
Embedded-physics machine learning for coarse-graining and collective variable discovery without data

Markus Schöberl*
Technical University of Munich
University of Notre Dame
mschoeberl@gmail.com

Nicholas Zabarás†
Center for Informatics and Computational Science
University of Notre Dame
nzabarás@gmail.com

Phaedon-Stelios Koutsourelakis‡
Professorship of Continuum Mechanics
Technical University of Munich
p.s.koutsourelakis@tum.de

Abstract

We present a novel learning framework that consistently embeds underlying physics while bypassing a significant drawback of most modern, data-driven coarse-grained approaches in the context of molecular dynamics (MD), i.e., the availability of big data. The generation of a sufficiently large training dataset poses a computationally demanding task, while complete coverage of the atomistic configuration space is not guaranteed. As a result, the explorative capabilities of data-driven coarse-grained models are limited and may yield biased “predictive” tools. We propose a novel objective based on reverse Kullback–Leibler divergence that fully incorporates the available physics in the form of the atomistic force field. Rather than separating model learning from the data-generation procedure - the latter relies on simulating atomistic motions governed by force fields - we query the atomistic force field at sample configurations proposed by the predictive coarse-grained model. Thus, learning relies on the evaluation of the force field but does not require any MD simulation. The resulting generative coarse-grained model serves as an efficient surrogate model for predicting atomistic configurations and estimating relevant observables. Beyond obtaining a predictive coarse-grained model, we demonstrate that in the discovered lower-dimensional representation, the collective variables (CVs) are related to physicochemical properties, which are essential for gaining understanding of unexplored complex systems. We demonstrate the algorithmic advances in terms of predictive ability and the physical meaning of the revealed CVs for a bimodal potential energy function and the alanine dipeptide.

Boltzmann densities, which are ensemble representations of equilibrium atomistic systems, are usually explored by molecular dynamics (MD) [1] or Monte Carlo-based (MC) techniques [2]. These versatile and general simulation techniques asymptotically guarantee unbiased estimates of observables [3]. However, these simulation techniques become computationally impractical in cases where the atomistic interaction potential exhibits several distinct minima or wells. Such complex potentials imply multimodal Boltzmann densities. Escaping such a well is rare and requires overcoming high free-energy barriers, resulting in impractically long simulation times or biased trajectories [4].

*<https://mjschoeberl.github.io/>, <http://orcid.org/0000-0001-7559-2619>

†311 Cushing Hall, Notre Dame, IN 46556, USA. <http://orcid.org/0000-0003-3144-8388>

‡Boltzmannstraße 15, 85748 Garching, Germany. <http://orcid.org/0000-0002-9345-759X>

Key to exploring such multimodal Boltzmann densities is the recognition of appropriate slow coordinates or collective variables (CVs) that exhibit sensitivity in transition regions between modes. This requires tremendous physicochemical insight, which is not available per se. CVs, which provide an effective lower-dimensional description of high-dimensional atomistic systems, are key to accelerating the exploration of multimodal densities by biasing the dynamics to escape deep free-energy wells [5].

Identifying expressive CVs governing major conformational changes in the absence of physical insight requires data-driven strategies. However, in many cases, the identification of CVs requires a dense sample of the target Boltzmann distribution, as well as unbiased simulation trajectories. This creates a contradiction, as it is computationally impractical to obtain unbiased trajectories in the presence of multiple modes [6], restricting the potential use-cases of such approaches. Few nonlinear dimensionality reduction methods coping with low data provide CVs that are not differentiable with respect to their atomistic counterparts [7–9]; however, this is required for biasing the dynamics [10–21]. Deep learning approaches providing flexible and efficiently differentiable functions have also influenced research on the efficient exploration of multimodal Boltzmann distributions. However, these build on previously acquired reference data and do not account directly for the interaction potential that actually drives the MD or MC simulation.

This work provides a novel and fundamentally different perspective on data-driven deep learning approaches. Instead of relying on two separate process, acquiring data and then employing statistical learning of a model, we synthesize and embed physics, i.e., the Boltzmann density, with a machine learning objective. The advocated learning methodology proposes (atomistic) configurations to which the model is attracted to learn from the potential energy and associated interatomic forces. The proposed machine learning algorithm does not require any simulation of the Boltzmann density but only queries the physical model, i.e., the potential and forces, to gain relevant information by evaluating rather than simulating the Boltzmann density. The proposed learning algorithm includes a versatile nonlinear dimensionality-reduction routine, which simultaneously discovers relevant CVs while learning the Boltzmann density. We demonstrate the procedure using a double well potential and the alanine dipeptide.

The present work differs clearly from recent developments on Boltzmann generators [22] that rely on invertible neural networks such as RealNVP [23] and NICE [24]. As it employs invertible neural networks, the dimensionality of the latent generator must equal the dimensionality of the atomistic configuration, which detains a consistent dimensionality reduction. Generated atomistic realizations of the employed model in [22] do not reflect the statistics of the reference Boltzmann distribution and serve instead as an input to a subsequent re-weighting importance sampling step. However, importance sampling is difficult to monitor if the variance in the importance weights is low, implying a large effective sample size, when none of the proposed realizations yield relatively high probabilities as evaluated by the target density [25]. Furthermore, a good guess of CVs is provided in Boltzmann generators, which depict physical insights that may not be available. By contrast, the proposed approach (similar double well example) reveals the effective CVs *and* also provides a generator to produce samples that yield the correct statistics of the target.

In the following Section 1, we develop the proposed learning approach based on KL divergence minimization and derive a tractable upper bound based on hierarchical variational models [26]. We discuss the required gradient computation and provide a physically interpretable underpinning of the components involved. After introducing a general model parametrization, we provide an adaptive tempering scheme facilitating a robust machine learning procedure at the end of Section 1. The proposed physics-embedding learning procedure for revealing CVs and obtaining a coarse-grained (CG) model is numerically validated in Section 2 with a double well potential and the alanine dipeptide. We close this paper in Section 3, summarizing the main findings of this work and outlining simple but effective further extensions and research directions. These include the generalization of the obtained predictive distribution for predictive purposes at any temperature.

1 Methodology

After introducing the notation in Section 1.1, we describe the general proposed framework in Section 1.2. A tractable optimization objective is provided in Section 1.3. We compare the proposed approach with data-driven objectives in Section 1.4. Relevant model specifications and gradient computations

for training are discussed in Section 1.5, and we close with some notes on the actual training procedure in Section 1.6.

1.1 Equilibrium statistical mechanics

In equilibrium statistical mechanics, we seek to estimate ensemble averages of observables $a(\mathbf{x})$ with respect to the Boltzmann density,

$$\langle a \rangle_{p_{\text{target}}(\mathbf{x};\beta)} = \int_{\mathcal{M}_f} a(\mathbf{x}) p_{\text{target}}(\mathbf{x};\beta) d\mathbf{x}. \quad (1)$$

We denote the Boltzmann distribution, for which we aim to learn an efficient approximation, by $p_{\text{target}}(\mathbf{x};\beta)$:

$$\begin{aligned} p_{\text{target}}(\mathbf{x}) &= \frac{1}{Z(\beta)} \underbrace{e^{-\beta U(\mathbf{x})}}_{\pi(\mathbf{x};\beta)} \\ &= \frac{\pi(\mathbf{x};\beta)}{Z(\beta)}. \end{aligned} \quad (2)$$

In Equation 2, $Z(\beta) = \int e^{-\beta U(\mathbf{x})} d\mathbf{x}$ is the partition function or normalization constant, and $\beta = \frac{1}{k_B T}$ is the reciprocal or inverse temperature with the Boltzmann constant k_B and the temperature T . The interatomic potential $U(\mathbf{x})$ depends on generalized atomistic coordinates denoted by $\mathbf{x} \in \mathcal{M}_f \subset \mathbb{R}^{n_f}$, with $n_f = \dim(\mathbf{x})$. In equilibrium statistical mechanics, we are usually interested in phase averages at distinct constant temperatures; however, we will also demonstrate how to utilize the temperature to introduce an auxiliary sequence of target distributions to facilitate learning the actual target distribution. The auxiliary sequence stabilizes the parameter learning inspired by annealing [27, 28] and adaptive sequential MC [29].

1.2 Coarse-graining through probabilistic generative models

Data-driven coarse-graining methodologies are based on a limited set of N realizations obtained from the target density $p_{\text{target}}(\mathbf{x})$. The realizations $\mathbf{x}^{(i)}$ are produced by drawing samples from $p_{\text{target}}(\mathbf{x})$: $\mathbf{x}^{(i)} \sim p_{\text{target}}(\mathbf{x})$ with Markov Chain MC (MCMC) methods [30, 31] and/or, especially in the context of biochemical atomistic systems, by MD simulations [32, 33]. Both methodologies yield a dataset $\mathbf{x}^{\mathcal{D}_N} = \{\mathbf{x}^{(i)}\}_{i=1}^N$, which approximates the target distribution with

$$\begin{aligned} p_{\text{target}}(\mathbf{x}) &\approx \tilde{p}(\mathbf{x}) \\ &\propto \prod_{i=1}^N \delta(\mathbf{x}^{(i)} - \mathbf{x}). \end{aligned} \quad (3)$$

The above approximation, given independent and identically distributed samples, may sufficiently resemble simple systems. However, atomistic many-body systems exhibit higher-order and long-range interactions [34, 35] involving multiple free energy modes separated by high barriers [36, 37]. Therefore, the collection of sufficient data becomes an insurmountable task: a protein folding process may take microseconds versus a time discretization of femtoseconds [38]. Given limited computational power, the relevant conformations and transitions are not guaranteed to be reflected by the reference simulation [39].

The quality of data-driven learning approaches depends strongly on the quality of the available set of reference data $\mathbf{x}^{\mathcal{D}_N}$. If, e.g., in the case of peptides, certain conformations are missed, it is an almost insurmountable challenge to obtain a data-driven model exploring such missed configurations [40, 41]. Enhanced sampling methods [10–21] can support the exploration of the configuration space, while the efficiency crucially depends on the quality of utilized CVs [42, 43].

Instead of relying on reference data, which may be a distorted representation of $p_{\text{target}}(\mathbf{x})$, or gradually exploring the configuration space by enhanced sampling, we present a variational approach that learns the target distribution $p_{\text{target}}(\mathbf{x})$ by querying the unnormalized distribution $\pi(\mathbf{x})$ or the corresponding potential energy $U(\mathbf{x})$ (see Equation 2).

We are first interested in identifying latent CVs \mathbf{z} ($\mathbf{z} \in \mathcal{M}_c \subset \mathbb{R}^{n_c}$) depending on the fully atomistic picture \mathbf{x} , which encode physically relevant characteristics (e.g., coordinates along transition regions

between conformations) and provide insight into the unknown atomistic system we seek to explore. Second, we seek to identify a CG model expressed in terms of the latent CVs \mathbf{z} that is predictive but nevertheless facilitates reasoning about all-atom coordinates \mathbf{x} [44]. The obtained CG model is expected to serve as an approximation to $p_{\text{target}}(\mathbf{x})$ to enable the efficient computation of the expectations of observables (Equation 1) and most importantly to capture relevant configurations in the free energy landscape that were inaccessible by brute-force MD or MCMC approaches [45]. Fulfilling the latter requirement also implies capturing statistics of $p_{\text{target}}(\mathbf{x})$.

The CVs \mathbf{z} serve as latent generators of the higher-dimensional generalized coordinates \mathbf{x} , where we seek $\dim(\mathbf{z}) \ll \dim(\mathbf{x})$. This generative process is expressed with two components,

- (i) the conditional density $q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})$, parametrized by θ_{cf} ,
- (ii) and the density over the latent CVs $q_{\theta_c}(\mathbf{z})$.

Combining both densities gives the following joint:

$$q_{\theta}(\mathbf{x}, \mathbf{z}) = q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_c}(\mathbf{z}). \quad (4)$$

Assuming we have obtained the optimal parameters θ^{opt} after a training process based on an objective, which we will discuss later in this section, we can utilize the model for predictive purposes. This can be done by ancestral sampling [25], i.e., first draw $\mathbf{z}^{(i)} \sim q_{\theta^{\text{opt}}}(\mathbf{z})$ and second $\mathbf{x}^{(i)} \sim q_{\theta^{\text{opt}}}(\mathbf{x}|\mathbf{z}^{(i)})$.

For obtaining optimal parameters θ , many methods rely on minimizing a distance from the target distribution $p_{\text{target}}(\mathbf{x})$ to the marginal distribution $q_{\theta}(\mathbf{x})$, which is given by:

$$q_{\theta}(\mathbf{x}) = \int q_{\theta}(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \int q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_c}(\mathbf{z}) d\mathbf{z}. \quad (5)$$

A commonly employed metric expressing the deviation between two densities is the Kullback–Leibler (KL) divergence, which belongs to the family of α -divergences [46–48]:

$$\begin{aligned} D_{\text{KL}}(p_{\text{target}}(\mathbf{x})||q_{\theta}(\mathbf{x})) &= - \int p_{\text{target}}(\mathbf{x}) \log \frac{q_{\theta}(\mathbf{x})}{p_{\text{target}}(\mathbf{x})} d\mathbf{x} \\ &= - \langle \log q_{\theta}(\mathbf{x}) \rangle_{p_{\text{target}}(\mathbf{x})} + \underbrace{\langle \log p_{\text{target}}(\mathbf{x}) \rangle_{p_{\text{target}}(\mathbf{x})}}_{-\mathbb{H}(p_{\text{target}})}. \end{aligned} \quad (6)$$

Minimizing Equation 6 with respect to θ leads to $q_{\theta}(\mathbf{x})$ being closer to $p_{\text{target}}(\mathbf{x})$. However, in practice, the expectations in Equation 6 are intractable:

- (i) the marginal $q_{\theta}(\mathbf{x})$ requires the integration with respect to \mathbf{z} which is intractable itself and
- (ii) the involved expectation with respect to $p_{\text{target}}(\mathbf{x})$, $\langle \cdot \rangle_{p_{\text{target}}(\mathbf{x})}$ is analytically intractable since the normalization constant of $p_{\text{target}}(\mathbf{x})$ is unavailable (which would require solving an integral with respect to \mathbf{x}).

Considering the above challenges, the latter could be addressed by approximating $p_{\text{target}}(\mathbf{x})$ with data or samples $\mathbf{x}^{\mathcal{D}_N}$ and thus approximating the corresponding expectations with MC estimators. However, as we deal with complex multimodal Boltzmann densities $p_{\text{target}}(\mathbf{x})$, the data generating process (MCMC or MD) may miss relevant modes. By employing a biased set of samples or data not approximating $p_{\text{target}}(\mathbf{x})$ [49], we learn a biased estimator not approximating $p_{\text{target}}(\mathbf{x})$. The generation of the training dataset is thus decoupled from the learning process.

To circumvent the data-generating process and thus sampling from $p_{\text{target}}(\mathbf{x})$, we propose employing the other extreme of the family of α -divergences (as compared to Equation 6), the *reverse* KL divergence:

$$\begin{aligned} D_{\text{KL}}(q_{\theta}(\mathbf{x})||p_{\text{target}}(\mathbf{x})) &= - \int q_{\theta}(\mathbf{x}) \log \frac{p_{\text{target}}(\mathbf{x})}{q_{\theta}(\mathbf{x})} d\mathbf{x} \\ &= \underbrace{\mathbb{E}_{q_{\theta}(\mathbf{x})} [\log q_{\theta}(\mathbf{x})]}_{-\mathbb{H}(q_{\theta}(\mathbf{x}))} - \mathbb{E}_{q_{\theta}(\mathbf{x})} [\log p_{\text{target}}(\mathbf{x})] \\ &= -\mathbb{E}_{q_{\theta}(\mathbf{x})} [\log p_{\text{target}}(\mathbf{x})] - \mathbb{H}(q_{\theta}(\mathbf{x})). \end{aligned} \quad (7)$$

Minimizing Equation 7 with respect to θ requires maximizing the log-likelihood $\log p_{\text{target}}(\mathbf{x})$ assessed under $q_{\theta}(\mathbf{x})$ (first component in Equation 7), and the maximization of the entropy of $q_{\theta}(\mathbf{x})$, $\mathbb{H}(q(\mathbf{x}))$ (second component in Equation 7). Minimizing the reverse KL divergence balances the two terms, as maximizing only the log-likelihood $\log p_{\text{target}}(\mathbf{x})$ assessed under $q_{\theta}(\mathbf{x})$ would result in a degenerate case where $q_{\theta}(\mathbf{x})$ would become a Dirac-delta placed at the (global) maximum of $p_{\text{target}}(\mathbf{x})$ obtained at the (global) minimum of $U(\mathbf{x})$. The second component implies a regularization favoring a parametrization θ such that the entropy of $q_{\theta}(\mathbf{x})$ is maximized.

1.3 Inference and learning

In what follows, we use the negative of the KL divergence in Equation 7 to be maximized, which we denote with \mathcal{L} for the sake of comparability with other learning approaches [50, 51]. At the end of this section, we compare the presented methodology with data-driven approaches relying on the forward KL divergence [51–53] and especially those addressing coarse-graining problems [54–57].

The objective to be maximized is

$$\mathcal{L}(\theta) = \mathbb{E}_{q_{\theta}(\mathbf{x})} [\log p_{\text{target}}(\mathbf{x}) - \log q_{\theta}(\mathbf{x})], \quad (8)$$

where we can draw samples from $q_{\theta}(\mathbf{x})$ as we can wisely select tractable hierarchical components composing to $q_{\theta}(\mathbf{x})$. The optimization of the first component in $\mathcal{L}(\theta)$ relating to the log-likelihood is tractable as the normalization of $p_{\text{target}}(\mathbf{x})$ does not depend on the parameters θ and thus being able to evaluate $\pi(\mathbf{x})$ or $U(\mathbf{x})$ suffices. However, the entropy term is not tractable ad-hoc as it involves the marginal $q_{\theta}(\mathbf{x}) = \int q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_{\text{c}}}(\mathbf{z}) d\mathbf{z}$, posing in most cases an intractable or least cumbersome task.

Therefore, we seek to construct a tractable lower bound on $\mathbb{H}(\mathbf{x})$ as presented in [26] by introducing an auxiliary density $r_{\phi}(\mathbf{z}|\mathbf{x})$ parametrized by ϕ and write:

$$\begin{aligned} -\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x})] &= -\mathbb{E}_{q(\mathbf{x})} \left[\log q(\mathbf{x}) + \underbrace{D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|q(\mathbf{z}|\mathbf{x}))}_{=0} \right] \\ &\geq -\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x}) + D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|r_{\phi}(\mathbf{z}|\mathbf{x}))] \\ &= -\mathbb{E}_{q(\mathbf{x})} [\mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log q(\mathbf{x}) + \log q(\mathbf{z}|\mathbf{x}) - \log r_{\phi}(\mathbf{z}|\mathbf{x})]]. \end{aligned} \quad (9)$$

Adding $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|q(\mathbf{z}|\mathbf{x}))$ in the first line of Equation 9 has no influence as the term is equal to zero. It involves the posterior distribution over the latent variables \mathbf{z} , $q(\mathbf{z}|\mathbf{x}) = \frac{q(\mathbf{x},\mathbf{z})}{q(\mathbf{x})}$, which is intractable. By utilizing an auxiliary distribution $r_{\phi}(\mathbf{z}|\mathbf{x})$, the equality becomes an inequality as a consequence of $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|r_{\phi}(\mathbf{z}|\mathbf{x})) \geq 0$ for $r_{\phi}(\mathbf{z}|\mathbf{x})$ deviating from $q(\mathbf{z}|\mathbf{x})$. Replacing the exact log-posterior $\log q(\mathbf{z}|\mathbf{x})$ by

$$\log q(\mathbf{z}|\mathbf{x}) = \log q(\mathbf{z}) + \log q(\mathbf{x}|\mathbf{z}) - \log q(\mathbf{x}), \quad (10)$$

it follows that

$$\begin{aligned} -\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x})] &\geq -\mathbb{E}_{q(\mathbf{x})} \left[\mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log q(\mathbf{x}) + \log q(\mathbf{z}) + \log q(\mathbf{x}|\mathbf{z}) \right. \\ &\quad \left. - \log q(\mathbf{x}) - \log r_{\phi}(\mathbf{z}|\mathbf{x})] \right] \\ &= -\mathbb{E}_{q(\mathbf{x})} [\mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log q(\mathbf{z}) + \log q(\mathbf{x}|\mathbf{z}) - \log r_{\phi}(\mathbf{z}|\mathbf{x})]]. \end{aligned} \quad (11)$$

Rewriting the expectation $\mathbb{E}_{q(\mathbf{x})} [\mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\cdot]]$ as $\mathbb{E}_{q(\mathbf{x},\mathbf{z})} [\cdot]$, Equation 11 depicts a tractable lower bound on the entropy term. Maximizing the lower bound in Equation 11 with respect to ϕ minimizes $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|r(\mathbf{z}|\mathbf{x};\phi))$ and thus tightens the bound on the entropy term. As mentioned earlier, the optimum⁴ is obtained when we identify the exact posterior of the latent CVs $q(\mathbf{z}|\mathbf{x})$ with $r(\mathbf{z}|\mathbf{x};\phi^{\text{opt}}) = q(\mathbf{z}|\mathbf{x})$, thus $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x})\|r(\mathbf{z}|\mathbf{x};\phi^{\text{opt}})) = 0$. Utilizing the obtained bound in Equation 11, the objective $\mathcal{L}(\phi, \theta)$ from Equation 8 becomes:

$$\mathcal{L}(\phi, \theta) = \mathbb{E}_{q(\mathbf{x},\mathbf{z};\theta)} [\log p_{\text{target}}(\mathbf{x}) - \log q_{\theta_{\text{c}}}(\mathbf{z}) - \log q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z}) + \log r_{\phi}(\mathbf{z}|\mathbf{x})]. \quad (12)$$

⁴The optimum with respect to $r_{\phi}(\mathbf{z}|\mathbf{x})$ and thus ϕ that tightens the lower bound.

The following shows the connection between the obtained objective and the KL divergence defined between the joint $q(\mathbf{x}|\mathbf{z})q(\mathbf{z})$ and $p_{\text{target}}(\mathbf{x})r(\mathbf{z}|\mathbf{x})$ acting on the extended probability space:

$$\begin{aligned}\mathcal{L}(\phi, \theta) &= \mathbb{E}_{q(\mathbf{x}, \mathbf{z}; \theta)} [\log p_{\text{target}}(\mathbf{x}) - \log q(\mathbf{z}) - \log q(\mathbf{x}|\mathbf{z}) + \log r_{\phi}(\mathbf{z}|\mathbf{x})] \\ &= \mathbb{E}_{q(\mathbf{x}, \mathbf{z}; \theta)} \left[\log \frac{p_{\text{target}}(\mathbf{x})r_{\phi}(\mathbf{z}|\mathbf{x})}{q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_{\text{c}}}(\mathbf{z})} \right]\end{aligned}\quad (13)$$

$$= -D_{\text{KL}}(q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_{\text{c}}}(\mathbf{z}) \| p_{\text{target}}(\mathbf{x})r_{\phi}(\mathbf{z}|\mathbf{x})). \quad (14)$$

Based on Equation 13 and Equation 14, we show how the objective separates into two KL divergence terms:

$$\begin{aligned}D_{\text{KL}}(q(\mathbf{x}|\mathbf{z})q(\mathbf{z}) \| p_{\text{target}}(\mathbf{x})r(\mathbf{z}|\mathbf{x})) &= -\mathbb{E}_{q(\mathbf{z})} \left[\mathbb{E}_{q(\mathbf{x}|\mathbf{z})} \left[\log \frac{p_{\text{target}}(\mathbf{x})r_{\phi}(\mathbf{z}|\mathbf{x})}{q(\mathbf{z}|\mathbf{x})q(\mathbf{x})} \right] \right] \\ &= -\mathbb{E}_{q(\mathbf{x})} \left[\log \frac{p_{\text{target}}(\mathbf{x})}{q(\mathbf{x})} \right] - \mathbb{E}_{q(\mathbf{z})q(\mathbf{x}|\mathbf{z})} \left[\log \frac{r_{\phi}(\mathbf{z}|\mathbf{x})}{q(\mathbf{z}|\mathbf{x})} \right] \\ &= D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x})) + D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| r_{\phi}(\mathbf{z}|\mathbf{x})) \\ &\geq D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}))\end{aligned}\quad (15)$$

As mentioned earlier, the lower bound on $\mathcal{L}(\phi, \theta)$ or upper bound on $D_{\text{KL}}(q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})q_{\theta_{\text{c}}}(\mathbf{z}) \| p_{\text{target}}(\mathbf{x})r(\mathbf{z}|\mathbf{x}))$ becomes tight when $r(\mathbf{z}|\mathbf{x}; \phi^{\text{opt}}) = q(\mathbf{z}|\mathbf{x})$, which is Equation 15. Suboptimal ϕ imply bounds on the objective owing to the positivity of $D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| r_{\phi}(\mathbf{z}|\mathbf{x})) \geq 0$.

The advantage of the proposed method for identifying CVs and learning a predictive coarse-graining model becomes clearer when we directly utilize the reference potential energy $U(\mathbf{x})$ (which we assume to be available in this paper). The objective $\mathcal{L}(\phi, \theta)$, which is the negative KL divergence defined by the joint distributions, is subject to maximization with respect to the parameters θ and ϕ :

$$\begin{aligned}\mathcal{L}(\phi, \theta) &= -D_{\text{KL}}(q_{\theta_{\text{c}}}(\mathbf{z})q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z}) \| p_{\text{target}}(\mathbf{x})r_{\phi}(\mathbf{z}|\mathbf{x})) \\ &= \langle \log p_{\text{target}}(\mathbf{x}) \rangle_{q_{\theta}(\mathbf{x}, \mathbf{z})} + \left\langle \log \frac{r_{\phi}(\mathbf{z}|\mathbf{x})}{q_{\theta_{\text{c}}}(\mathbf{z})q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})} \right\rangle_{q_{\theta}(\mathbf{z}, \mathbf{x})} \\ &= -\beta \langle U(\mathbf{x}) \rangle_{q_{\theta}(\mathbf{x}, \mathbf{z})} + \left\langle \log \frac{r_{\phi}(\mathbf{z}|\mathbf{x})}{q_{\theta_{\text{c}}}(\mathbf{z})q_{\theta_{\text{cf}}}(\mathbf{x}|\mathbf{z})} \right\rangle_{q_{\theta}(\mathbf{z}, \mathbf{x})}\end{aligned}\quad (16)$$

Maximizing Equation 16 solely involves expectations with respect to the generative model, from which it is easy to draw samples from. Explicitly there are no expectations with respect to the target density $p_{\text{target}}(\mathbf{x})$, which would require an approximation with data. Instead of data, the target density $p_{\text{target}}(\mathbf{x})$ contributes to the learning of the parameters (ϕ, θ) through the interatomic potential energy $U(\mathbf{x})$ assessed for samples of the generative model $q(\mathbf{x}|\mathbf{z})q(\mathbf{z})$. Note that the normalization constant of $p_{\text{target}}(\mathbf{x})$ is independent of ϕ and θ and has been omitted in Equation 16. We are aware that the method requires a potential energy function $U(\mathbf{x})$, which can be assessed at \mathbf{x} . This is always the case for systems where we can set up MD or MCMC simulations, although we do circumvent the need to simulate a trajectory or draw reference samples by directly incorporating the available physics expressed by the potential energy.

1.4 Reverse or forward KL divergence?

In the following, we point out commonalities and differences between the proposed approach relying on the reverse KL divergence as introduced in Equation 7 and the forward KL divergence (Equation 6). The latter has been successfully employed for the development of coarse-graining methodologies [54–56] and with a focus on CV discovery in combination with predictive coarse-graining in [57].

The data-driven objective is based on minimizing the following KL divergence:

$$D_{\text{KL}}(p_{\text{target}}(\mathbf{x}) \| q_{\theta}(\mathbf{x})). \quad (17)$$

Reformulating the minimization of Equation 17 to a maximization problem, the lower bound based on the summation over terms corresponding to each datum $\mathbf{x}^{(i)}$ of a set of $\mathbf{x}^{\mathcal{D}_N} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ is

written as:

$$\begin{aligned}
\mathcal{L}^{\text{forward}}(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}^{\mathcal{D}_N}) &= \sum_{i=1}^N \mathbb{E}_{r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})} \left[-\log r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)}) + \log q_{\boldsymbol{\theta}}(\mathbf{x}^{(i)}, \mathbf{z}^{(i)}) \right] \\
&= - \sum_{i=1}^N D_{\text{KL}} \left(r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)}) \| q_{\boldsymbol{\theta}}(\mathbf{z}^{(i)}) \right) \\
&\quad + \sum_{i=1}^N \mathbb{E}_{r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})} \left[\log q_{\boldsymbol{\theta}}(\mathbf{x}^{(i)}|\mathbf{z}^{(i)}) \right]. \tag{18}
\end{aligned}$$

The objective above depicts the lower bound on the marginal log-likelihood and has been constructed in the context of data-driven variational inference [51, 53, 58]. The first component in Equation 18 implies minimizing $D_{\text{KL}}(r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)}) \| q_{\boldsymbol{\theta}}(\mathbf{z}^{(i)}))$ in an aggregation of all considered $\mathbf{x}^{(i)}$. Hence, in aggregation the pre-images of $\mathbf{x}^{(i)}$, expressed by the approximate posterior, should resemble the generative component $q_{\boldsymbol{\theta}}(\mathbf{z})$, whereas the latter term in Equation 18 accounts for the reconstruction loss of encoded pre-images $\mathbf{z}^{(i)}$ (encoded through $r_{\boldsymbol{\phi}}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})$) to its origin $\mathbf{x}^{(i)}$.

Minimizing the reverse KL divergence as introduced in Equation 7 with

$$D_{\text{KL}}(q_{\boldsymbol{\theta}}(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}))$$

implies a tractable maximization with respect to $(\boldsymbol{\phi}, \boldsymbol{\theta})$ of the following objective based on [26]:

$$\mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\theta}) = \underbrace{-\beta \langle U(\mathbf{x}) \rangle_{q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})}}_* + \underbrace{\mathbb{E}_{q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})} [\log r_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})]}_{\dagger} + \underbrace{\mathbb{H}(q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}))}_{\ddagger}. \tag{19}$$

We comment on the meaning of the indicated terms in Equation 19; however, note that the optimization always needs to be regarded in the composition of all terms.

- *) Maximizing $\mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\theta})$ seeks to minimize $\beta \langle U(\mathbf{x}) \rangle_{q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})}$, which corresponds to the average potential energy of the system evaluated under the generative model $q_{\boldsymbol{\theta}}(\mathbf{x})$.
- †) Maximize the expected log-probability that a given fine-scale realization $\mathbf{x}^{(i)}$ (with the corresponding latent pre-image $\mathbf{z}^{(i)}$) drawn from the joint of the generative model, $q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})$, can be reconstructed by $r_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})$ based on $\mathbf{x}^{(i)}$.
- ‡) Maximize the entropy of the generative model $\mathbb{H}(q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}))$.

Note that all aforementioned contributions must be seen in the composition, and maximizing $\mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\theta})$ with respect to $(\boldsymbol{\phi}, \boldsymbol{\theta})$ maximizes the balance of all. Most important is that the involved objective in the reverse KL divergence does not encompass any expectations with respect to $p_{\text{target}}(\mathbf{x})$, which need to be approximated by data as is the case in $\mathcal{L}^{\text{forward}}(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x}^{\mathcal{D}_N})$.

We discuss in the next section the particulars of the optimization with respect to $\boldsymbol{\theta}, \boldsymbol{\phi}$, and also specify the form of the densities involved, i.e., $q_{\boldsymbol{\theta}}$ and $r_{\boldsymbol{\phi}}$.

1.5 Model specification and gradient derivation

In the sequel we introduce a general approach for parametrizing distributions $q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})$ and $r_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})$ and provide an approach for optimizing parameters with variance-reduction methods, enabling accelerated convergence.

1.5.1 Model specification

We base the model specification on previous work in the context of data-driven CVs discovery [57]. The model involves two components, ($q(\mathbf{x}|\mathbf{z})$ and $q(\mathbf{z})$), with respect to the generative path and the encoder $r(\mathbf{z}|\mathbf{x})$ in the recognition path.

As we seek to obtain a set of lower-dimensional coordinates representing characteristic and slow coordinates of the system, we aim to provide complexity in the mapping and thus the encoder and decoder components $r(\mathbf{z}|\mathbf{x})$ and $q(\mathbf{x}|\mathbf{z})$, respectively, and simple descriptions of the CVs through

$q(\mathbf{z})$. Pushing complexity to the involved mappings and assuming simple correlations in $q(\mathbf{z})$ yields CVs capturing the most relevant features of the atomistic system compressed in low dimensions [59, 60].

The distribution $q_{\theta_c}(\mathbf{z})$, which the obtained CVs are supposed to follow and which we desire to be simple, is represented as a standard Gaussian with unit diagonal variance:

$$q_{\theta_c}(\mathbf{z}) = q(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I}). \quad (20)$$

The simplicity induced by Equation 20 is balanced by employing a flexible mapping given latent CVs \mathbf{z} to fine-scale atomistic coordinates \mathbf{x} (probabilistic decoder) with

$$q_{\theta_{cf}}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_{\theta_{cf}}(\mathbf{z}), \mathbf{S}_{\theta_{cf}}), \quad (21)$$

where the nonlinear mapping

$$\boldsymbol{\mu}_{\theta_{cf}}(\mathbf{z}) = f_{\theta_{cf}}^{\boldsymbol{\mu}}(\mathbf{z}), \quad (22)$$

with $\mathbf{z} \mapsto f_{\theta_{cf}}^{\boldsymbol{\mu}}(\mathbf{z})$ ($f_{\theta_{cf}}^{\boldsymbol{\mu}}: \mathbb{R}^{n_c} \mapsto \mathbb{R}^{n_f}$) is expressed by a flexible (multilayer) neural network [61–63]. The Gaussian in Equation 21 with the flexible mean $\boldsymbol{\mu}_{\theta_{cf}}(\mathbf{z})$ is then fully defined by considering a diagonal covariance matrix with $\mathbf{S}_{\theta_{cf}} = \text{diag}(\boldsymbol{\sigma}_{\theta_{cf}}^2)$ [64]. We omit the subscripts of $\boldsymbol{\theta}$, as the latent generator $q(\mathbf{z})$ does not depend on parameters. Thus, we write $\boldsymbol{\theta} = \boldsymbol{\theta}_{\theta_{cf}}$. We treat the entries $\sigma_{\theta,j}^2$ directly as parameters without dependence on latent CVs \mathbf{z} . Maintaining $\sigma_{\theta,j}^2 > 0$ is ensured by optimizing $\log \sigma_{\theta,j}^2$ instead.

In a similar fashion, compared to the model of $q_{\theta_{cf}}(\mathbf{x}|\mathbf{z})$, we express the encoder that approximates the actual posterior distribution $p(\mathbf{z}|\mathbf{x})$ as follows:

$$r_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_{\phi}(\mathbf{x}), \mathbf{S}_{\phi}(\mathbf{x})), \quad (23)$$

with the diagonal covariance matrix $\mathbf{S}_{\phi}(\mathbf{x}) = \text{diag}(\boldsymbol{\sigma}_{\phi}^2(\mathbf{x}))$. Likewise, $\boldsymbol{\mu}_{\phi}(\mathbf{x})$ and $\log \boldsymbol{\sigma}_{\phi}^2(\mathbf{x})$ are obtained from encoding neural networks $f_{\phi}^{\boldsymbol{\mu}}(\mathbf{x})$ and $f_{\phi}^{\boldsymbol{\sigma}}(\mathbf{x})$, respectively:

$$\boldsymbol{\mu}_{\phi}(\mathbf{x}) = f_{\phi}^{\boldsymbol{\mu}}(\mathbf{x}) \quad \text{and} \quad \log \boldsymbol{\sigma}_{\phi}^2(\mathbf{x}) = f_{\phi}^{\boldsymbol{\sigma}}(\mathbf{x}). \quad (24)$$

The actual but intractable posterior $q(\mathbf{z}|\mathbf{x})$ will differ from a multivariate normal distribution, for which we compensate by providing a flexible mean in $r_{\phi}(\mathbf{z}|\mathbf{x})$. Structural correlations revealed by a full rank covariance matrix represent an interesting avenue to be explored [65]; however, this is not part of this paper. The employed models resemble those developed earlier in the context of CV discovery. Therefore, we refer to the discussion in [57] justifying the use of the neural networks.

We utilize the following general structure for the decoding neural network $f_{\theta_{cf}}^{\boldsymbol{\mu}}(\mathbf{z})$:

$$f_{\theta}^{\boldsymbol{\mu}, K_q}(\mathbf{z}) = \left(l_{\theta_{cf}}^{(K_q+1)} \circ \tilde{a}^{(K_q)} \circ l_{\theta}^{(K_q)} \circ \dots \circ \tilde{a}^{(1)} \circ l_{\theta}^{(1)} \right) (\mathbf{z}). \quad (25)$$

with K_q hidden layers. In a similar manner, we define the encoding networks for $\boldsymbol{\mu}_{\phi}(\mathbf{x})$ and $\boldsymbol{\sigma}_{\phi}^2(\mathbf{x})$ of $r_{\phi}(\mathbf{z}|\mathbf{x})$:

$$f_{\phi}^{K_r}(\mathbf{x}) = \left(a^{(K_r)} \circ l_{\phi}^{(K_r)} \circ \dots \circ a^{(1)} \circ l_{\phi}^{(1)} \right) (\mathbf{x}), \quad (26)$$

which leads to the mean and diagonal terms of the covariance matrix with

$$f_{\phi}^{\boldsymbol{\mu}}(\mathbf{x}) = l_{\phi}^{(K_r+1)} \left(f_{\phi}^{K_r}(\mathbf{x}) \right) \quad \text{and} \quad f_{\phi}^{\boldsymbol{\sigma}}(\mathbf{x}) = l_{\phi}^{(K_r+2)} \left(f_{\phi}^{K_r}(\mathbf{x}) \right). \quad (27)$$

The linear layers used in the above expressions are denoted as $l^{(i)}$, e.g., mapping a variable \mathbf{y} to the output with $l^{(i)}(\mathbf{y}) = \mathbf{W}^{(i)}\mathbf{y} + \mathbf{b}^{(i)}$. The nonlinearities in $f_{\phi}^{(\cdot)}$ are implied by activation $a(\cdot)$. Encoding and decoding functions are indicated by the superscripts ϕ and θ , respectively. Activation functions belonging to the encoder are $a^{(i)}$, and those involved in decoding \mathbf{z} are $\tilde{a}^{(i)}$. The size of $\mathbf{W}^{(i)}$ is specified by the input dimension, which could be the output of a precedent layer $l^{(i-1)}(\mathbf{y})$, and the output dimension, which we specify with $d_{l^{(i)}}$. This leads to a matrix $\mathbf{W}^{(i)}$ of dimension $d_{l^{(i)}} \times d_{l^{(i-1)}}$. The corresponding parametrization details with depth K and activations of the networks are specified with the corresponding numerical illustrations in Section 2.

1.5.2 Gradient computation and reparametrization

This section is devoted to deriving relevant gradients of the objective $\mathcal{L}(\phi, \theta)$ in Equation 12, which involve the fine-scale potential energy $U(\mathbf{x})$. We show a noise-reducing gradient estimator by utilizing reparametrization [66, 67].

The focus is on the first component in Equation 12, which depends only on the parameters θ . We write for the corresponding derivative:

$$\begin{aligned} -\beta \frac{\partial}{\partial \theta} \langle U(\mathbf{x}) \rangle_{q_\theta(\mathbf{x}|\mathbf{z})} &= -\beta \frac{\partial}{\partial \theta} \int \int q_\theta(\mathbf{x}|\mathbf{z}) q(\mathbf{z}) U(\mathbf{x}) d\mathbf{x} d\mathbf{z} \\ &= -\beta \int q(\mathbf{z}) \underbrace{\frac{\partial}{\partial \theta} \left(\int q_\theta(\mathbf{x}|\mathbf{z}) U(\mathbf{x}) d\mathbf{x} \right)}_{\nabla_\theta \mathbb{E}_{q_\theta(\mathbf{x}|\mathbf{z})} [U(\mathbf{x})]} d\mathbf{z}. \end{aligned} \quad (28)$$

In the last line of the above equation, we note the expression $\nabla_\theta \mathbb{E}_{q_\theta(\mathbf{x}|\mathbf{z})} [U(\mathbf{x})]$; this is for the case of using approximate MC estimators, highly affected by noise, as discussed in [53]. This would hamper the optimization even when employing stochastic techniques. The variance of the approximate estimator of $\nabla_\theta \mathbb{E}_{q_\theta(\mathbf{x}|\mathbf{z})} [U(\mathbf{x})]$ can be reduced by the reparametrization of $q_\theta(\mathbf{x}|\mathbf{z})$. This is done by introducing an auxiliary random variable ϵ , which gives rise to \mathbf{x} by a differentiable transformation:

$$\mathbf{x} = g_\theta(\epsilon; \mathbf{z}) \text{ with } \epsilon \sim p(\epsilon). \quad (29)$$

With the mapping, $g_\theta : \epsilon \rightarrow \mathbf{z}$, the following holds by change of variables:

$$q_\theta(\mathbf{x}|\mathbf{z}) = p(g_\theta^{-1}(\mathbf{x}; \mathbf{z})) \left| \frac{\partial g_\theta^{-1}(\mathbf{x}; \mathbf{z})}{\partial \mathbf{x}} \right|, \quad (30)$$

where the inverse function of g_θ , $g_\theta^{-1} : \mathbf{x} \rightarrow \epsilon$ leads to $\epsilon = g_\theta^{-1}(\mathbf{x}; \mathbf{z})$. Different possibilities of auxiliary distributions and invertible transformations are discussed in more detail in [68]. With the introduced transformation, we can rewrite the derivative with:

$$\begin{aligned} \nabla_\theta \mathbb{E}_{q_\theta(\mathbf{x}|\mathbf{z})} [U(\mathbf{x})] &= \mathbb{E}_{p(\epsilon)} [\nabla_\theta U(g_\theta(\epsilon; \mathbf{z}))] \\ &= \mathbb{E}_{p(\epsilon)} \left[\frac{\partial U(g_\theta(\epsilon; \mathbf{z}))}{\partial \mathbf{x}} \frac{\partial g_\theta(\epsilon; \mathbf{z})}{\partial \theta} \right]. \end{aligned} \quad (31)$$

The auxiliary random variables ϵ follow a Gaussian with $\epsilon^{(l)} \sim p(\epsilon) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. The corresponding transformation for representing the random variables \mathbf{x} is:

$$\mathbf{x} = g_\theta(\epsilon; \mathbf{z}) = \boldsymbol{\mu}_\theta(\mathbf{z}) + \boldsymbol{\sigma}_\theta(\mathbf{z}) \odot \epsilon. \quad (32)$$

Replacing the expression $\nabla_\theta \mathbb{E}_{q_\theta(\mathbf{x}|\mathbf{z})} [U(\mathbf{x})]$ in Equation 28 with Equation 31 leads to:

$$\begin{aligned} -\beta \frac{\partial}{\partial \theta} \langle U(\mathbf{x}) \rangle_{q_\theta(\mathbf{x}|\mathbf{z})} &= -\beta \left\langle \mathbb{E}_{p(\epsilon)} \left[\frac{\partial U(g_\theta(\epsilon; \mathbf{z}))}{\partial \mathbf{x}} \frac{\partial g_\theta(\epsilon; \mathbf{z})}{\partial \theta} \right] \right\rangle_{q(\mathbf{z})} \\ &= -\beta \left\langle \underbrace{\frac{\partial U(g_\theta(\epsilon; \mathbf{z}))}{\partial \mathbf{x}}}_{=-\mathbf{F}(\mathbf{x})} \frac{\partial g_\theta(\epsilon; \mathbf{z})}{\partial \theta} \right\rangle_{p(\epsilon)q(\mathbf{z})}. \end{aligned} \quad (33)$$

First, the physically less interesting part in Equation 33 is the contribution $\frac{\partial g_\theta(\epsilon; \mathbf{z})}{\partial \theta}$, which can be estimated by employing efficient backpropagation and automatic differentiation algorithms for neural networks [61, 69]. However, the more physically relevant component, the gradient of the atomistic potential, $\frac{\partial U(g_\theta(\epsilon; \mathbf{z}))}{\partial \mathbf{x}}$, is involved in Equation 33. The gradient of the potential $U(\mathbf{x})$ with respect to \mathbf{x} equals the negative interatomic force $\mathbf{F}(\mathbf{x})$, evaluated at \mathbf{x} , where $\mathbf{x} = g_\theta(\epsilon; \mathbf{z})$. This latter term incorporates physics into the gradient computation of $\mathcal{L}(\phi, \theta)$ in the form of interatomic forces. This is the source from which physics are embedded into our proposed model and drives the optimization of $\mathcal{L}(\phi, \theta)$ by querying the force field at samples of $q_\theta(\mathbf{x})$. Notably, the forces are incorporated at atomistic positions \mathbf{x} , which are determined by sampling as follows.

- (i) Draw a sample from the generative distributions: $\mathbf{z}^{(i)} \sim q(\mathbf{z})$ which is simple to sample from.
- (ii) Then obtain a sample from the auxiliary distribution: $\boldsymbol{\epsilon}^{(j)} \sim p(\boldsymbol{\epsilon})$.
- (iii) Determine the corresponding atomistic representation of $(\mathbf{z}^{(i)}, \boldsymbol{\epsilon}^{(j)})$ with: $\mathbf{x}^{(i,j)} = g_{\theta}(\boldsymbol{\epsilon}^{(j)}; \mathbf{z}^{(i)}) = \boldsymbol{\mu}_{\theta}(\mathbf{z}^{(i)}) + \boldsymbol{\sigma}_{\theta}(\mathbf{z}^{(i)}) \odot \boldsymbol{\epsilon}^{(j)}$.

This means we evaluate the force \mathbf{F} at samples $\mathbf{x}^{(i,j)}$; no reference data are required in this process.

The force evaluation at atomistic coordinates \mathbf{x} is the heart of common MD software such as LAMMPS [70], GROMACS [71–77], and OpenMM [78]. The MD simulators are highly sophisticated in terms of efficiency and allow us to employ this optimized force evaluation function in our development.

In this work, we develop a PyTorch module that incorporates OpenMM [78] in the backward pass, which enables efficient optimization by querying the forces computed by OpenMM at input positions governed by $q_{\theta}(\mathbf{x})$. We are continuously developing the software on GPU platforms, and it will be made available⁵.

1.6 Training

Training the model parameters (ϕ, θ) requires some attention as variational models tend to be mode-focusing [79, 80]. If parameters update too rapidly, in terms of configurations of $p_{\text{target}}(\mathbf{x})$ that have been explored by $q_{\theta}(\mathbf{x})$ thus far, relevant conformations could be missed. However, compared with the data-driven approach, the proposed variational coarse-graining methodology offers strategies ensuring that relevant conformations are captured and incorporates querying of the potential $U(\mathbf{x})$ into the learning procedure. In data-driven schemes, once the data is obtained, there is no control on exploring unobserved conformations [6, 60]. Remedy, next to employing stochastic optimization with adaptive step size control [81], provide tempering approaches [82, 83]. These start at high initial temperatures or low inverse temperatures, with, e.g., $0 \leq \beta_1$, whereas $\beta = 0$ resembles a uniform target distribution. A sequence of K temperatures and related inverse temperatures $0 \leq \beta_1 \cdots \leq \beta_k \leq \dots \leq \beta_K$ yields a sequence of target distributions with [84–86]

$$p_{\text{target}}(\mathbf{x}; \beta_k) = \frac{1}{Z(\beta_k)} e^{-\beta_k U(\mathbf{x})}, \quad \forall k \in \{1, \dots, K\}, \quad (34)$$

while β_K equals the target simulation temperature β_{target} .

Instead of directly minimizing $D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_{\text{target}}))$, we minimize subsequent $D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_k))$ while we obtain optimal (ϕ_k, θ_k) , which we use as initial parameters for minimizing $D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_{k+1}))$. However, the size of the increment between two subsequent temperature steps $\Delta\beta_k = \beta_{k+1} - \beta_k$ is a difficult choice.

Therefore, we develop an adaptive scheme for gradually increasing β_k , which adjusts the proposed $\Delta\beta_k$ such that the relative difference in subsequent KL divergences estimated at β_k and β_{k+1} does not exceed a threshold c_{max} . We define the relative increase of the KL divergence between β_k and β_{k+1} with:

$$\frac{D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_{k+1})) - D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_k))}{D_{\text{KL}}(q(\mathbf{x}) \| p_{\text{target}}(\mathbf{x}; \beta_k))}. \quad (35)$$

By employing the derived upper bound on the KL divergence, which is defined in Equation 16, we can rewrite Equation 35 as

$$c_k = \frac{\log(Z(\beta_{k+1})) - \log(Z(\beta_k)) + (\beta_{k+1} - \beta_k) \langle U(\mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})}}{\log Z(\beta_k) + \beta_k \langle U(\mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})} - \langle \log r(\mathbf{z} | \mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})} - \mathbb{H}(q(\mathbf{x}, \mathbf{z}))}. \quad (36)$$

Besides the (log-)difference of the normalization constants, $\log(Z(\beta_{i+1})) - \log(Z(\beta_i))$, and $\log(Z(\beta_k))$, all remaining components in Equation 36 are accessible through MC estimators. The

⁵Software available upon publication on <https://github.com/mjs.../...>

supporting material in Appendix B includes an approximation of $\log(Z(\beta_{k+1})) - \log(Z(\beta_k))$ and $\log(Z(\beta_k))$. The procedure for updating the temperature is summarized in Algorithm 1.

Algorithm 1: Tempering scheme for updating β_k . We set $\Delta\beta_{\max} = 1.0 \times 10^{-3}$ and $c_{\max} = 1.0$.

Input: Converged model with its parameter (ϕ_k, θ_k) at current inverse temperature β_k ; c_{\max} , maximal relative increase in D_{KL} ; $\Delta\beta_{\max}$, the temperature increment; Current step k .

Output: β_{k+1}

- 1 **Initialize:** $s := 0, f_s := 1.0$.
 - 2 **while** $c_k^s > c_{\max}$ **do**
 - 3 **Propose new inverse temperature** β_{k+1}^s :
 $\beta_{k+1}^s = \beta_k + f_s \Delta\beta_{\max}$
 - 4 **Estimate rel. increase** c_k^s **with proposed** β_{k+1}^s :
 See computation in Equation 36.
 - 5 **Update** f_s **for proposing a new maximal increase in** β :
 $f_{s+1} = 0.6f_s$
 - 6 **Update step:**
 $s = s + 1$.
 - 7 **Set:** $\beta_{k+1} = \beta_{k+1}^s$
 - 8 **Update:** $k = k + 1$
 - 9 **Continue optimization with:** $\log p_{\text{target}}(\mathbf{x}; \beta_k) \propto e^{\beta_k U(\mathbf{x})}$
-

2 Numerical illustrations

The following section demonstrates the developed methodology based on a double well potential in Section 2.1 and an alanine dipeptide in Section 2.2.

2.1 Double well

This section shows the capabilities of the proposed method in the context of a two-dimensional double well potential energy function $U(\mathbf{x})$ ($\dim(\mathbf{x}) = 2$) that exhibits two distinct modes distinguishable in the x_1 direction. One of the modes is favorably explored owing to its lower potential energy. The potential is quadratic in the x_2 direction, as depicted in Figure 1:

$$U(\mathbf{x}) = \frac{1}{4}x_1^4 - 3 \cdot x_1^2 + x_1 + \frac{1}{2}x_2^2. \quad (37)$$

The double well potential in Equation 37 and the implied target distribution $p_{\text{target}}(\mathbf{x}; \beta = 1) \propto e^{-\beta U(\mathbf{x})}$ result in a distribution that is challenging to explore with purely random walk MCMC and without performing extensive fine-tuning of the proposal step. A test MCMC estimator, which was as fair as possible, did not discover the second mode for $x_1 > 0$ after 1×10^5 steps. The natural CV of the potential $U(\mathbf{x})$ and thus of $p_{\text{target}}(\mathbf{x}) \propto e^{-U(\mathbf{x})}$ is the x_1 coordinate. The x_1 direction distinguishes the two modes that $p_{\text{target}}(\mathbf{x})$ exhibits. We expect our algorithm to reveal CVs z “equal” to x_1 or having high correlation with x_1 . We put “equal” in quotes as we work in a probabilistic framework. The dimensionality of $\dim(\mathbf{z})$ is 1.

The functional form and parameters have been taken from [22] to ensure comparability. However, note that we seek to identify simultaneously the lower-dimensional characteristics revealing the relevant physics, encoded in CVs, and obtain a generative CG model for predictive purposes. In [22], the focus was on the generative component. The CVs utilized for learning are *selected* rather than revealed from the physics. The latent CVs \mathbf{z} have the same dimensionality as \mathbf{x} owing to the use of *invertible* neural networks that require $\dim(\mathbf{z}) = \dim(\mathbf{x})$ [23].

We employ the model as introduced in Section 1.5.1 and define the unspecified options such as the number of layers, layer dimensions, and activation functions used in the encoder and decoder as in Tables 1 and 2, respectively. To train the parameters (ϕ, θ) , we employ a tempering scheme as introduced in Section 1.6 and specified in Algorithm 1 with initial $\beta_0 = 1 \times 10^{-10}$, while the target is defined with $\beta_K = 1$. For all numerical illustrations, we employ ADAM stochastic optimization [88] with $\alpha = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon_{\text{ADAM}} = 1.0 \times 10^{-8}$. The expectations with respect to $q(\mathbf{x}, \mathbf{z})$ are computed based on $J = 1000$ samples.




Figure 1: Reference potential energy $U(\mathbf{x})$. The color relates to the value of $U(\mathbf{x})$ quantified by the *logarithmic* color bar on the right. Most MCMC random walk approaches will discover only one of the depicted potential energy basins.

Linear layer	Input dimension	Output dimension	Activation layer	Activation function
$l_\phi^{(1)}$	$\dim(\mathbf{x}) = 2$	d_1	$a^{(1)}$	SeLu ⁶
$l_\phi^{(2)}$	d_1	d_2	$a^{(2)}$	SeLu
$l_\phi^{(3)}$	d_2	d_3	$a^{(3)}$	TanH
$l_\phi^{(4)}$	d_3	$\dim(\mathbf{z})$	None	-
$l_\phi^{(5)}$	d_3	$\dim(\mathbf{z})$	None	-

Table 1: Network specification of the encoding neural network with $d_{\{1,2,3\}} = 100$.

Linear layer	Input dimension	Output dimension	Activation layer	Activation function
$l_\theta^{(1)}$	$\dim(\mathbf{z}) = 1$	d_3	$\tilde{a}^{(1)}$	Tanh
$l_\theta^{(2)}$	d_3	d_2	$\tilde{a}^{(2)}$	Tanh
$l_\theta^{(3)}$	d_1	$\dim(\mathbf{x})$	None	-

Table 2: Network specification of the decoding neural network with $d_{\{1,2,3\}}$ as defined in Table 1.

We will assess the trained model with respect to its predictive power and the obtained CVs in the following.

2.1.1 Predictive CG model

Figures 2 and 3 show intermediate results obtained while training the model. The left columns depict a two-dimensional (2D) histogram containing the target histogram based on a long reference simulation obtained by employing the Metropolis-adjusted Langevin algorithm [89] at $\beta = 1$. Next to the histogram of $p_{\text{target}}(\mathbf{x}; \beta = 1)$, we provide 2D histograms of intermediate predictions at β_k , as indicated in the sub-caption. The predictive histograms are obtained by drawing J samples from the predictive distribution $q_{\theta}(\mathbf{x})$. The latter is very simple and computationally efficient owing to the use of ancestral sampling [25] of the generative model, as explained in the Section 1.2. The right columns of Figures 2 and 3 provide the reference potential energy $U(x_1, x_2 = 0)$, the intermediate target potential $\beta_k U(x_1, x_2 = 0)$, and the predicted potential $U_k^{\text{pred}}(x_1, x_2 = 0)$ after convergence of (ϕ, θ) at temperature β_k . For the intermediate steps, we estimate $U_k^{\text{pred}}(x_1, x_2 = 0)$ as follows:

$$U_k^{\text{pred}}(x_1, x_2 = 0) \propto -\frac{1}{\beta_k} \log q_{\theta}(x_1, x_2 = 0). \quad (38)$$

We note that the evaluation of $\log q_{\theta}(x_1, x_2 = 0)$ requires approximation of the integral $\log q_{\theta}(x_1, x_2 = 0) = \int q(\mathbf{x}|\mathbf{z})q(\mathbf{z}) d\mathbf{z}$, which induces noise. The aforementioned integral has been approximated by $N = 5000$ samples drawn from $q(\mathbf{z})$.


Figure 4(a) shows the overall convergence of the model, expressed in the form of the reverse KL divergence (Equation 7) and the forward KL divergence (Equation 6); the latter, which relies on the data, is only used for illustrative purposes. Data for evaluating $D_{\text{KL}}(p_{\text{target}}(\mathbf{x})||q_{\theta}(\mathbf{x}))$ were not used in the training process. We compare reference statistics (again based on data which were not used during training) with statistics estimated based on the efficient predictive distribution $q_{\theta}(\mathbf{x})$ in Figure 4(b)

2.1.2 Predictive collective variables

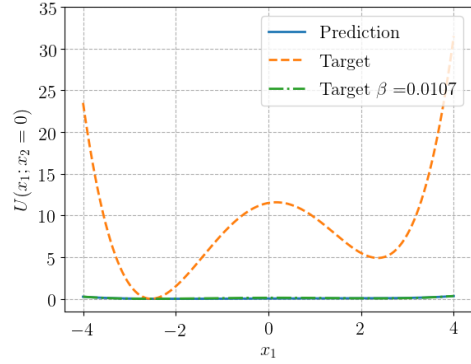
The proposed approach provides an efficient CG model that can be employed for predictive purposes, as described in the previous section. We claim that in addition to obtaining a CG model, we can provide relevant insights by identifying CVs of the system. In the double well example, one would expect the CV to be the x_1 coordinate that separates the two modes, where conformational changes are implied by moving along x_1 .

To visualize the assigned CVs given samples $\mathbf{x}^{(i)} \sim q_{\theta}(\mathbf{x})$, we plot samples as dots in Figure 5, while the color of the $\mathbf{x}^{(i)}$ is assigned based on the corresponding value of the CV. We note that for every $\mathbf{x}^{(i)}$ there exists a whole distribution of CVs $r_{\phi}(z|\mathbf{x}^{(i)})$, as we work in a probabilistic framework. The assigned color in Figure 5 is based on the mean of $r_{\phi}(z|\mathbf{x}^{(i)})$, which is obtained by evaluating $\mu_{\phi}(\mathbf{x}^{(i)})$.


The (color) gradient of z with respect to \mathbf{x} is almost exactly parallel to the x_1 -direction, which implies that the revealed CV z is (probabilistically) parallel to the x_1 axis and thus meets our expectations. The proposed approach reveals the relevant, slow, CV x_1 solely by evaluating $U(\mathbf{x})$ under $q_{\theta}(\mathbf{x})$.




(a) Histogram of $q_\theta(\mathbf{x})$ at $\beta \approx 0$ and of $p_{\text{target}}(\mathbf{x})$.




(b) $U(x_1, x_2 = 0)$ at $\beta \approx 0$.




(c) Histogram of $q_\theta(\mathbf{x})$ at $\beta \approx 0.2$ and of $p_{\text{target}}(\mathbf{x})$.



(d) $U(x_1, x_2 = 0)$ at $\beta \approx 0.2$.




(e) Histogram of $q_\theta(\mathbf{x})$ at $\beta \approx 0.36$ and of $p_{\text{target}}(\mathbf{x})$.




(f) $U(x_1, x_2 = 0)$ at $\beta \approx 0.36$.


Figure 2: The left column shows histograms of the target $p_{\text{target}}(\mathbf{x})$ (at $\beta = 1$) and predictions based on $q_\theta(\mathbf{x})$ at the indicated temperature β in the subcaptions. The right column shows a 1D slice through the potential energy $U(\mathbf{x})$ at $x_2 = 0$, emphasizing the two distinct modes. The figures include the reference potential for the indicated temperature β_k with $\beta_k U(\mathbf{x})$ and an estimation of $U_k^{\text{pred}}(\mathbf{x})$ based on $q_\theta(\mathbf{x})$.




(a) Histogram of $q_{\theta}(\mathbf{x})$ at $\beta \approx 0.7$ and of $p_{\text{target}}(\mathbf{x})$.



(b) $U(x_1, x_2 = 0)$ at $\beta \approx 0.7$.




(c) Histogram of $q_{\theta}(\mathbf{x})$ at $\beta \approx 1$ and of $p_{\text{target}}(\mathbf{x})$.

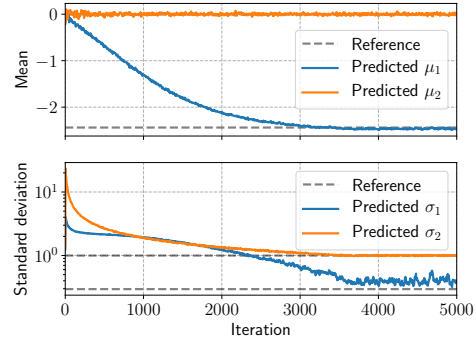


(d) $U(x_1, x_2 = 0)$ at $\beta \approx 1$.

Figure 3: The left column shows histograms of the target $p_{\text{target}}(\mathbf{x})$ (at $\beta = 1$) and predictions based on $q_{\phi}(\mathbf{x})$ at the indicated temperature β in the subcaptions. The right column shows a 1D slice through the potential energy $U(\mathbf{x})$ at $x_2 = 0$, emphasizing the two distinct modes. The figures include the reference potential for the indicated temperature β_k with $\beta_k U(\mathbf{x})$ and an estimation of $U_k^{\text{pred}}(\mathbf{x})$ based on $q_{\theta}(\mathbf{x})$.



(a) Upper and lower bounds of the training objective.



(b) Predicted mean and standard deviations compared with reference data-based estimates. The subscripts of μ and σ indicate the corresponding x_1 and x_2 directions.

Figure 4: Convergence of the KL divergences (left) and predicted statistics compared with reference estimates (right).




Figure 5: Samples $\mathbf{x}^{(i)} \sim q_{\theta}(\mathbf{x})$ at the indicated temperature β are depicted as dots, whereas the assigned color of $\mathbf{x}^{(i)}$ corresponds to its latent CV obtained by the mean of $r_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})$. The color bar below the images shows the color corresponding to the assigned value of the CV \mathbf{z} given \mathbf{x} . The figure is based on $N = 1 \times 10^4$ samples $\mathbf{x}^{(i)}$.




Figure 6: Dihedral angles (left) and (ϕ, ψ) statistics of a reference simulation with characteristic modes (right).

2.2 ALA-2

After demonstrating the functionality of the proposed scheme for a double well potential energy, we are interested in addressing an atomistic system. The following is devoted to the CV discovery of alanine dipeptide (ALA-2) in the context of an implicit solvent. Characteristic coordinates of the ALA-2 peptide include the dihedral angles (ϕ, ψ) , as shown in Figure 6(a). Distinct combinations of the dihedral angles characterize three distinguishable $(\alpha, \beta-1, \beta-2)$ conformations, as provided in the Ramachandran plot [90] in Figure 6(b) [91]. The peptide consisting of 22 atoms can be described by 60 effective degrees of freedom (rigid body motion removed); however, we store the complete Cartesian coordinate vector \mathbf{x} with $\dim(\mathbf{x}) = 66$, where six degrees of freedom are fixed. The exact representation of ALA-2 in \mathbf{x} with coordinate bookkeeping is given in the Appendix C.

2.2.1 Reference model setting

Applying the proposed methodology does not require the production of any reference atomistic trajectories. However, we are interested in comparing our obtained predictions from the generative CG model to reference observables estimated by a reference MD simulation. We refer to Appendix D for all necessary details obtaining the MD trajectory. Nevertheless, for the sake of evaluating forces, we need to specify system properties such as the force field, which in this case is AMBER ff96 [92–94]. We employ an implicit water model based on the generalized Born approach [95, 96], which serves as a solvent. The temperature of interest is $T = 330$ K.

2.2.2 Model specification

The general model structure introduced earlier in Section 1.5.1 is also employed in the context of the ALA-2 setting. We mostly rely on findings in [57], where an identical system was explored on the basis of data-driven forward KL divergence minimization. All required details for the model are specified in Tables 3 and 4 for the encoder $(r_\phi(\mathbf{z}|\mathbf{x}))$ and decoder $(q_\theta(\mathbf{x}|\mathbf{z}))$, respectively. Similar to the previous example in Section 2.1, we employ a tempering scheme as introduced in Section 1.6 and specified in Algorithm 1 with initial $\beta_0 = 1 \times 10^{-14} \cdot \beta_K$, while the target temperature is defined by $\beta_K = \frac{1}{k_B T}$ and $T = 330$ K. The inverse temperature β_0 occurs as a multiplicative factor, multiplying the potential energy $U(\mathbf{x})$. For gradient estimation, the interatomic force $\mathbf{F}(\mathbf{x})$ is multiplied by β_k . In the variational approach presented in this work, we evaluate the force field under samples from $q_\theta(\mathbf{x})$. However, when $q_\theta(\mathbf{x})$ has not yet learned, samples $\mathbf{x}^{(i)}$ will potentially yield high-energy states associated with large forces. According to experimental results, the magnitude of $\mathbf{F}(\mathbf{x})$ in early training stages reaches $\pm 1 \times 10^{-18}$. Therefore, the initial inverse temperature is chosen such that $\beta_0 \mathbf{F}(\mathbf{x})$ evaluated under $q_\theta(\mathbf{x})$ yields values of $\pm 1 \times 10^1$. This implies that the embedded physics,

Linear layer	Input dimension	Output dimension	Activation layer	Activation function
$l_\phi^{(1)}$	$\dim(\mathbf{x}) = 66$	d_1	$a^{(1)}$	SeLu ⁷
$l_\phi^{(2)}$	d_1	d_2	$a^{(2)}$	SeLu
$l_\phi^{(3)}$	d_2	d_3	$a^{(3)}$	Log Sigmoid ⁸
$l_\phi^{(4)}$	d_3	$\dim(\mathbf{z})$	None	-
$l_\phi^{(5)}$	d_3	$\dim(\mathbf{z})$	None	-

Table 3: Network specification of the encoding neural network with $d_{\{1,2,3\}} = 170$.

Linear layer	Input dimension	Output dimension	Activation layer	Activation function
$l_\theta^{(1)}$	$\dim(\mathbf{z}) = 2$	d_3	$\tilde{a}^{(1)}$	Tanh
$l_\theta^{(2)}$	d_3	d_2	$\tilde{a}^{(2)}$	Tanh
$l_\theta^{(3)}$	d_2	d_1	$\tilde{a}^{(3)}$	Tanh
$l_\theta^{(4)}$	d_1	$\dim(\mathbf{x})$	None	-

Table 4: Network specification of the decoding neural network with $d_{\{1,2,3\}} = 120$.

expressed by interatomic forces $\mathbf{F}(\mathbf{x})$, are weak in the early training stages and are emphasized as the learning process proceeds with increasing β_k . For further details, refer to Appendix F.

The stochastic optimization algorithm is ADAM [88] with $\alpha = 0.001, \beta_1 = 0.9, \beta_2 = 0.999, \epsilon_{\text{ADAM}} = 1.0 \times 10^{-8}$. We employ $J = 2000$ samples for computing expectations with respect to $q_\theta(\mathbf{x})$ throughout the training process. Initially, in the early training stages, using fewer samples does not influence the training. The number of samples should be increased once the model has been refined and comes closer to $p_{\text{target}}(\mathbf{x})$.

2.2.3 Collective variables

When training the model with its encoder and decoder components $r_\phi(\mathbf{z}|\mathbf{x})$ and $q_\theta(\mathbf{x}|\mathbf{z})$, it is important that these consistently map a generated sample $\mathbf{z}^{(i)} \sim q(\mathbf{z})$ to $\mathbf{x}^{(i)}$ through the decoder $q_\theta(\mathbf{x}|\mathbf{z})$, and from the decoded atomistic configuration $\mathbf{x}^{(i)}$ back to its origin, the value of the CV $\mathbf{z}^{(i)}$ it has been generated from. The projection from $\mathbf{x}^{(i)}$ to the according CV is enabled through the encoder $r_\phi(\mathbf{z}|\mathbf{x})$. After some initial iterations optimizing (ϕ, θ) , the encoder and decoder work consistently as depicted in Figure 7.

In Figure 8 we utilize the identified encoder $r_\phi(\mathbf{z}|\mathbf{x})$, which assigns CVs to an input atomistic configuration, for encoding a reference test dataset. This dataset has not been used for training and is used here solely for visualization purposes. The test data (generated according to Appendix D) contains atomistic configurations from multiple characteristic modes based on their dihedral angle values (ϕ, ψ) as shown in Figure 6(b). Given a datum from the test dataset $\mathbf{x}^{(i)}$, we can assign the corresponding value of its CV by employing the mean $\mu_\phi(\mathbf{x}^{(i)})$ of the approximate posterior over the latent variables $r_\phi(\mathbf{z}|\mathbf{x})$. The assigned CV depicts the pre-image of the atomistic configuration in the reduced CV space. It is important to note in Figure 8 that atomistic configurations belonging to characteristic conformations $(\alpha, \beta-1, \beta-2)$ are identified by $r_\phi(\mathbf{z}|\mathbf{x})$ and form clusters in the CV space. We note that the conformations $\beta-1, \beta-2$ interleave with each other in regions around $z_1 = 0$. An explanation for this overlap is the similarity of (ϕ, ψ) combinations in the Ramachandran plot in Figure 6(b). Separate from the β configurations is the cluster associated with α configurations in the CV space. The latter differ significantly with respect to the (ϕ, ψ) pairs from the β conformations.

The implied similarity of, e.g., β conformations in the CV space is in accordance with the expectations on dimensionality reduction methods. Similar atomistic—or, in general, observed—coordinates should map to similar regions in their latent lower-dimensional embedding, as emphasized in [5]. This is achieved in, e.g., multidimensional scaling [97] or isomap [10]. The presented dimensionality




Figure 7: Samples from $q_\theta(\mathbf{z})$ (blue, filled) are decoded with $q_\theta(\mathbf{x}|\mathbf{z})$ and encoded with $r_\phi(\mathbf{z}|\mathbf{x})$ (orange, no facecolor). We consider the means of aforementioned distributions for performing the decoding and encoding processes. Re-encoding the decoded $\mathbf{z}^{(i)}$ matches its origin.




Figure 8: Representation of \mathbf{z} -coordinates of test data assigned by the mean of $r_\phi(\mathbf{z}|\mathbf{z})$, which we learn by minimizing the reverse KL divergence without reference data. Characteristic conformations of ALA-2 are indicated in: α black, $\beta-1$ blue, and $\beta-2$ red color. Without any prior physical information and in the complete absence of any data, the encoder identifies physically relevant coordinates, which are related to ϕ, ψ values.




Figure 9: *Predicted* configurations \mathbf{x} (including dihedral angle values) with $\mu_{\theta}(\mathbf{z})$ of $p_{\theta}(\mathbf{x}|\mathbf{z})$. As one moves along the z_1 axis, we obtain for the given CVs atomistic configurations \mathbf{x} reflecting the conformations α , β -1, and β -2. All rendered atomistic representations in this work were created by VMD [98].

reduction relies solely on evaluating the force field $\mathbf{F}(\mathbf{x})$ at generated samples from $q_{\theta}(\mathbf{x})$, without using any data, and is differentiable with respect to \mathbf{x} .

The hidden and lower-dimensional physically characteristic generative process is emphasized further in Figure 9. We illustrate *predicted* atomistic configurations \mathbf{x} given the marked (circle) values of the CVs \mathbf{z} . The change of characteristic (ϕ, ψ) dihedrals can be observed by moving from the red (β -2) to the blue region (β -1) in the CV space and observing the configurational change in the predicted atomistic configurations, given the indicated CVs. The depicted atomistic configurations are obtained using the input CV \mathbf{z} and the mean of $q_{\theta}(\mathbf{x}|\mathbf{z})$, which is expressed as a neural network with $\mu_{\theta}(\mathbf{z})$. The probabilistic decoder $q_{\theta}(\mathbf{x}|\mathbf{z})$ is a distribution, implying that given one value of the CV, several atomistic realizations can be produced. For illustrative reasons we represent the mean $\mu_{\theta}(\mathbf{z})$.

To obtain a better understanding of the meaning of identified CVs in terms of the dihedral angles (ϕ, ψ) , we visualize them by mapping given values of \mathbf{z} to atomistic configurations and compute the (ϕ, ψ) values assigned to the corresponding \mathbf{z} , as shown in Figure 10. Again, in the probabilistic model, we draw multiple atomistic realizations \mathbf{x} given one CG representation \mathbf{z} . The realization for a given \mathbf{z} fluctuates in terms of bonded vibrations rather than any change in the dihedrals (ϕ, ψ) . We observe a strong correlation between (ϕ, ψ) and the CVs \mathbf{z} .

In addition to the visual assessment given in Figure 9, we show quantitatively that the structural properties of atomistic configurations generated through $q_{\theta}(\mathbf{x})$ truly capture those of a reference trajectory at $T = 330$ K, as shown in Figure 11. Figure 11 provides histograms over bonding distances over all bonded atoms in the system. Reference statistics of bond lengths are compared with those based on generated samples of the predictive distribution $q_{\theta}(\mathbf{x})$. Figure 12 provides estimated observables based on the predictive model and a reference trajectory. The observables are computed as explained in Appendix E.




Figure 10: Predicted dihedrals (ϕ, ψ) for the latent CVs. The depicted (ϕ, ψ) values were obtained from atomistic configurations given a CV value \mathbf{z} through the mean of $q_\theta(\mathbf{x}|\mathbf{z})$, $\mu_\theta(\mathbf{z})$.

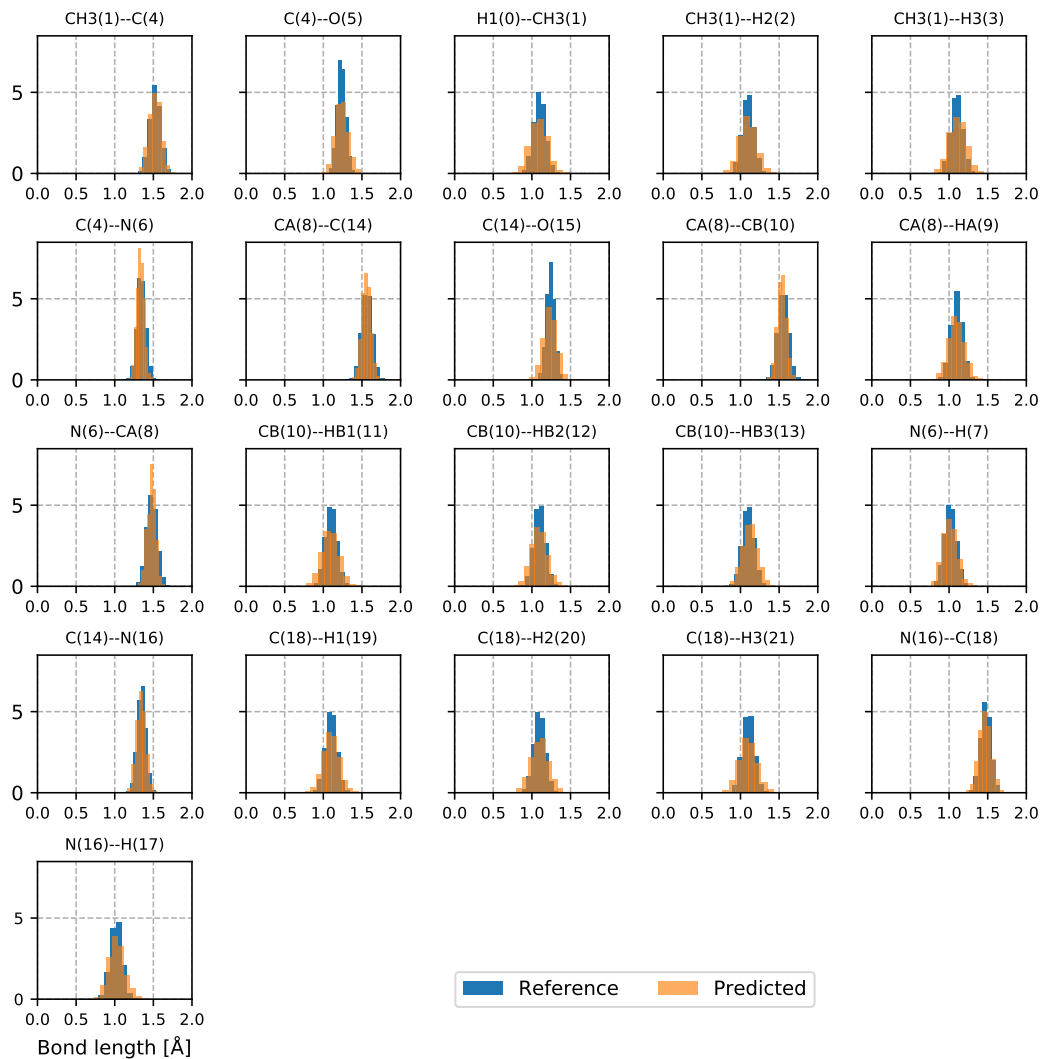


Figure 11: Bonding distance statistics. In ALA-2, bonded atoms of a reference simulation (blue) compared with histograms of the bond lengths of the *predicted* atomistic ensemble based on $q_{\theta}(\mathbf{x})$ (semi-transparent in the foreground in orange). The titles of the subplots indicate the relevant atom names, and the corresponding atom id of the structure file of ALA-2 as provided in Appendix C is shown in brackets. The physics, in the form of bonding distances, is well maintained in the generated realizations. Predictive estimates are obtained by employing $J = 2000$ samples of $q_{\theta}(\mathbf{x})$, and the reference is based on $N = 4000$ MD snapshots.




Figure 12: Predicted observables compared with reference estimates. Radius of gyration (left) and root-mean-squared deviation (right). Predictive estimates are obtained by employing $J = 2000$ samples of $q_\theta(\mathbf{x})$ and the reference by $N = 4000$ MD snapshots. Observables are estimated according to Appendix E.

3 Summary and outlook

We have presented a variational approach for revealing CVs of atomistic systems that additionally yield a predictive CG model. We circumvent the need for reference data, which is supposed to provide an approximation of the target distribution $p_{\text{target}}(\mathbf{x})$. The simulation of complex biochemical systems and thus the obtained data may suffer from bias owing to insufficiently exploring all relevant conformations. Conformations are separated by high free-energy barriers, which hamper efficient exploration with brute-force MD [6]. The presented variational coarse-graining and CV discovery approach is guided by evaluating interatomic forces under the predictive distribution $q_{\theta}(\mathbf{x})$, where sampling is computationally efficient. By embedding the atomistic force components, $q_{\theta}(\mathbf{x})$ learns from the target distribution $p_{\text{target}}(\mathbf{x})$. We derived an upper bound on the reverse KL divergence, in which all terms are tractable, and discuss the physical underpinning of the components involved. The derived upper bound is subject to minimization with respect to all model parameters. We provide a variance-reducing gradient estimator based on reparametrization. Whereas variational approaches are known for being mode focusing, remedy provides the introduced consistent tempering scheme, alleviating the simultaneous learning of modes. We demonstrate the proposed algorithmic advances with a double well potential energy and the ALA-2 peptide. Characteristic CVs have been identified by the introduced optimization objective.

The following steps will be pursued in continuation of this work. Atomistic forces and thus gradients span many orders of magnitude at initial iterations. This could lead to numerical instabilities. Thus, we are interested in synthesizing the advantages of the forward and reverse KL divergences in the context of atomistic systems. We propose an adaptive learning scheme that may rely in its early training stages on a few data points. These are not required to reflect the whole phase space, but it is important to have a basis for learning, e.g., the structure of the atomistic system with its approximate bond lengths. This eases the problem of evaluating $\mathbf{F}(\mathbf{x})$ for the un-physical realizations that may be predicted by $q_{\theta}(\mathbf{x})$ in an early training phase:

$$\mathcal{E} = \gamma D_{\text{KL}}(q_{\theta}(\mathbf{x}) \| p_{\text{target}}(\mathbf{x})) + (1 - \gamma) D_{\text{KL}}(p_{\text{target}}(\mathbf{x}) \| q_{\theta}(\mathbf{x})). \quad (39)$$

With $\gamma \in [0, 1]$ weighting the overall contribution from the reverse and forward KL divergences, we use an adaptive weight, $\gamma(k)$, which implies dependence on the current iteration k . With the proceeding learning process (increasing k), γ could increase up to $\gamma = 1$, so that we fully rely on the variational approach and thus the associated physics expressed by the potential $U(\mathbf{x})$ and forces $\mathbf{F}(\mathbf{x})$. Minimizing the above objective Equation 39 with respect to model distributions synthesizes the findings of this work and those of [57].

We furthermore propose the employment of the obtained $q_{\theta}(\mathbf{x})$ for predictive purposes of systems at different temperatures. This can be achieved by obtaining the implicitly learned predictive potential expressed in terms of fine-scale coordinates at β_{target} based on $q_{\theta}(\mathbf{x})$:

$$U_{\theta}^{\text{pred}}(\mathbf{x}) = -\frac{1}{\beta_{\text{target}}} \log q_{\theta}(\mathbf{x}) + \text{const}. \quad (40)$$

Assuming we are interested in simulating the same system at β_{new} where $\beta_{\text{new}} \neq \beta_{\text{target}}$, we can readily provide a generalized predictive distribution for any β_{new} ,

$$\tilde{q}_{\theta}(\mathbf{x}) \propto e^{-\beta_{\text{new}} U_{\theta}^{\text{pred}}(\mathbf{x})}, \quad (41)$$

by employing the predictive potential $U_{\theta}^{\text{pred}}(\mathbf{x})$ defined in Equation 40.

Finally, we emphasize the most relevant findings of this work. We have reformulated the identification of CVs as an optimization problem, which additionally provides a predictive CG model. CVs are revealed in the absence of any prior physical knowledge or insight, and thus in the absence of any system-dependent assumptions. Instead of relying on reference data, we employ the minimization of the reverse KL divergence and develop an inference scheme in the context of atomistic systems. Thus, the optimization is solely guided by the *evaluation* of the potential $U(\mathbf{x})$ and/or forces $\mathbf{F}(\mathbf{x})$ at samples of the predictive distributions $q_{\theta}(\mathbf{x})$. We have also developed an adaptive tempering scheme based on findings of [29].

A Relation with Expectation-Propagation

This section emphasizes the relationship of hierarchical variational models with expectation-propagation (EP) [99].

The following is not directly relevant to optimization of the objective Equation (7), but it shows the existence of an upper bound of the entropy term $-\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x})]$. Similar to Equation (9) one denotes,

$$\begin{aligned}
-\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x})] &= -\mathbb{E}_{q(\mathbf{x})} [\log q(\mathbf{x}) - D_{KL}(q(\mathbf{z}|\mathbf{x})||q(\mathbf{z}|\mathbf{x}))] \\
&\leq \mathbb{E}_{q(\mathbf{x})} [-\log q(\mathbf{x}) + D_{KL}(r(\mathbf{z}|\mathbf{x})||q(\mathbf{z}|\mathbf{x}))] \\
&= \mathbb{E}_{q(\mathbf{x})} [\mathbb{E}_{r(\mathbf{z}|\mathbf{x})} [-\log q(\mathbf{x}) - \log q(\mathbf{z}|\mathbf{x}) + \log r(\mathbf{z}|\mathbf{x})]] \\
&= \mathbb{E}_{q(\mathbf{x})} \left[\mathbb{E}_{r(\mathbf{z}|\mathbf{x})} \left[-\log q(\mathbf{x}) - \log \frac{q(\mathbf{x}|\mathbf{z})q(\mathbf{z})}{q(\mathbf{x})} + \log r(\mathbf{z}|\mathbf{x}) \right] \right] \\
&= \mathbb{E}_{q(\mathbf{x})} [\mathbb{E}_{r(\mathbf{z}|\mathbf{x})} [-\log q(\mathbf{x}|\mathbf{z}) - \log q(\mathbf{z}) + \log r(\mathbf{z}|\mathbf{x})]] . \tag{42}
\end{aligned}$$

The bound in Equation (42) is tractable if sampling from $q(\mathbf{x})$ and $r(\mathbf{z}|\mathbf{x})$ is feasible. Both bounds (Eqs. 42 and 9) show similarities to the derivation of EP [99] and variational Bayesian inference [100]. However, note that the lower bound in Equation (11) is connected to the objective in EP, although EP only minimizes $D_{KL}(q||r)$ locally. The bound derived with $q(\mathbf{x})$ results in a tighter bound compared with variational autoencoders with $q(\mathbf{x}|\mathbf{z})$, as $\mathbb{H}[q(\mathbf{x})] \geq \mathbb{H}[q(\mathbf{x}|\mathbf{z})]$ (for details, see [101]).

B Estimating the relative increase of the KL divergence

The relative increase of the KL divergence induced by decreasing the temperature is denoted as in Equation (36), with

$$c_k = \frac{\log(Z(\beta_{k+1})) - \log(Z(\beta_k)) + (\beta_{k+1} - \beta_k) \langle U(\mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})}}{\log Z(\beta_k) + \beta_k \langle U(\mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})} - \langle \log r(\mathbf{z}|\mathbf{x}) \rangle_{q(\mathbf{x}, \mathbf{z})} - \mathbb{H}(q(\mathbf{x}, \mathbf{z}))}.$$

The following addresses the estimation of $\log(Z(\beta_{k+1})) - \log(Z(\beta_k))$ with $\Delta\beta_k = \beta_{k+1} - \beta_k$:

$$\begin{aligned}
Z(\beta_k + \Delta\beta_k) &= \int e^{-(\beta_k + \Delta\beta_k)U(\mathbf{x})} d\mathbf{x} \tag{43} \\
&= \int \frac{e^{-(\beta_k + \Delta\beta_k)U(\mathbf{x})} e^{-\beta_k U(\mathbf{x})}}{\frac{e^{-\beta_k U(\mathbf{x})}}{Z(\beta_k)}} d\mathbf{x} \\
&= Z(\beta_k) \int e^{-\Delta\beta_k U(\mathbf{x})} p_{\text{target}}(\mathbf{x}; \beta_k) d\mathbf{x} \\
&= Z(\beta_k) \int e^{-\Delta\beta_k U(\mathbf{x})} \frac{e^{-\beta_k U(\mathbf{x})} r(\mathbf{z}|\mathbf{x})}{q(\mathbf{x}|\mathbf{z}) q(\mathbf{z})} q(\mathbf{x}, \mathbf{z}) d\mathbf{x} d\mathbf{z}.
\end{aligned}$$

We are interested in $\log(Z(\beta_{i+k})) - \log(Z(\beta_k))$. Therefore, we write:

$$\begin{aligned}
\log(Z(\beta_{i+1})) - \log(Z(\beta_i)) &= \log \int e^{-\Delta\beta U(\mathbf{x})} \underbrace{\frac{e^{-\beta_i U(\mathbf{x})} r(\mathbf{z}|\mathbf{x})}{q(\mathbf{x}|\mathbf{z}) q(\mathbf{z})}}_w q(\mathbf{x}, \mathbf{z}) d\mathbf{x} d\mathbf{z} \tag{44} \\
&\approx \log \sum_{i=1}^N e^{-\Delta\beta U(\mathbf{x}^{(i)})} W^{(i)}.
\end{aligned}$$

Equation (44) depicts a noisy Monte Carlo estimator for $\log(Z(\beta_{k+1})) - \log(Z(\beta_k))$ based on importance sampling [102] with the following normalized weights:

$$W^{(i)} = \frac{w^{(i)}}{\sum w^{(i)}} \quad \text{with} \quad w^{(i)} \propto \frac{e^{-\beta_i U(\mathbf{x})} r(\mathbf{z}|\mathbf{x})}{q(\mathbf{x}|\mathbf{z}) q(\mathbf{z})}. \tag{45}$$

As $e^{-\beta_i U(\mathbf{x})} r(\mathbf{z}|\mathbf{x})$ may be small for samples $(\mathbf{x}^{(i)}, \mathbf{z}^{(i)}) \sim q(\mathbf{x}, \mathbf{z})$, we use instead $\log \bar{w}^{(i)}$ with $\log \bar{w}^{(i)} = \log w^{(i)} - a$ and $a = \max\{\log w^{(i)}\}$ to avoid numerical issues.

Whereas above we showed an approximate estimator for $\log(Z(\beta_{k+1})) - \log(Z(\beta_k))$, the following addresses $\log(Z(\beta_k))$. To estimate the relative increase in the KL divergence, one requires the normalization constant as mentioned in Equation (36). Multistage sampling [103] provides a way to approximate $Z(\beta_i)$, given all previous $Z(\beta_k)$ with $k < i$ and $\beta_i > \beta_{i-1}$:

$$\frac{Z(\beta_i)}{Z(0)} = \frac{Z(\beta_1)}{Z(\beta_0)} \cdot \frac{Z(\beta_2)}{Z(\beta_1)} \dots \frac{Z(\beta_{i-1})}{Z(\beta_{i-2})}. \quad (46)$$

The ratios $\frac{Z(\beta_{k-1})}{Z(\beta_{k-2})}$ are given by Equation (44). The remaining component to be estimated is $Z(0)$, as we utilize the expression from Equation (44) to estimate the ratios of normalization factors. To avoid learning parametrizations θ yielding almost uniform $q(\mathbf{x})$ on an infinite domain, which occurs in the limit when $\beta = 0$, we define the following auxiliary potential to restrict the domain:

$$U_{aux}(\mathbf{x}) = \begin{cases} U(\mathbf{x}), & \text{if } \mathbf{x} \in [-b, b]^{\dim(\mathbf{x})} \\ -\frac{u}{\beta} \mathbf{x}, & \mathbf{x} < -b \\ \frac{u}{\beta} \mathbf{x}, & \mathbf{x} > b, \end{cases} \quad (47)$$

with $u = 10 \times 10^2$. The above extension does not influence the potential energy $U(\mathbf{x})$ at relevant temperatures.

The initial $Z(\beta_0)$ is computed with importance sampling. This is done only once upon convergence of (θ, ϕ) for β_0 :

$$\begin{aligned} Z(\beta_0) &= \int e^{-\beta_0 U(\mathbf{x})} d\mathbf{x} \\ &= \int e^{-\beta_0 U(\mathbf{x})} r(\mathbf{z}|\mathbf{x}) d\mathbf{x} d\mathbf{z} \\ &= \int \underbrace{\frac{e^{-\beta_0 U(\mathbf{x})} r(\mathbf{z}|\mathbf{x})}{q(\mathbf{x}, \mathbf{z})}}_w q(\mathbf{x}, \mathbf{z}) d\mathbf{x} d\mathbf{z}. \end{aligned} \quad (48)$$

With samples $(\mathbf{x}^{(i)}, \mathbf{z}^{(i)}) \sim q(\mathbf{x}, \mathbf{z})$, we obtain the following unnormalized weights:

$$w^{(i)} = \frac{e^{-\beta_0 U(\mathbf{x}^{(i)})} r(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})}{q(\mathbf{x}^{(i)}, \mathbf{z}^{(i)})}, \quad (49)$$

or $\log w^{(i)} = -\beta_0 U(\mathbf{x}^{(i)}) + \log r(\mathbf{z}^{(i)}|\mathbf{x}^{(i)}) - \log q(\mathbf{x}^{(i)}, \mathbf{z}^{(i)})$. Then,

$$\log Z(\beta_0) = -\log N + \log \sum_{i=1}^N e^{\log w^{(i)} - c} + c, \quad (50)$$

with $c = \max(\log w^{(i)})$.

C ALA-2 coordinate representation

We show the structure of the ALA-2 peptide in Figure 13. The numbers in the circles, which depict the involved atoms of ALA-2, correspond to the order in which we assemble block-wise the Cartesian coordinates (x_i, y_i, z_i) of atom i to

$$\mathbf{x} = (x_1, y_1, z_1, x_2, y_2, \dots, x_{22}, y_{22}, z_{22})^T,$$

where i is the atom number as depicted in Figure 13. For removing rigid-body motion, we fix the Cartesian coordinates (x_6, y_6, z_6) of atom 6, (x_9, y_9) of atom 9, and (y_{15}) of atom 15. The employed PDB structure file is available online at https://github.com/cics-nd/predictive-cvs/blob/master/data_peptide/ala-2/ala2_adopted.pdb.




Figure 13: ALA-2 structure with numbered atoms as used for decomposing \mathbf{x} .

D Simulation of ALA-2

The procedure for generating a reference trajectory for computing reference observables of ALA-2 is identical to that employed in [57], whereas the data generation approach relies on [104]. The utilized interaction force field is AMBER ff96 [92–94], resolved by an implicit water model based on the generalized Born model [95, 96]. Incorporating an explicit water model would, obviously, provide trajectories that would yield observables closer to the *experimental* reference. An Andersen thermostat is used to maintain fluctuations around the desired temperature $T = 330$ K. All reference simulations are carried out using Gromacs [71–77]. The time step is $\Delta t = 1$ fs, with a preceding equilibration phase of 50 ns. Thereafter, a trajectory snapshot is taken every 10 ps. Rigid-body motions have been removed from the Cartesian coordinates.

E Observable estimation for ALA-2

We are interested in estimating observables based on predictive models, in contrast to those obtained through reference MD simulations. In general, observables are evaluated as ensemble (MC) or phase (MD) averages, $\int a(\mathbf{x})p_{\text{target}}(\mathbf{x}) d\mathbf{x}$, by making use of $q_{\theta}(\mathbf{x})$ and samples drawn by ancestral sampling. We illustrate the radius of gyration (Rg) [104, 105], given as:

$$a_{\text{Rg}}(\mathbf{x}) = \sqrt{\frac{\sum_p m_p \|\mathbf{x}_p - \mathbf{x}_{\text{COM}}\|^2}{\sum_p m_p}}. \quad (51)$$

The sum in Equation (51) considers all system atoms $p = 1, \dots, P$, with the atom mass m_p and Cartesian coordinate \mathbf{x}_p of each atom. The center of mass of the peptide is denoted by \mathbf{x}_{COM} . A histogram of $a_{\text{Rg}}(\mathbf{x})$ reflects the statistics of the peptide’s average size, which characterize its various conformations [105].

F Gradient normalization

During optimization of the objective in the context of atomistic systems, we encounter significant forces, $\mathbf{F}(\mathbf{x})$. These differ in magnitude owing to sampling atomistic realizations, which induce, e.g., relatively small distances between bonded atoms. This leads to extreme force components. Gradient normalization [106, 107] circumvents disruption of the current set of learned parameters (θ, ϕ) via a single component attached with extreme magnitudes, owing to, e.g., short bonded distances. Once training proceeds, and predicted atomistic realizations are closer to reasonable ones, the gradient normalization becomes redundant, affecting only gradients in extreme settings where the absolute values of $\mathbf{F}(\mathbf{x}) \geq 1 \times 10^{15}$. After an initial learning phase, such extreme magnitudes do not occur, and thus the normalization does not affect or distort the physics induced by evaluating the force field $\mathbf{F}(\mathbf{x})$.

Given a batch of I samples $\{\mathbf{x}^{(i)}\}_{i=1}^I$ obtained from $q_{\theta}(\mathbf{x})$, we estimate the gradient of the objective, $\mathbf{g}^i(\mathbf{x}^{(i)})$ and calculate its ℓ -2 norm:

$$l^i = \|\mathbf{g}^i\|_2. \quad (52)$$

The average gradient norm is $\bar{l} = 1/I \sum_{i=1}^I l^i$, and we allow a maximal gradient norm based on the mean with $l_{\max} = \kappa \bar{l}$, $\kappa = 3.0$. κ was determined by an empirical study. Those gradients with $l^i > l_{\max}$ are normalized such that

$$\mathbf{g}_n^i = \frac{l_{\max}}{l^i} \mathbf{g}^i. \quad (53)$$

As mentioned earlier, realistic atomistic systems at relevant temperatures are not exposed to ℓ -2 norms of gradients differing more as twice as compared with the gradient with the lowest ℓ -2 norm. Thus, the gradient normalization is inactive when learning realistic atomistic configurations.

References

- [1] B. J. Alder and T. E. Wainwright, *The Journal of Chemical Physics* **31**, 459 (1959).
- [2] N. Metropolis and S. Ulam, *J. Am. Stat. Assoc.* **44**, 335 (1949).
- [3] J. Weiner, *Statistical mechanics of elasticity; 2nd ed.* Dover books on physics (Dover Publications, New York, NY, 2012).
- [4] G. G. Maisuradze, A. Liwo, and H. A. Scheraga, *Journal of chemical theory and computation* **6**, PMC3633568[pmcid], 583 (2010).
- [5] M. A. Rohrdanz, W. Zheng, and C. Clementi, *Annual Review of Physical Chemistry* **64**, PMID: 23298245, 295 (2013).
- [6] F. Sittel and G. Stock, *The Journal of Chemical Physics* **149**, 150901 (2018).
- [7] M. Balasubramanian and E. L. Schwartz, *Science* **295**, 7 (2002).
- [8] D. L. Donoho and C. Grimes, *Proceedings of the National Academy of Sciences* **100**, 5591 (2003).
- [9] H. Risken and T. Frank, *The fokker-planck equation: methods of solution and applications (springer series in synergetics)* (Springer, 1996).
- [10] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, *Science* **290**, 2319 (2000).
- [11] M. Ceriotti, G. A. Tribello, and M. Parrinello, *Proceedings of the National Academy of Sciences* **108**, 13023 (2011).
- [12] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, *The Journal of Chemical Physics* **134**, 124116 (2011).
- [13] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, *Proceedings of the National Academy of Sciences* **102**, 7426 (2005).
- [14] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, *Proceedings of the National Academy of Sciences* **102**, 7432 (2005).
- [15] A. L. Ferguson, A. Z. Panagiotopoulos, P. G. Debenedetti, and I. G. Kevrekidis, *The Journal of Chemical Physics* **134**, 135103 (2011).
- [16] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis, *Applied and Computational Harmonic Analysis* **21**, Special Issue: Diffusion Maps and Wavelets, 113 (2006).
- [17] R. R. Coifman, I. G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, *Multiscale Modeling & Simulation* **7**, 842 (2008).
- [18] F. Noé and F. Nüske, *Multiscale Modeling & Simulation* **11**, 635 (2013).
- [19] F. Noé, R. Banisch, and C. Clementi, *Journal of Chemical Theory and Computation* **12**, PMID: 27696838, 5620 (2016).
- [20] J. McCarty and M. Parrinello, *The Journal of Chemical Physics* **147**, 204109 (2017).
- [21] W. Zheng, A. V. Vargiu, M. A. Rohrdanz, P. Carloni, and C. Clementi, *The Journal of Chemical Physics* **139**, 145102 (2013).
- [22] F. Noé, S. Olsson, J. Köhler, and H. Wu, *Science* **365** (2019).
- [23] L. Dinh, J. Sohl-Dickstein, and S. Bengio, in 5th international conference on learning representations, ICLR 2017, toulon, france, april 24-26, 2017, conference track proceedings (2017).

- [24] L. Dinh, D. Krueger, and Y. Bengio, in 3rd international conference on learning representations, ICLR 2015, san diego, ca, usa, may 7-9, 2015, workshop track proceedings (2015).
- [25] J. Liu, *Monte carlo strategies in scientific computing*, edited by J. Liu (Springer Verlag, New York, Berlin, Heidelberg, 2008), p. 344.
- [26] R. Ranganath, D. Tran, and D. Blei, in Proceedings of the 33rd international conference on machine learning, Vol. 48, edited by M. F. Balcan and K. Q. Weinberger, Proceedings of Machine Learning Research (20–22 Jun 2016), pp. 324–333.
- [27] A. Lecchini-visintini, J. Lygeros, and J. Maciejowski, in *Advances in neural information processing systems 20*, edited by J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis (Curran Associates, Inc., 2008), pp. 865–872.
- [28] Y. Li, V. A. Protopopescu, N. Arnold, X. Zhang, and A. Gorin, *Appl. Math. Comput.* **212**, 216 (2009).
- [29] I. Bilionis and P. Koutsourelakis, *Journal of Computational Physics* **231**, 3849 (2012).
- [30] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *The Journal of Chemical Physics* **21**, 1087 (1953).
- [31] U. Grenander and M. I. Miller, *Journal of the Royal Statistical Society. Series B (Methodological)* **56**, 549 (1994).
- [32] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids: second edition* (Oxford University Press, Oxford, 2017), p. 640.
- [33] D. Frenkel and B. Smit, eds., *Understanding molecular simulation: from algorithms to applications*, 1st (Academic Press, Inc., Orlando, FL, USA, 1996).
- [34] M. E. Tuckerman and G. J. Martyna, *The Journal of Physical Chemistry B* **104**, 159 (2000).
- [35] D. E. Shaw, J. P. Grossman, J. A. Bank, B. Batson, J. A. Butts, J. C. Chao, M. M. Deneroff, R. O. Dror, A. Even, C. H. Fenton, A. Forte, J. Gagliardo, G. Gill, B. Greskamp, C. R. Ho, D. J. Ierardi, L. Iserovich, J. S. Kuskin, R. H. Larson, T. Layman, L. Lee, A. K. Lerer, C. Li, D. Killebrew, K. M. Mackenzie, S. Y. Mok, M. A. Moraes, R. Mueller, L. J. Nociolo, J. L. Peticolas, T. Quan, D. Ramot, J. K. Salmon, D. P. Scarpazza, U. B. Schafer, N. Siddique, C. W. Snyder, J. Spengler, P. T. P. Tang, M. Theobald, H. Toma, B. Towles, B. Vitale, S. C. Wang, and C. Young, in *Sc '14: proceedings of the international conference for high performance computing, networking, storage and analysis* (Nov. 2014), pp. 41–53.
- [36] R. C. Bernardi, M. C. Melo, and K. Schulten, *Biochimica et Biophysica Acta (BBA) - General Subjects* **1850**, Recent developments of molecular dynamics, 872 (2015).
- [37] M. KARPLUS and J. A. MCCAMMON, *Nature* **277**, 578 (1979).
- [38] P. L. Freddolino, C. B. Harrison, Y. Liu, and K. Schulten, *Nature physics* **6**, PMC3032381[pmcid], 751 (2010).
- [39] D. L. Ensign and V. S. Pande, *Biophysical journal* **96**, S0006-3495(09)00495-0[PII], L53 (2009).
- [40] Y.-L. Chen and M. Habeck, *PLOS ONE* **12**, 1 (2017).
- [41] A. Laio and M. Parrinello, *Proceedings of the National Academy of Sciences* **99**, 12562 (2002).
- [42] M. A. Rohrdanz, W. Zheng, and C. Clementi, *Annual Review of Physical Chemistry* **64**, PMID: 23298245, 295 (2013).
- [43] R. T. McGibbon, B. E. Husic, and V. S. Pande, *The Journal of Chemical Physics* **146**, 044109 (2017).
- [44] M. Schöberl, N. Zabarar, and P.-S. Koutsourelakis, *Journal of Computational Physics* **333**, 49 (2017).
- [45] H. Grubmüller, *Phys. Rev. E* **52**, 2893 (1995).
- [46] A. Cichocki and S.-i. Amari, *Entropy* **12**, 1532 (2010).
- [47] S.-H. Cha, in, Vol. 1 (2007).
- [48] J. Hernandez-Lobato, Y. Li, M. Rowland, T. Bui, D. Hernandez-Lobato, and R. Turner, in Proceedings of the 33rd international conference on machine learning, Vol. 48, edited by M. F. Balcan and K. Q. Weinberger, Proceedings of Machine Learning Research (20–22 Jun 2016), pp. 1511–1520.
- [49] P. Koutsourelakis, N. Zabarar, and M. Girolami, *Journal of Computational Physics* **321**, 1252 (2016).

- [50] D. Rezende and S. Mohamed, in Proceedings of the 32nd international conference on machine learning, Vol. 37, edited by F. Bach and D. Blei, Proceedings of Machine Learning Research (July 2015), pp. 1530–1538.
- [51] D. P. Kingma and M. Welling, in 2nd international conference on learning representations, ICLR 2014, banff, ab, canada, april 14-16, 2014, conference track proceedings (2014).
- [52] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, *J. Mach. Learn. Res.* **14**, 1303 (2013).
- [53] R. Ranganath, S. Gerrish, and D. Blei, in Proceedings of the seventeenth international conference on artificial intelligence and statistics, Vol. 33, edited by S. Kaski and J. Corander, Proceedings of Machine Learning Research (22–25 Apr 2014), pp. 814–822.
- [54] C. X. Hernández, H. K. Wayment-Steele, M. M. Sultan, B. E. Husic, and V. S. Pande, in (2017).
- [55] C. Wehmeyer and F. Noé, *The Journal of Chemical Physics* **148**, 241703 (2018).
- [56] M. M. Sultan, H. K. Wayment-Steele, and V. S. Pande, *Journal of Chemical Theory and Computation* **14**, PMID: 29529369, 1887 (2018).
- [57] M. Schöberl, N. Zabarar, and P.-S. Koutsourelakis, *The Journal of Chemical Physics* **150**, 024109 (2019).
- [58] M. J. Beal, “Variational algorithms for approximate bayesian inference”, PhD thesis (Gatsby Computational Neuroscience Unit, University College London, 2003).
- [59] A. Amadei, A. B. M. Linssen, and H. J. C. Berendsen, *Proteins: Structure, Function, and Bioinformatics* **17**, 412 (1993).
- [60] A. L. Ferguson, A. Z. Panagiotopoulos, I. G. Kevrekidis, and P. G. Debenedetti, *Chemical Physics Letters* **509**, 1 (2011).
- [61] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, in , edited by D. E. Rumelhart, J. L. McClelland, and C. PDP Research Group (MIT Press, Cambridge, MA, USA, 1986) Chap. Learning Internal Representations by Error Propagation, pp. 318–362.
- [62] C. Van Der Malsburg, in *Brain theory*, edited by G. Palm and A. Aertsen (1986), pp. 245–248.
- [63] S. Haykin, *Neural networks: a comprehensive foundation*, 2nd (Prentice Hall PTR, Upper Saddle River, NJ, USA, 1998).
- [64] P.-A. Mattei and J. Frellsen, in *Advances in neural information processing systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018), pp. 3859–3870.
- [65] J. C. Pinheiro and D. M. Bates, *Statistics and Computing* **6**, 289 (1996).
- [66] D. Kingma and M. Welling, in Proceedings of the 31st international conference on machine learning, Vol. 32, edited by E. P. Xing and T. Jebara, Proceedings of Machine Learning Research 2 (22–24 Jun 2014), pp. 1782–1790.
- [67] D. J. Rezende, S. Mohamed, and D. Wierstra, in Proceedings of the 31st international conference on machine learning, Vol. 32, edited by E. P. Xing and T. Jebara, Proceedings of Machine Learning Research 2 (22–24 Jun 2014), pp. 1278–1286.
- [68] F. R. Ruiz, M. Titsias RC AUEB, and D. Blei, in *Advances in neural information processing systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), pp. 460–468.
- [69] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Nature* **323**, 533 (1986).
- [70] S. Plimpton, *Journal of Computational Physics* **117**, 1 (1995).
- [71] H. Berendsen, D. van der Spoel, and R. van Drunen, *Computer Physics Communications* **91**, 43 (1995).
- [72] E. Lindahl, B. Hess, and D. van der Spoel, *Molecular modeling annual* **7**, 306 (2001).
- [73] D. V. D. Spoel, E. Lindahl, B. Hess, G. Groenhof, A. E. Mark, and H. J. C. Berendsen, *Journal of Computational Chemistry* **26**, 1701 (2005).
- [74] B. Hess, C. Kutzner, D. van der Spoel, and E. Lindahl, *Journal of Chemical Theory and Computation* **4**, 435 (2008).
- [75] S. Pronk, S. Páll, R. Schulz, P. Larsson, P. Bjelkmar, R. Apostolov, M. R. Shirts, J. C. Smith, P. M. Kasson, D. van der Spoel, B. Hess, and E. Lindahl, *Bioinformatics* **29**, 845 (2013).
- [76] S. Páll, M. J. Abraham, C. Kutzner, B. Hess, and E. Lindahl, in *Solving software challenges for exascale*, edited by S. Markidis and E. Laure (2015), pp. 3–27.

- [77] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, *SoftwareX* **1-2**, 19 (2015).
- [78] P. Eastman, J. Swails, J. D. Chodera, R. T. McGibbon, Y. Zhao, K. A. Beauchamp, L.-P. Wang, A. C. Simmonett, M. P. Harrigan, C. D. Stern, R. P. Wiewiora, B. R. Brooks, and V. S. Pande, *PLoS computational biology* **13**, PCOMPBIOL-D-16-01696[PII], e1005659 (2017).
- [79] C. M. Bishop, *Pattern recognition and machine learning (information science and statistics)* (Springer-Verlag, Berlin, Heidelberg, 2006).
- [80] R. Shu, H. H. Bui, S. Zhao, M. J. Kochenderfer, and S. Ermon, in *Advances in neural information processing systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018), pp. 4393–4402.
- [81] J. C. Spall, *Introduction to stochastic search and optimization*, 1st ed. (John Wiley & Sons, Inc., USA, 2003).
- [82] E. Marinari and G. Parisi, *Europhysics Letters (EPL)* **19**, 451 (1992).
- [83] R. Chandra, K. Jain, R. V. Deo, and S. Cripps, (2018).
- [84] A. Doucet, N. de Freitas, and N. Gordon, “An introduction to sequential monte carlo methods”, in *Sequential monte carlo methods in practice*, edited by A. Doucet, N. de Freitas, and N. Gordon (Springer New York, New York, NY, 2001), pp. 3–14.
- [85] P. Del Moral, A. Doucet, and A. Jasra, *J. R. Stat. Soc. Series B Stat. Methodol.* **68**, 411 (2006).
- [86] P. Del Moral, *Feynman-Kac formulae: genealogical and interacting particle systems with applications* (Springer New York, Dec. 2011).
- [87] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, in *Advances in neural information processing systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017), pp. 971–980.
- [88] D. P. Kingma and J. Ba, in 3rd international conference on learning representations, ICLR 2015, san diego, ca, usa, may 7-9, 2015, conference track proceedings (2015).
- [89] G. O. Roberts and R. L. Tweedie, *Bernoulli* **2**, 341 (1996).
- [90] G. Ramachandran, C. Ramakrishnan, and V. Sasisekharan, *Journal of Molecular Biology* **7**, 95 (1963).
- [91] R. Vargas, J. Garza, B. P. Hay, and D. A. Dixon, *The Journal of Physical Chemistry A* **106**, 3213 (2002).
- [92] E. J. Sorin and V. S. Pande, *Biophys J* **88**, 2472 (2005).
- [93] A. J. DePaul, E. J. Thompson, S. S. Patel, K. Haldeman, and E. J. Sorin, *Nucleic Acids Res* **38**, 4856 (2010).
- [94] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids* (Clarendon Press, New York, NY, USA, 1989).
- [95] A. Onufriev, D. Bashford, and D. A. Case, *Proteins: Structure, Function, and Bioinformatics* **55**, 383 (2004).
- [96] W. C. Still, A. Tempezyk, R. C. Hawley, and T. Hendrickson, *Journal of the American Chemical Society* **112**, 6127 (1990).
- [97] J. M. Troyer and F. E. Cohen, *Proteins: Structure, Function, and Bioinformatics* **23**, 97 (1995).
- [98] W. Humphrey, A. Dalke, and K. Schulten, *Journal of Molecular Graphics* **14**, 33 (1996).
- [99] T. P. Minka, “A family of algorithms for approximate bayesian inference”, PhD thesis (Massachusetts Institute of Technology, Cambridge, MA, USA, 2001).
- [100] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, *Mach. Learn.* **37**, 183 (1999).
- [101] J. A. T. Thomas M. Cover, *Elements of information theory*, 2nd ed, Wiley Series in Telecommunications and Signal Processing (Wiley-Interscience, 2006).
- [102] H. Kahn and T. E. Harris, *Journal of Research of the National Bureau of Standards. Applied Mathematics Series* **12**, 27 (1951).
- [103] J. P. Valleau and D. N. Card, *The Journal of Chemical Physics* **57**, 5457 (1972).
- [104] S. P. Carmichael and M. S. Shell, *The Journal of Physical Chemistry B* **116**, PMID: 22300263, 8383 (2012).
- [105] A. M. Fluitt and J. J. de Pablo, *Biophysical Journal* **109**, 1009 (2015).
- [106] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, (2017).

[107] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*, [http : / / www .
deeplearningbook.org](http://www.deeplearningbook.org) (MIT Press, 2016).