathbf{s}) \right\rangle_{p_\theta(\mathbf{s})} \tag{S9}$$

C. Active mode-exploration in VADE

Since VADE yields the (unnormalized) probability $p_\theta(\mathbf{s})$ for a training sample $\mathbf{s}$, we can make use of this property to encourage VADE to learn the modes of the training samples which might have been ignored, in the spirit of importance sampling. Specifically, we can assign each training sample an importance weight $w(\mathbf{s})$ (normalized over mini-batch of samples),

$$w(\mathbf{s}) \propto \left[p_\theta(\mathbf{s})\right]^{\gamma - 1} \tag{S10}$$

where $\gamma \geq 1$ is a factor determining to what extent active mode-exploration is performed, and we re-write Eq. (S9) into,

$$L(\theta) = \left\langle w(\mathbf{s}) \beta F_\theta(\mathbf{s}) \right\rangle_{p(\mathbf{s})} - \left\langle \beta F_\theta(\mathbf{s}) \right\rangle_{p_\theta(\mathbf{s})} \tag{S11}$$

When $\gamma = 1$, Eq. (S11) reduces to the original VADE objective, i.e., Eq. (S9).

II. Training objectives of VADE-NS

A. Deep bijective models as neural samplers

Given a random variable $\mathbf{z} \sim q(\mathbf{z})$ and a function $f_\psi : \mathbb{R}^D \to \mathbb{R}^D$ that bijectively maps $\mathbf{z}$ into $\mathbf{s}$, where $D = \mathrm{Dim}(\mathbf{s}) = \mathrm{Dim}(\mathbf{z})$, then the probability of the random variable $\mathbf{s}$ can be computed according to the change-of-variables formula:

$$\log q_\psi(\mathbf{s}) = \log q(\mathbf{z}) - \log \left| \det \left(\frac{\partial f_\psi}{\partial \mathbf{z}} \right) \right| \tag{S12}$$

Calculating $q_\psi(\mathbf{s})$ using Eq. (S12) admits two key premises: (i) $q(\mathbf{z})$ is tractable so that drawing samples of $\mathbf{z}$ can be achieved with ease; (ii) the bijective function f_ψ has a tractable Jacobian whose determinant can be computed efficiently. In practice, one can simply choose a normal distribution as $q(\mathbf{z})$ to meet the first requirement. However, the computational cost of the determinant of the Jacobian of a bijective function generally scales as $\mathcal{O}(D^3)$, so ANNs of special architectures are needed for the second requirement.

Deep generative models that satisfy these two requirements include but not limited to normalizing-flow models,[2-4] auto-regressive flow models,[5-6] and we refer readers interested in this topic to the references.

B. Training objectives of VADE-NS

VADE-NS proposes a nested optimization problem as in GANs: Given a fixed VADE potential F_θ, a NS is optimized to approximate the equilibrium distribution of F_θ; while given an optimized NS, F_θ can be updated to fit density of the training data. Specifically, the NS is optimized following the gradient with F_θ fixed,

$$\begin{aligned}\nabla_\psi L(\psi) = &\left\langle \nabla_\psi \log q\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})} \\ &+\left\langle \beta \nabla_\psi F_\theta\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})}\end{aligned} \tag{S13}$$

where $L(\psi)$ denotes $D_{\mathrm{KL}}(q_\psi||p_\theta)$ in Eq. (4) of the main text, and f_ψ is a deep bijective model, e.g., a normalizing flow model taking the form of Eq. (S12). For continuous normalizing flows, $q\left(f_\psi(\mathbf{z})\right)$ is computed via black-box solvers for ordinary differential equation.[7-8]

On the other hand, when f_ψ is fixed, the VADE potential F_θ is optimized following the gradient,

$$\begin{aligned}\nabla_\theta L(\theta) = &\left\langle \beta \nabla_\theta F_\theta(\mathbf{s})\right\rangle_{p_{\mathrm{FG}}(\mathbf{s})} \\ &-\left\langle \beta \nabla_\theta F_\theta\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})}\end{aligned} \tag{S14}$$

where $L(\theta)$ is formulated by Eq. (S9).

Such nested optimization of VADE-NS through Eqs. (S13-S14) enjoys the merits that parametrization of the VADE potential is in free-form and independent of how f_ψ is chosen, so that F_θ can be designed arbitrarily expressive or variationally flexible. Besides, VADE-NS in this form can admit any prescribed additive restraint energies in the VADE potential F_θ.

In many cases where the functional form of F_θ is not prescribed, so that one can choose any black-box function approximator as F_θ, we can reduce the nested optimization problem by parameter-tying trick (or Eq. (6) in the main text),

$$\beta F_\theta(\mathbf{s}) = -\log q\left(f_\theta(\mathbf{z})\right) \tag{S15}$$

where $f_\theta: \mathbb{R}^D \to \mathbb{R}^D$ is a bijective function transforms $\mathbf{z}$ into $\mathbf{s}$. Then F_θ can be simply optimized following the gradient,

$$\begin{aligned}\nabla_\theta L(\theta) = &\left\langle \nabla_\theta \log q\left(f_\theta(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})} \\ &-\left\langle \nabla_\theta \log q\left(f_\theta^{-1}(\mathbf{s})\right)\right\rangle_{p_{\mathrm{FG}}(\mathbf{s})}\end{aligned} \tag{S16}$$

where f_θ^{-1} denotes the inverse function of f_θ that transforms $\mathbf{s}$ into $\mathbf{z}$.

Compared to the nested VADE-NS optimization problem Eqs. (S13-S14), Eq. (S16) is more convenient to implement in practice. Besides, the parameter-tying trick also reduces the number of necessary model parameters thus expedite the training of VADE potentials.

C. Incorporating parameter-free restraints in VADE-NS

Consider a VADE potential consisting of extra restraints, where the restraints take in the form of additive energies, F_{prior},

$$F_{\theta'}(\mathbf{s}) = F_\theta(\mathbf{s}) + F_{\mathrm{prior}}(\mathbf{s}) \tag{S17}$$

In the scenario of nested optimization, the gradient for the VADE potential F_θ remains the same as Eq. (S14), but the gradient for the NS becomes,

$$\begin{aligned}\nabla_\psi L(\psi) = &\left\langle \nabla_\psi \log q\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})} \\ &+\left\langle \beta \nabla_\psi F_\theta\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})} \\ &+\left\langle \beta \nabla_\psi F_{\mathrm{prior}}\left(f_\psi(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})}\end{aligned} \tag{S18}$$

To reduce the nested optimization problem into a standalone optimization problem, one can also adopt the parameter-tying trick (Eq. (S15) or Eq. (6) in the main text). Incorporating Eq. (S15), the first two terms on the R.H.S. of Eq. (S18) cancels out, and the gradient of the VADE-NS turns out to be,

$$\begin{aligned}\nabla_\theta L(\theta) = &\left\langle \nabla_\theta \log q\left(f_\theta(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})} \\ &-\left\langle \nabla_\theta \log q\left(f_\theta^{-1}(\mathbf{s})\right)\right\rangle_{p_{\mathrm{FG}}(\mathbf{s})} \\ &+\left\langle \beta \nabla_\theta F_{\mathrm{prior}}\left(f_\theta(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})}\end{aligned} \tag{S19}$$

The only difference of Eq. (S19) compared to Eq. (S19) is the inclusion of an additional term $\left\langle \beta \nabla_\theta F_{\mathrm{prior}}\left(f_\theta(\mathbf{z})\right)\right\rangle_{q(\mathbf{z})}$. This term can be regarded as a regularization over the VADE potential F_θ that encourages the generated CG samples to have low restraint energies (F_{prior}), similar to the effect of a bias potential widely adopted in restraint dynamics simulations.[9]

III. Deep energy functions in (RE-)VADE

There is relatively large margin for the specific functional forms for F_θ (and/or V_ϕ). For example, it can be a linear expansion of certain basis functions as in VES where θ are

the expansion coefficients [10]; or can be a non-linear neural network, where θ are the built-in parameters of neural network [11].

A. Orthonormal polynomials or functions

If $\mathrm{Dim}(\mathbf{s}) \leq 3$, we recommend orthonormal basis functions as F_θ and V_ϕ, and the expansion coefficients are the learnable parameters. For periodic CG variables $\mathbf{s}$, Fourier expansions can be adopted. For non-periodic $\mathbf{s}$, Legendre or Chebyshev polynomials can be used.[10] Orthonormal polynomials usually yield smooth energy function, so generally no additional regularization is needed.

B. Artificial neural networks

ANNs are expressive parametric models which can be used as F_θ and V_ϕ. Modern deep learning models have evolved based on specially designed architecture. Here we briefly introduce some building blocks which may be useful for constructing the deep models in (RE-)VADE. Multi-layer perceptrons (MLP) are most commonly seen ANNs consisting of fully connected hidden layer. The input of a MLP, $\mathbf{s}$, should trans-rotational invariant features of the molecular system. More importantly, each dimension of the input vector $\mathbf{s}$ should be indexible (or non-degenerate w.r.t. permutations). MLP will transform the input vector to hidden features, and finally yields an output vector.

C. Neural allocative potentials (NAP)

We note here that it is not the absolute value but the difference of energy makes physical sense. If we do not exploit this property, the absolute value of the energy produced by ANNs may grow rapidly but the energy difference does not. This is harmful to our purpose, because it was recently revealed that the smoothness and continuity of ANNs become vulnerable if the absolute values of built-in parameters are too large hence the weight matrices of ANNs become ill-conditioned.[12] Given the above considerations, we can first choose a lower bound and an upper bound for the bias potential, and "quantized" this energy range into K fixed levels $\{E_k\}_{k=1..K}$, based on which we propose the following functional form,

$$F_\theta(\mathbf{s}) = \sum_{k=1}^{K} \alpha_k(\mathbf{s};\theta) E_k \tag{S20}$$

$$\sum_{k=1}^{K} \alpha_k(\mathbf{s};\theta) = 1 \tag{S21}$$

where $\alpha_k(\mathbf{s};\theta)$ corresponds to the output of an ANN with a soft-max output layer. The associated force takes the following form,

$$-\nabla_{\mathbf{s}} F_\theta(\mathbf{s}) = \sum_{k=1}^{K} E_k \nabla_{\mathbf{s}} \alpha_k(\mathbf{s};\theta) \tag{S22}$$

In other words, we transform the problem of learning a scalar into learning a simplex $\{\alpha_k\}_{k=1..K}$. Following this form, NAP is trained to allocate proper amount of energy to configuration $\mathbf{s}$, rather than estimate the absolute value of the bias potential. Therefore, the training of $F_\theta(\mathbf{s})$ is more robust and the gradient will be well-behaving.

D. Gradient regularization

In order to ensure the to-be-optimized function $F_\theta(\mathbf{s})$ to be smooth enough for integration, we introduce several techniques to effectively regularize the gradient, $\nabla_{\mathbf{s}} F_\theta(\mathbf{s})$.

1. Weight clipping

One simple and intuitive approach is to clip the weights [1] of the neural network F_θ (e.g., $\theta \in \mathcal{W}[-0.01, 0.01]$) after each gradient update. The reason behind is that, if the parameter space $\mathcal{W}$ for θ is compact, then all the functions F_θ will be K-Lipschitz for some K that only depends on $\mathcal{W}$ and not the individual weights. In practice, the clipping range is a very important hyper-parameter, which can neither be too large (reduced compactness) nor too small (reduced capacity of F_θ). Besides, usually this method requires batch-normalization[13] in order to achieve a robust and good performance.

2. Spectral normalization

One more recent approach to stabilize the training of GAN can also be applied in (RE-)VADE, which is to perform spectral normalization for the weight matrices θ in the energy function F_θ. The basic idea is to constrain the Lipschitz constant of the function by restricting the spectral norm of each layer. Compared to other normalization techniques, spectral normalization does not require extra hyper-parameter tuning (setting the spectral norm of all weight layers to 1 consistently performs well in practice). Moreover, the computational cost is also relatively small. We refer readers interested in this technique to the reference.[12]

3. Drift removal

In experiments we found that removing the net drift of the energy function will help stabilize training. To be specific, we added a regularization term to the overall loss function, to restrict the mean value of the energy distribution of FG samples to be close to zero. This trick works possibly because the output of the network is anchored to a nearly zero-mean value, so the absolute value of the output (and the built-in parameters) will not be too large.

IV. Target distribution for RE-VADE

A. Well-Tempered (WT) distribution

Given the density estimation, i.e., $F_\theta(\mathbf{s})$, we can define the target distribution in WT form,[14] which is equivalent to a partially flattened FES,

$$p_{\mathrm{T}}(\mathbf{s}) = \frac{e^{-\frac{\beta}{\gamma} F_\theta(\mathbf{s})} d\mathbf{s}}{\int e^{-\frac{\beta}{\gamma} \beta F_\theta(\mathbf{s})} d\mathbf{s}} \tag{S23}$$

where $\gamma > 1$ is the WT factor. Noteworthy, with p_{T} in the form of Eq. (S23), the optimal F_θ and V_ϕ can be proved to take the following form,

$$\begin{aligned}\arg\min_{F_\theta} D_{\mathrm{KL}}\left(p \,\|\, p_\theta\right) &= F(\mathbf{s}) \\ \arg\min_{V_\phi} D_{\mathrm{KL}}\left(p_{\mathrm{T}} \,\|\, p_\phi\right) &= \left(\frac{1}{\gamma}-1\right)F(\mathbf{s})\end{aligned} \tag{S24}$$

B. Lorentzian-coupled distribution

As recently proposed by Debnath et al.,[15] the target distribution may take the following form:

$$p_{\mathrm{T}}(\mathbf{s}) = \frac{l(\mathbf{s})\, p_{\mathrm{WT}}(\mathbf{s};\gamma)}{\int l(\mathbf{s})\, p_{\mathrm{WT}}(\mathbf{s};\gamma)\, d\mathbf{s}} \tag{S25}$$

$$l(\mathbf{s}) = \frac{\zeta}{\zeta^2 + \left[\beta \dfrac{dF_\theta(\mathbf{s})}{d\mathbf{s}}\right]^2} \tag{S26}$$

where ζ is a scaling factor and $p_{\mathrm{WT}}(\mathbf{s};\gamma)$ is a pre-defined WT distribution with the WT-factor γ (Eq. S23). More introduction to this kind of target distribution can be found in the reference [15]. The advantage of such a Lorentzian-type target is that, the transition-state region (where $dF_\theta/d\mathbf{s} \approx 0$) can be further enhanced along with metastable states. This property is very useful in some scenarios where the sampling of transition state is crucial as in transition path sampling.

V. Connection between RE-VADE and other sampling methods

A. Metadynamics

Metadynamics[16] is a powerful tool in enhanced sampling of rare events if low-dimensional CV $\mathbf{s}$ can be defined *a priori*. However, metadynamics cannot be readily applied to cases where $\mathrm{Dim}(\mathbf{s}) \geq 3$. As introduced in the main text, RE-VADE can be viewed as a generalization of metadynamics into large $\mathrm{Dim}(\mathbf{s})$ and parametric bias potential functions.

(1) Metadynamics exploits KDE to coin the target distribution which will fail in large $\mathrm{Dim}(\mathbf{s})$. In contrast, RE-VADE uses parametric models to perform density estimation which is even applicable to large $\mathrm{Dim}(\mathbf{s})$.

(2) Metadynamics accumulated non-parametric Gaussian as bias potential which will be intractable for large $\mathrm{Dim}(\mathbf{s})$; while RE-VADE implements parametric learning to construct the bias potential, thus being more flexible and scalable for complex systems.

In summary, RE-VADE will be more useful if one wants to boost the sampling efficiency but does not have enough expert knowledge to determine low-dimensional CV.

B. VES

The mathematical form of VADE training objective is almost identical to VES [10], although it is derived from a different starting point. Both VES and RE-VADE aims to minimize the KL-divergence between the sampled distribution and a target distribution. Nevertheless, they can be distinguished by several remarkable differences:

(1) The target distribution in VES has to be defined carefully because VES would suffer from gradient vanishing or exploding issues when the distributions scarcely overlap. Usually this is a very demanding task. In contrast, RE-VADE follows actor-critic learning and coins a target distribution in a metadynamics-like fashion, thus always ensures a useful gradient for optimization.

(2) Optimization in VES relies on higher-order derivatives thus involving higher computational cost, while RE-VADE employs the state-of-the-art first-order optimization techniques from deep learning thus being robust, fast and economic.

(3) The functional form of bias potentials supported by vanilla VES is limited to linear expansions, thus the dimension of the CV space cannot be too large. In contrast, RE-VADE supports any differentiable functions (such as neural networks) as valid forms and can be easily extended to ultra-large dimensions.

C. TALOS

Both TALOS [17] and RE-VADE are closely connected to actor-critic reinforcement learning. They both parametrize a pair of value function and policy function in order to boost the rare events. But there are several key differences demarcating the two approaches:

(1) In TALOS, the value network is a critic (or discriminator) D_w as in GAN; whereas in RE-VADE, the value network F_θ is a density estimator.

(2) TALOS optimizes the policy network (i.e., the bias potential) to minimize the Wasserstein-1 distance between the sampled distribution and the target; while RE-VADE minimizes the KL-divergence for the same purpose.

(3) The target distribution in TALOS is manually selected *a priori*, the choice of which is sometimes tricky. While in RE-VADE, the target distribution is automatically determined and adaptively updated according to the current density estimation.

(4) The critic and policy networks in TALOS can be defined on different vector spaces; while in RE-VADE, both must operate on the same space.

We mark here that the strengths of TALOS and RE-VADE are highly complementary. It remains as an interesting research direction that one may combine both methods, that is, to minimize the Wasserstein-1 distance between the sampled distribution and an adaptively determined target distribution.

Algorithm S1. Reinforced Variational Adversarial Density Estimation (RE-VADE)

Input: Initialize value network F_θ and policy network V_ϕ; Learning rate α_θ and α_ϕ for θ and ϕ, respectively.

While ϕ do not converge, **do**

  Run MD under V_ϕ, collect samples to calculate $\langle V_\phi(\mathbf{s})\rangle_{p_\phi}$ ▷ run FG simulations

  Define $p_{\mathrm{FG}}(\mathbf{s}) \propto \exp\left(\beta V_\phi(\mathbf{s})\right)$ and calculate $\langle F_\theta(\mathbf{s})\rangle_{p_{\mathrm{FG}}}$ ▷ reweight biased samples

  For $t = 0 \rightarrow m$ ▷ train F_θ for m iterations

    Draw samples from F_θ, calculate $\langle F_\theta(\mathbf{s})\rangle_{p_\theta}$. ▷ run CG simulations

    Calculate $\mathcal{L}(\theta) = \langle \beta F_\theta(\mathbf{s})\rangle_{p_{\mathrm{FG}}} - \langle \beta F_\theta(\mathbf{s})\rangle_{p_\theta}$ ▷ Eq. (S9)

    $\theta \leftarrow \mathrm{Adam}\left(\nabla_\theta \mathcal{L}(\theta), \theta, \alpha_\theta\right)$ ▷ update value network

  End For

  Define target distribution p_{T} according to F_θ

  Calculate $\mathcal{L}(\phi) = \langle V_\phi(\mathbf{s})\rangle_{p_{\mathrm{T}}} - \langle V_\phi(\mathbf{s})\rangle_{p_\phi}$ ▷ Eq. (S9)

  $\phi \leftarrow \mathrm{Adam}\left(\nabla_\phi \mathcal{L}(\phi), \phi, \alpha_\phi\right)$ ▷ update bias potential

End While

PART II. SIMULATION DETAILS

I. VADE for numerical model: 2-dimensional 3-well potential

A. Simulation setup

The potential energy function of the 2D 3-well potential [18] takes the following form given the inverse temperature β:

$$\beta U(x,y) = -16\exp\left\{-2(x+0.5)^2 - 2(y-0.5)^2\right\}$$
$$-18\exp\left\{-2(x-0.8)^2 - 2(y-1.2)^2\right\}$$
$$-16\exp\left\{-2(x-0.5)^2 - 2(y+0.3)^2\right\} + 0.5\left(x^6 + y^6\right)$$

For the overdamped Langevin simulation, we chose $\beta = 1$, and the diffusion tensor was set to be:

$$\mathbf{D} = \begin{pmatrix} D_{xx} & 0 \\ 0 & D_{yy} \end{pmatrix} = \begin{pmatrix} (5\beta)^{-1} & 0 \\ 0 & (5\beta)^{-1} \end{pmatrix}$$

The resulting white-noised Langevin dynamics was simulated with a discrete time integration step of 0.01. 100,000 samples in total were collected every 100 integration steps. We randomly selected 50,000 samples as the training set, and 5,120 out of the rest of samples as validation set.

B. Model setups

For VADE-MC, F_θ is constructed on the 2D (x, y) space via a MLP. The MLP contains 2 hidden layers, each consisting of 64 units with softplus as the activation function. The hidden layers were also regularized by Spectral Normalization (see SI Part I, Section III) to smooth the gradients. The hidden layers were terminated by a one-unit linear output layer.

For VADE-NS, F_θ takes the form of a continuous normalizing flow (CNF) model with a free-form Jacobian (FFJORD).[8] The state transform function of FFJORD was modeled by a MLP consisting of 3 hidden layers, each composed of 64 units with softplus activation function, followed by a linear output layer composed of 2 units (the dimension of the output should be the same as the input). We stacked two CNF's of such architecture to improve the expressivity of the resulting model as suggested by the reference.[8] The base distribution for the latent variable $\mathbf{z}$, $q(\mathbf{z})$ (Eq. (S12) or Eq. (6) in the main text), was chosen to be the multi-variate normal distribution.

C. Training details

For VADE-MC, we trained F_θ with a mini-batch size of 500 and 100 epochs in total. Samples from F_θ was drawn through importance MC sampling. F_θ was optimized w.r.t. Eq. (S8) (or Eq. (3) in the main text). We further regularized the model by removing the drift of mean energy of the FG samples. The default Adam optimizer[19] with a learning rate of 10^{-4} was adopted. Optimization was performed on Tensorflow.

For VADE-NS, we trained F_θ with a mini-batch size of 500 and 100 epochs in total. Samples from F_θ was generated by the reversed dynamics of FFJORD CNF model. F_θ was optimized w.r.t. the reduced VADE-NS objective, i.e. Eq. (S16). The default Adam optimizer with a learning rate of 10^{-3} was adopted. Optimization was performed on Tensorflow.

II. VADE for Chignolin

A. Data source

The all-atom simulation data of chignolin is taken from reference [20], which contains over 500,000 samples in total. These samples were used to approximate p_{FG}.

B. Model setups

We imposed an exclusion energy F_{prior}, which penalizes if the distance between the *i*-th and *j*-th particles, r_{ij}, is shorter than a prescribed cutoff r_{ex}, and stays zero otherwise

$$F_{\mathrm{prior}}(\mathbf{s}) = \sum_{\mathrm{i}\neq\mathrm{j}} \left(r_{\mathrm{ex}} - r_{\mathrm{ij}}\right)^2 \qquad \text{(S27)}$$

and we chose $r_{\mathrm{ex}} = 0.1$ nm in our experiment.

We performed VADE-NS by employing a Masked Auto-regressive Flow (MAF) model[6] as the VADE potential F_θ. We first converted the input Cartesian coordinates of all C_α atoms, denoted by $\mathbf{s}$, into the trans-rotational invariant internal coordinates $\mathbf{s}_z$. To remove the non-invertibility caused by the periodicity of angles (and torsions) in $\mathbf{s}_z$, denoted by $-\pi \le \mathbf{s}_{\mathrm{rad}} < \pi$, we first performed a logit-transformation over the periodic variables,

$$\mathbf{s}_{z'} = \log\frac{\mathbf{x}}{1-\mathbf{x}} \qquad \text{(S28)}$$

$$\mathbf{x} = \alpha + (1-\alpha)\frac{\mathbf{s}_{\mathrm{rad}} + \pi}{2\pi} \qquad \text{(S29)}$$

where $\alpha = 0.05$ in Eq.(S29) to prevent overflow.

MAF was then operated on these transformed coordinates. The base distribution for the latent variable $\mathbf{z}$, $q(\mathbf{z})$ (Eq. (S12) or Eq. (5) in the main text), was chosen to be the multi-variate normal distribution. Code of MAF is adapted from GitHub: https://github.com/spinaotey/maf_tf.

C. Training details

For VADE-NS, we trained F_θ with a mini-batch size of 512 and 10 epochs in total. Samples from F_θ was generated according to Eq. (S12), that is, by the sampling $\mathbf{z}$ from the base distribution $q(\mathbf{z})$ and transforms $\mathbf{z}$ into $\mathbf{s}$ through the MAF model. We optimized the VADE potential F_θ w.r.t. the regularized VADE-NS objective i.e. Eq. (S19). The default

Adam optimizer with a learning rate of 10^{-4} was adopted. Optimization was performed on Tensorflow.

III. RE-VADE for Alanine dipeptide (Ala2)

A. Simulation setup

For the alanine dipeptide in aqueous solution, no ions were added since the terminal of the alanine was neutrally blocked (namely, ACE-ALA-NME) surrounded by 384 SPCE water molecules.[21] All the simulations were executed on AMBER17 package[22] using FF99SB force field parameters. [23] The aqueous solution system was put in a rectangular simulation box with periodic boundaries on. SHAKE algorithm [24] was adopted to constrain all covalent bonds involving hydrogen atoms, and a 2 fs time step was permitted. The system underwent a standard relaxation procedure and equilibrated to an NTP ensemble (300 K, 1atm). To equilibrate the system to the appropriate volume, the pressure of the system was adjusted to 1 atm by the Berendsen weak-coupling algorithm [25] with the relaxation time constants of 0.2 ps under another 1 ns long normal MD. For production run, samples were collected every 1 ps from 10 parallel walkers [26].

B. Model setups

Both F_θ and V_ϕ are functions of the 2D (ϕ,φ) space. Since ϕ and φ are periodic variables, we construct both functions using the Fourier polynomial:

$$F_\theta = \sum_{i=0}^{i=N}\sum_{j=0}^{j=N}\theta_{ij}\left[\cos(i\phi)+\sin(i\phi)\right]\left[\cos(j\varphi)+\sin(j\varphi)\right]$$

where $N=8$ for both F_θ and V_ϕ.

C. Target distribution

In order to train V_ϕ, we chose a WT-form target distribution with $1/\gamma = 0.4$ based on F_θ according to Eq. (S23).

D. Training details

In each RE-VADE iteration, we ran MD simulations for 40 ps to collect samples as p_{FG}. Before every training step of V_ϕ, F_θ was first trained for 5 steps (i.e., $m = 5$ in Algorithm S1). We adopted the Averaged SGD algorithm as recommended by Refs. [27] and [10] for each optimization step, and the update step-size for V_ϕ and F_θ is 0.1. Codes for optimization was inherited from Ref. [17] and [10]. In total, RE-VADE was performed for 200 iterations, which is equivalent to 8-ns MD simulation.

IV. RE-VADE for Claisen rearrangement

A. Simulation setup

The simulation was performed at the QM/MM interface on AMBER14 MD platform. The self-consistent charge density functional tight-binding (SCC-DFTB) method [28] was adopted to approximate the quantum mechanical Hamiltonian of the reactant molecule. The solvent is a kind of ionic liquid, containing a pair of soluble ion pairs termed as $[C_2mim]^+[NTf_2]^-$. We adopted the classical force field developed by Sieffert and Wipff [29-30] to describe the solvent molecules (or ions) and SHAKE was imposed on the solvent. No additional ions were added.

The system underwent a standard relaxation procedure and equilibrated to an NTP ensemble (300 K, 1 atm) lasting for 1-ns long normal MD. A cutoff of 10.0 Å was applied for calculating nonbonding interactions. All the simulations were performed with a 1-fs time integration step (no SHAKE on QM-treated molecule) and with periodic boundary condition. In production MD run, samples were collected every 0.5 ps from 8 parallel walkers.

B. Functional forms

In reference of the work by *Zhang et al.* [31], we set $s = 0.82d_2 - 0.18d_1$ as a one dimensional CV for enhanced sampling.

So both F_θ and V_ϕ are functions of s, and they are expanded by Legendre polynomials up to the same order. For instance, F_θ takes the following form,

$$F_\theta(s) = \sum_{i=1}^{i=N}\theta_i f_i(\bar{s})$$

where f_i denotes the i-th order Legendre polynomial, $N=50$, and $\bar{s}$ linearly rescales s to the range of [0,1], with $s_{\min}=0$ and $s_{\max}=0.35$ nm,

$$\bar{s} = \left(s - \frac{s_{\min}+s_{\max}}{2}\right)\Big/\left(\frac{s_{\max}-s_{\min}}{2}\right)$$

C. Target distribution

In order to enhance the chemical transition, we optimized V_ϕ according to a Lorentzian-coupled target distribution (Eqs. S25 and S26), where $p_{\mathrm{WT}}(s;\gamma)$ in Eq. (S25) takes a WT factor $1/\gamma = 0.2$ and $\zeta = 500$.

D. Training details

In each RE-VADE iteration, we ran MD simulations for 60 ps to collect samples as p_{FG}. Before every training step of V_ϕ, F_θ was first trained for 5 steps (i.e., $m = 5$ in Algorithm S1). Similarly to Ala2, we also adopted the Averaged SGD algorithm as for each optimization step, and the update step-size for V_ϕ and F_θ is 0.5. In total, RE-VADE was performed for 100 iterations, which is equivalent to 6-ns MD simulation.

E. Characterizing the chirality

We define the chirality order parameter q_{C} as follows [32]. First we pick the position of central carbon (corresponding to the chirality site) as $\mathbf{R}_{\mathrm{C}}$, and order the attached four atoms from $\mathbf{R}_0$ to $\mathbf{R}_3$ according to conventions (Fig. S3). We then define several geometric vectors (Fig. S3): $\boldsymbol{v}_i = \mathbf{R}_i - \mathbf{R}_{\mathrm{C}}$. In order to judge the chirality of the central carbon atom, we first calculate the cross product of two in-plane vectors,

$$\boldsymbol{u} = (\boldsymbol{v}_2 - \boldsymbol{v}_3)\times(\boldsymbol{v}_1 - \boldsymbol{v}_2)$$

Finally, the chirality order parameter q_C is defined as the cosine of the angle between $\boldsymbol{u}$ and the pseudo-normal vector $\boldsymbol{v}_0$, e vectors,

$$q_C = \frac{\boldsymbol{u} \cdot \boldsymbol{v}_0}{\|\boldsymbol{u}\|\|\boldsymbol{v}_0\|}$$

PART III. SUPPLEMENTAL FIGURES

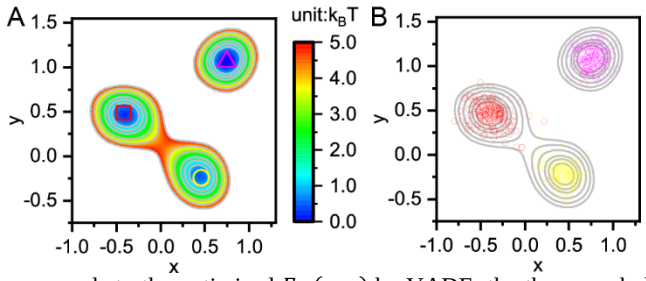

Figure S1. (A) The contour map corresponds to the optimized $F_{\theta^*}(x, y)$ by VADE, the three symbols (magenta triangle, red square and yellow circle) represent three free energy minima found in F_{θ^*}. (B) Hollow circles are simulation samples produced by Langevin dynamics on $U(x, y)$, colored by which local minimum they are minimized to over F_{θ^*}. Colors are in line with the symbols in panel a. The contour map of $F_{\theta^*}(x, y)$ (as in Panel A) is shown in grey as background.

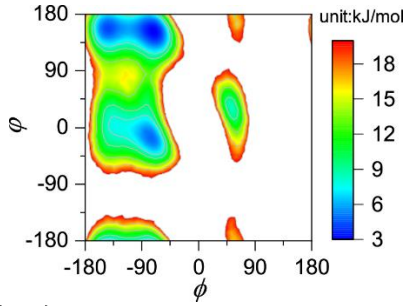

Figure S2. Reference FES for torsional angles (ϕ, φ) of Ala2. The reference FES is obtained via kernel density estimation over 750 ns MD simulation of Ala2 in explicit water. The data is accessible to the public from:
https://github.com/markovmodel/mdshare

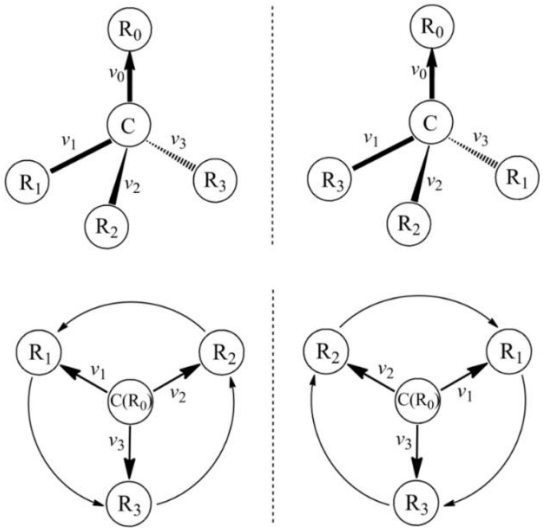

Figure S3. The chiral carbon atoms in the product. We order the attached atoms to the central carbon by the ordering R_3>R_2>R_1>R_0. Upper panels: Illustration of the ordering of the central chiral carbon atom and the geometric vectors. Lower panels: Clockwise sequence of the attached groups corresponds to R-configuration, and a counter-clockwise sequence to S-configuration. This figure is adapted from Ref. 32 with permission.

REFERENCES

1. Arjovsky, M.; Chintala, S.; Bottou, L. In *Wasserstein generative adversarial networks*, International conference on machine learning, 2017; pp 214-223.
2. Rezende, D. J.; Mohamed, S., Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770* **2015**.
3. Dinh, L.; Krueger, D.; Bengio, Y., Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516* **2014**.
4. Dinh, L.; Sohl-Dickstein, J.; Bengio, S., Density estimation using Real NVP. In *International Conference on Learning Representations*, 2017.
5. Germain, M.; Gregor, K.; Murray, I.; Larochelle, H., MADE: Masked Autoencoder for Distribution Estimation. In *International Conference on Machine Learning*, 2015; pp 881-889.
6. Papamakarios, G.; Pavlakou, T.; Murray, I., Masked Autoregressive Flow for Density Estimation. In *Neural Information Processing Systems*, 2017; pp 2338-2347.
7. Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D., Neural Ordinary Differential Equations. In *Neural Information Processing Systems*, 2018; pp 6572-6583.
8. Grathwohl, W.; Chen, R. T. Q.; Bettencourt, J.; Sutskever, I.; Duvenaud, D., FFJORD: Free-Form Continuous Dynamics for Scalable Reversible Generative Models. In *International Conference on Learning Representations*, 2019.
9. Torrie, G. M.; Valleau, J. P., Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics* **1977,** *23* (2), 187-199.
10. Valsson, O.; Parrinello, M., Variational approach to enhanced sampling and free energy calculations. *Physical review letters* **2014,** *113* (9), 090601.
11. Schneider, E.; Dai, L.; Topper, R. Q.; Drechsel-Grau, C.; Tuckerman, M. E., Stochastic neural network approach for learning high-dimensional free energy surfaces. *Physical review letters* **2017,** *119* (15), 150601.
12. Miyato, T.; Kataoka, T.; Koyama, M.; Yoshida, Y. In *Spectral Normalization for Generative Adversarial Networks*, International Conference on Learning Representations, 2018; 2018.
13. Ioffe, S.; Szegedy, C., Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* **2015**.
14. Barducci, A.; Bussi, G.; Parrinello, M., Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Physical review letters* **2008,** *100* (2), 020603.
15. Debnath, J.; Invernizzi, M.; Parrinello, M., Enhanced sampling of transition states. *Journal of chemical theory and computation* **2019,** *15* (4), 2454-2459.
16. Laio, A.; Parrinello, M., Escaping free-energy minima. *Proceedings of the National Academy of Sciences* **2002,** *99* (20), 12562-12566.
17. Zhang, J.; Yang, Y. I.; Noé, F., Targeted Adversarial Learning Optimized Sampling. **2019**.
18. Tiwary, P.; Berne, B., Predicting reaction coordinates in energy landscapes with diffusion anisotropy. *The Journal of chemical physics* **2017,** *147* (15), 152701.
19. Kingma, D. P.; Ba, J., Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* **2014**.
20. Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E., How fast-folding proteins fold. *Science* **2011,** *334* (6055), 517-520.
21. Berendsen, H.; Grigera, J.; Straatsma, T., The missing term in effective pair potentials. *Journal of Physical Chemistry* **1987,** *91* (24), 6269-6271.
22. Case, D.; Cerutti, D.; Cheatham III, T.; Darden, T.; Duke, R.; Giese, T.; Gohlke, H.; Goetz, A.; Greene, D.; Homeyer, N., AMBER 2017, 2017. *San Francisco: University of California*.
23. Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C., Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins: Structure, Function, and Bioinformatics* **2006,** *65* (3), 712-725.
24. Ryckaert, J.-P.; Ciccotti, G.; Berendsen, H. J. C., Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes. *Journal of Computational Physics* **1977,** *23* (3), 327-341.
25. Berendsen, H. J. C.; Postma, J. P. M.; Gunsteren, W. F. v.; DiNola, A.; Haak, J. R., Molecular dynamics with coupling to an external bath. *The Journal of Chemical Physics* **1984,** *81* (8), 3684-3690.
26. Raiteri, P.; Laio, A.; Gervasio, F. L.; Micheletti, C.; Parrinello, M., Efficient reconstruction of complex free energy landscapes by multiple walkers metadynamics. *The journal of physical chemistry B* **2006,** *110* (8), 3533-3539.
27. Bach, F.; Moulines, E. In *Non-strongly-convex smooth stochastic approximation with convergence rate O(1/n)*, neural information processing systems, 2013; pp 773-781.
28. Elstner, M.; Porezag, D.; Jungnickel, G.; Elsner, J.; Haugk, M.; Frauenheim, T.; Suhai, S.; Seifert, G., Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties. *Phys. Rev. B* **1998,** *58*, 7260-7268.

29. Sieffert, N.; Wipff, G., The [BMI][Tf2N] Ionic Liquid/Water Binary System: A Molecular Dynamics Study of Phase Separation and of the Liquid−Liquid Interface. *The Journal of Physical Chemistry B* **2006,** *110* (26), 13076-13085.
30. Fu, J.; Yang, Y. I.; Zhang, J.; Chen, Q.; Shen, X.; Gao, Y. Q., Structural Characteristics of Homogeneous Hydrophobic Ionic Liquid–HNO3–H2O Ternary System: Experimental Studies and Molecular Dynamics Simulations. *The Journal of Physical Chemistry B* **2016,** *120* (23), 5194-5202.
31. Zhang, J.; Zhang, Z.; Yang, Y. I.; Liu, S.; Yang, L.; Gao, Y. Q., Rich Dynamics Underlying Solution Reactions Revealed by Sampling and Data Mining of Reactive Trajectories. *ACS Central Science* **2017,** *3* (5), 407-414.
32. Zhang, J.; Yang, Y. I.; Yang, L.; Gao, Y. Q., Dynamics and Kinetics Study of "In-Water" Chemical Reactions by Enhanced Sampling of Reactive Trajectories. *The Journal of Physical Chemistry B* **2015,** *119* (45), 14505-14514.