

Acceleration with Interpretability: Surrogate Model Based Collective Variable for Enhanced Sampling

Sompriya Chatterjee^{†,‡} and Dhiman Ray^{*,†,‡}

[†]*Department of Chemistry and Biochemistry, University of Oregon, Eugene, Oregon 97403,
USA*

[‡]*Materials Science Institute, University of Oregon, Eugene, Oregon 97403, USA*

E-mail: dray@uoregon.edu

Abstract

Most enhanced sampling methods facilitate the exploration of molecular free energy landscapes by applying a bias potential along a reduced dimensional collective variable (CV) space. The success of these methods depends on the ability of the CVs to follow the relevant slow modes of the system. Intuitive CVs, such as distances or contacts, often prove inadequate, particularly in biological systems involving many coupled degrees of freedom. Machine learning algorithms, especially neural networks (NN), can automate the process of CV discovery by combining a large number of molecular descriptors and often outperform intuitive CVs in sampling efficiency. However, their lack of interpretability and high cost of evaluation during trajectory propagation make NN-CVs difficult to apply to large biomolecular processes. Here, we introduce a surrogate model approach using lasso regression to express the output of a neural network as a linear combination of an automatically chosen subset of the input descriptors. We demonstrate successful applications of our surrogate model CVs in the enhanced sampling simulation of the conformational landscape of alanine dipeptide and chignolin mini-protein. In addition to providing mechanistic insights due to their explainable nature, the surrogate model CVs showed a negligible loss in efficiency and accuracy, compared to the NN-CVs, in reconstructing the underlying free energy surface. Moreover, due to their simplified functional forms, these CVs are better at extrapolating to unseen regions of the conformational space, e.g., saddle points. Surrogate model CVs are also less expensive to evaluate compared to their NN counterparts, making them suitable for enhanced sampling simulation of large and complex biomolecular processes.

1 Introduction

Molecular dynamics (MD) simulations have made it possible, in the past few decades, to study biomolecular processes in atomistic detail.¹ However, significant challenges remain in simulating rare events where high free energy barriers preclude the adequate sampling of the molecular conformational space. Enhanced sampling methods such as umbrella sampling (US),² metadynamics (MetaD),³ Gaussian accelerated MD (GaMD),⁴ Adaptive Biasing Force (ABF),⁵ etc. can address this problem and have been successfully applied to a wide range of biomolecular systems. To accelerate the sampling of physiologically relevant processes, most enhanced sampling methods drive transitions across high energy barriers by applying a biasing force along a low dimensional collective variable (CV) space that captures the relevant dynamical modes of the system.^{2,3,5,6} Due to the inherent complexity of biomolecules, intuitive CVs, such as distances, contacts, or torsion angles, often cannot describe the slow modes of the system, resulting in reduced sampling efficiency and accuracy.

To address this issue, machine learning algorithms have been used, in recent days, to transform a high-dimensional geometric descriptor space (\mathbf{d}) into a low-dimensional CV space $\mathbf{s} = \mathbf{s}(\mathbf{d})$ in a data-driven manner.⁷⁻²⁰ Particularly notable are the deep neural network (NN) based CV discovery algorithms, which overcomes the limitation of conventional linear machine learning approaches by allowing the possibility of capturing non-linear coupled dynamics among various degrees of freedom. In addition, NN-CVs can handle a significantly larger descriptor space compared to conventional approaches (e.g. principal component analysis (PCA),²¹ time-lagged independent component analysis (TICA),^{22,23} harmonic linear discriminant analysis (HLDA),¹⁵ etc.) and are less sensitive to the correlations among the different descriptors.¹⁶ NN-CVs, therefore, eliminate the requirement of manual feature selection, significantly automating the CV discovery protocol. Previous studies have demonstrated improved sampling efficiency of NN-CVs over conventional ML CVs based on PCA, TICA, and HLDA.^{16,24}

The increased performance, however, comes at the cost of evaluating complex non-linear

functions with a large number of parameters ($10^3 - 10^5$) at every time step of the MD simulation to calculate the bias potential along the CV. This can considerably increase the computational cost of the already expensive MD simulation algorithm. Furthermore, due to the black-box nature of the NN-CV, the user has little control over the amount of bias deposited along different degrees of freedom. This increases the likelihood of biasing irrelevant degrees of freedom included in the blindly chosen input descriptor space, potentially reducing the sampling efficiency. Furthermore, neural networks have limited ability to extrapolate beyond the training data, which is often collected only in the metastable states.^{15,16,18} Therefore, they often do not perform well in saddle points and transition regions.¹⁹ It can often be prohibitively expensive, especially for large biological systems, to extensively sample the conformational landscape just to train a CV. In addition, the complex NN-CVs are less interpretable compared to conventional ML CVs e.g. HLDA, TICA, etc., that assign specific weights to each input descriptor. Therefore, extracting mechanistic information from NN-CVs can be difficult and often requires additional sensitivity analysis.^{16,25}

In this work, we attempt to address these limitations of NN-CVs without sacrificing their sampling efficiency. We employ an explainable artificial intelligence (XAI) framework known as surrogate models, which refers to a broad category of methods that utilizes a simple interpretable function to represent the input-output relationship of a complex and less interpretable model, e.g., neural network.²⁶ Surrogate models are used in many areas of science and engineering²⁷⁻³⁰ where computational efficiency and interpretability are key considerations, and gaining a deeper understanding of the relationship among input variables is necessary. A surrogate model approach called “Thermodynamics-inspired Explainable Representations of AI and other black-box Paradigms (TERP)” has recently been applied to molecular processes to interpret neural-network-based dynamical models.³¹ Here, we take a simpler approach by using the lasso regression algorithm using the NN-CV as our target function³² to express the output of a neural network as a linear combination of a reduced set of input descriptors. Lasso algorithms have been applied earlier to identify key molecu-

lar features that can discriminate between metastable states³³ and to evaluate the relative contribution of different descriptors in an NN-CV.^{34,35} Here, we use the output of the lasso regression model as a CV to perform enhanced sampling simulations. Being a considerably simpler function, the computational effort to evaluate the lasso regression CV is negligible compared to a full neural network. Furthermore, the CV is represented as a combination of a handful of key degrees of freedom, thereby directing the bias potential only along relevant degrees of freedom. In addition, the surrogate CV is significantly more interpretable due to its linear nature and the use of fewer descriptors, making it relatively straightforward to gain molecular-level mechanistic insights. Moreover, lasso regression models can provide a better description of the saddle points in the free energy surface as these simpler models are more likely to generalize to new instances where little training data is available.^{32,36–38} Unlike other linear ML-CVs, the user no longer needs to perform manual feature selection as this step is automated by the training of the initial neural network CV and k -fold cross-validation in the lasso regression step. Here, we provide the theoretical framework of our approach and demonstrate its effectiveness in studying biomolecular processes using the examples of alanine dipeptide and chignolin mini-protein.

2 Theory

In this section, we briefly discuss the theoretical underpinnings of the neural-network-based CV discovery protocols and the lasso regression approach for constructing surrogate models of deep-learning CVs.

2.1 Deep Neural Networks for Collective Variable Discovery

In this work, we constructed surrogate models for two forms of Deep Neural Network CVs: Deep Targeted Discriminant Analysis (Deep-TDA) and Deep Time-lagged Independent Component Analysis (Deep-TICA). The Deep-TDA algorithm aims to discriminate between the

metastable states of a system, while the Deep-TICA CV seeks to learn the slow degrees of freedom for a molecular process. Below we provide a brief overview of these two methods.

2.1.1 Deep-TDA

The Deep-TDA algorithm¹⁸ uses a feed-forward Neural Network (NN) to discriminate between different molecular conformations based on a large number of descriptors (e.g. distance, contacts, etc.) as inputs. The training data are collected by performing short unbiased simulations in each known metastable basin. Given a system with N_m metastable states that can be characterized by a set of N_d descriptors \mathbf{d} , the NN parameters are optimized to map the multi-dimensional space of descriptors \mathbf{d} into a N_s dimensional CV, \mathbf{s} ($N_s \ll N_d$). The sampled configurations from each metastable state, when projected along the output node (\mathbf{s}), are distributed as non-overlapping Gaussian functions located at pre-defined locations of the CV space. This is achieved by minimizing the loss function:

$$\mathcal{L}_{TDA} = \alpha \sum_k^{N_m} \sum_l^{N_d} (\mu_{k,l} - \bar{\mu}_{k,l})^2 + \beta \sum_k^{N_m} \sum_l^{N_d} (\sigma_{k,l} - \bar{\sigma}_{k,l})^2, \quad (1)$$

where the first term enforces the mean of the k -th metastable state distribution along the l component of \mathbf{s} to remain close to the center of the target Gaussians $\bar{\mu}_{k,l}$. Similarly, the second term ensures the variance $\sigma_{k,l}$ of the metastable states distributions follow the same of the target distribution $\bar{\sigma}_{k,l}$. The hyperparameters, α and β , determine the relative weights of the two components of the loss function. In the CV space, the target distributions are placed sufficiently far from each other to avoid overlap between the different metastable states and transition regions.

2.1.2 Deep-TICA

The Deep-TICA algorithm¹⁷ is a non-linear version of the time-lagged independent component analysis (TICA)^{22,23,39} where the non-linearity is introduced via a neural network. In

a conventional TICA approach, the slow modes of a system are extracted via singular value decomposition (SVD) of the time-lagged covariance matrix $\mathbf{C}(\tau)$ of molecular descriptors:

$$\mathbf{C}(\tau)\mathbf{U} = \mathbf{C}(0)\mathbf{U}\mathbf{\Lambda}, \quad (2)$$

where τ is the lag-time, \mathbf{U} is the eigenvector-matrix containing the time-lagged independent components (ICs), and $\mathbf{\Lambda}$ is a diagonal eigenvalue matrix. The elements of the matrix $\mathbf{C}(\tau)$ are given by:

$$C_{ij}(\tau) = \frac{1}{T - \tau - 1} \sum_{t=1}^{T-\tau} d_i(\mathbf{r}^{3N}(t))d_j(\mathbf{r}^{3N}(t + \tau)) \quad (3)$$

where T is the length of the time series and $d_i(\mathbf{r}^{3N}(t))$ is the value i -th molecular descriptor at time step t . The descriptors are functions of the $3N$ dimensional atomic coordinates \mathbf{r}^{3N} . After sorting the eigenvectors in decreasing order of eigenvalues, the first eigenvector points to the direction of the slowest degrees of freedom and so on.

The Deep-TICA algorithm¹⁷ modifies the TICA approach by creating non-linear combinations of the molecular descriptors (d) through the hidden layers of a feed-forward neural network. The time-lagged covariance matrix is then constructed in terms of these non-linear combinations. The neural network parameters are optimized by maximizing the largest eigenvalues of the covariance matrix. Following the variational approach to conformational dynamics (VAC),⁴⁰ the Deep-TICA CV can also be trained using biased enhanced sampling trajectories as long as the bias is converged and the timescales have been appropriately reweighted. The elements of the time-lagged covariance matrix will, therefore, be expressed as:

$$C_{ij}(\tau) = \frac{\int_0^{T-\tau} dt \exp(\beta V(s(\mathbf{r}^{3N}(t)))) h_i(\mathbf{r}^{3N}(t)) h_j(\mathbf{r}^{3N}(t + \tau))}{\int_0^{T-\tau} dt \exp(\beta V(s(\mathbf{r}^{3N}(t))))} \quad (4)$$

where V is the bias potential, s is the CV along which the bias V has been applied, h_i, h_j are the non-linear combinations of descriptors generated by the neural network, and β is inverse

temperature. The integral is evaluated over the discrete MD trajectory. After training the neural network, the singular value decomposition of the biased time-lagged covariance matrix is performed, and the slow degrees of freedom are identified as the eigenvectors with the largest eigenvalues.

2.2 Surrogate Models Using Lasso Regression

To make neural network CVs interpretable, we utilized one of the simplest forms of the surrogate models: the lasso regression. We aim to recover the input-output relationship of the neural network using a linear combination of a subset of the input descriptors (d_i). In lasso regression, the predicted values (y_{pred}) of the target variable (y_{target}) is given by:

$$y_{\text{pred}} = b + \sum_{i=1}^{N_d} w_i d_i \quad (5)$$

The weights (w_i) and bias (b) are optimized by minimizing the loss function:

$$\mathcal{L}_{\text{lasso}} = \frac{1}{2M} \|\mathbf{y}_{\text{target}} - \mathbf{y}_{\text{pred}}\|_2^2 + \lambda \|\mathbf{w}\|_1 \quad (6)$$

where M is the number of data points and $\mathbf{y}_{\text{target}}$, \mathbf{y}_{pred} , and \mathbf{w} are vectors containing the target values, predicted values, and the weights, respectively, while $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the $L1$ and $L2$ norms of a vector. Minimizing this loss function sets the value of some of the weights to zero, reducing the effective dimensionality of the descriptor space. The hyperparameter, λ , determines the number of coefficients (w) that will be set to zero. It, therefore, creates a sparse linear model as a function of a limited number of descriptors that are most relevant for reproducing the target variable. In our case, we use the CV space learned from the neural network as our target variable y_{target} and the descriptors supplied in the input node of the NN as the feature space of lasso regression. This idea is illustrated in Figure 1 where the output node of an 8-dimensional feedforward neural network

is approximated as a linear combination of three of the descriptors. In earlier work, the lasso algorithm has been used to classify metastable states in protein conformational space as a post-analysis of the enhanced sampling simulation.³³ It has also been proposed as a suitable method to interpret deep-learning CVs.³⁴ Here, we use the reduced descriptions learned by the lasso regression as collective variables and apply bias potential along them to perform enhanced sampling.

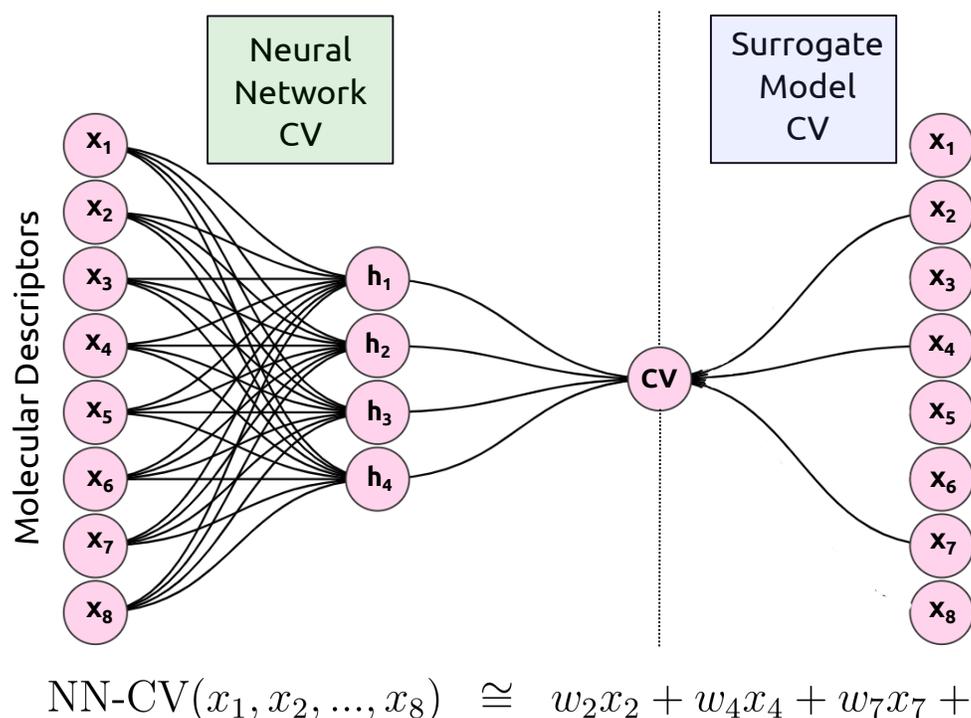


Figure 1: Schematic of a surrogate model CV constructed to reproduce the input-output relationship of a neural network CV. On the left, a neural network of 8-4-1 architecture is used to construct a CV from 8 input descriptors. On the right, the same descriptors are used in lasso regression to predict the outcome of the neural network (i.e. the CV). Due to the introduction of sparsity in the model, only 3 out of 8 descriptors contribute to the regression model. At the bottom, we show the mathematical expression of how the linear lasso regression model approximates the outcome of the neural network. This figure uses the following notations: $x_i = i$ -th descriptor, $h_i = i$ -th node of the hidden layer, $w_i =$ regression coefficient of the i -th descriptor in the surrogate model.

3 Computational Methods

3.1 Alanine Dipeptide

First, we tested our approach for the conformational transition between the $C7_{eq}$ and $C7_{ax}$ states of alanine dipeptide in vacuum. This is a prototypical system for testing new rare event sampling algorithms. This 22-atom system has been modeled using the AMBER99SB-ILDN force field.⁴¹ All simulations were performed using GROMACS 2024.3⁴² patched with PLUMED 2.11⁴³ using parameters identical to those in previous work.²⁴ As descriptors, we used the 45 interatomic distances introduced in Ref. 16. To train the Deep-TDA CV, we collected unbiased configurations from the two metastable states by performing 1 ns of unbiased MD simulation in each state. A neural network architecture of the form 45-30-15-1 was used where the 1D output node acted as the CV. The values of α and β of the loss function (Eq. 1) were set to 1 and 250, respectively. The target distributions of the $C7_{eq}$ and $C7_{ax}$ states were centered at $s = -5$ and $s = +5$ with a width of 0.2 arbitrary unit. The loss function was minimized using the ADAM optimizer with a learning rate of 10^{-3} . All NN-CVs were trained using the `mlcolvar` package.⁴⁴ To check the versatility of our approach, we trained another Deep-TDA CV using the 39 torsion angles introduced by Mendels et al.⁴⁵ as descriptors. A cosine transformation of the form: $f(\theta) = 0.5 + \cos(\theta - 1.25)$ is performed following Ref. 7 to address the periodicity of the dihedral angles before supplying them to the input node of the neural network.

For training the Deep-TICA CV, we performed a 10 ns On-the-fly Probability Enhanced Sampling (OPES) simulation⁴⁶ with bias deposition along the ϕ and ψ torsion angles. The convergence of the simulation is monitored by observing multiple back-and-forth transitions between the states and by ensuring the bias has reached the quasistatic regime. The last 8 ns of the OPES trajectory has been used for training the deep TICA CV after reweighting the configurations by the bias potential. The 45 pairwise distances have been used as descriptors within a neural network architecture of 45-15-15-5, and only the first eigenvector (slowest

degree of freedom) has been used as CV. The rest of the hyperparameters for the neural network optimization were identical to the Deep-TDA CV.

For each of the Deep-TDA CVs and the Deep-TICA CV, a surrogate model CV has been constructed using linear regression with L1 penalty (lasso regression) with 5-fold cross-validation to determine the optimal value of the lasso coefficient λ . The lasso regression model has been trained on the same configurations used for the corresponding neural network training. The values of the NN-CV have been used as the target values (y_{target}).

To assess the quality of the deep CVs and the lasso regression CVs, enhanced sampling simulations have been performed using the OPES algorithm with a BARRIER parameter (ΔE) of 40 kJ/mol. For each of the six CV combinations (Deep-TDA and Deep-TICA with distance descriptors, Deep-TDA with torsion descriptors, and the corresponding lasso regression CVs), three independent 10 ns OPES simulations have been performed starting from the $C7_{eq}$ basin. The first 1 ns of each simulation has been discarded, and the rest of the trajectory data is used to assess the free energy convergence. Kinetics of the transition from $C7_{eq}$ to $C7_{ax}$ state are also computed using the distance-based NN-CVs and their corresponding surrogate models using the OPES-flooding algorithm.²⁴ For these simulations, we used $\Delta E = 25$ kJ/mol, and the excluded region boundaries were chosen based on the free energy profiles to avoid bias deposition in the transition states.^{24,47}

3.2 Chignolin

Next, we evaluated the effectiveness of our surrogate model approach in sampling the conformational landscape of the fast-folding chignolin mini-protein. This 10-residue polypeptide has a funnel-like free energy landscape that mirrors standard protein folding pathways.⁴⁸ In its folded state, it forms a beta-hairpin conformation, which can unfold into a disordered structure within the microsecond timescale.⁴⁹ In previous work, the folding and unfolding dynamics of the CLN025 mutant of chignolin were investigated by performing an extended, unbiased MD simulation using the Anton supercomputer.⁵⁰ The simulation estimated the

unfolding and folding timescales to be $2.2 \pm 0.4 \mu\text{s}$ and $0.6 \pm 0.1 \mu\text{s}$, respectively, using the CHARMM22* force field⁵¹ for the protein and the TIP3P model⁵² for the solvent. Due to the availability of the unbiased reference simulation, this system is widely used to test new enhanced sampling methods.^{24,53–57} Moreover, deep learning CVs such as Deep-TDA and Deep-TICA have been demonstrated to achieve a converged free energy landscape for the chignolin system.^{17,19}

In the present study, we used a simulation setup identical to that of Ref. 50 to assess the quality of sampling of the free energy landscapes, independent of the force field accuracy. All simulations were performed using the CUDA-enabled version of GROMACS 2024.3⁴² patched with PLUMED 2.11.⁴³ We trained the Deep-TDA CV on 10 ns long unbiased MD simulations in the folded and unfolded states of chignolin. Contrarily, for the training of Deep-TICA CV, we utilized the whole $\sim 106 \mu\text{s}$ long unbiased trajectory from Ref. 50 as one needs to sample multiple back-and-forth transitions to learn the slow degrees of freedom via Deep-TICA analysis. It should be noted that in most practical systems, such an unbiased trajectory is not available, and one needs to utilize biased simulation data for training Deep-TICA CV following the protocol introduced in Ref. 17. To train both the neural network CVs, we used a descriptor set comprised of the 45 pairwise distances between the 10 α -carbon atoms of chignolin. Apart from the 45-dimensional input layer, each neural network contains two hidden layers, each with 30 hidden nodes. The one-dimensional CV is obtained in the output node of the neural network after training. The surrogate model CVs were trained using lasso regression using a protocol identical to the one used for alanine dipeptide.

To assess the quality of the CVs, three independent OPES simulations using a BARRIER parameter (ΔE) of 30 kJ/mol were carried out for each of the four CV combinations: Deep-TDA, Deep-TICA, and their corresponding surrogate models. The initial configurations for the independent simulations were sampled from a 10 ns long unbiased simulation performed in the folded state. Each OPES simulation was extended to $2.5 \mu\text{s}$, sampling multiple re-crossing events between the folded and unfolded states. The last $2.25 \mu\text{s}$ of each run are

used to compute the free energy profile via reweighting.⁴⁶ We also tested the ability of our CVs to recover the unfolding kinetics of chignolin by performing OPES-flooding simulations using each of the four CV combinations. A total of 15 independent OPES-flooding simulations were initiated from the folded configurations and were terminated upon reaching the unfolded state. A BARRIER parameter (ΔE) of 12 kJ/mol was used for all OPES-flooding simulations. The excluded region boundaries were chosen based on the free energy surfaces obtained from the OPES simulation, and they were located at $s_{\text{exc}} = -3.5$ for Deep-TDA and -4.0 for its surrogate model, whereas $s_{\text{exc}} = 0.6$ and 0.7 were used for the Deep-TICA and its corresponding surrogate model CVs.

4 Results and Discussions

4.1 Alanine Dipeptide

We could converge the free energy landscape of the conformational transition in alanine dipeptide using all the tested CV combinations. The surrogate model of the Deep-TDA CV led to noticeably more transitions and a better free energy convergence than the NN-CV. This counterintuitive result can be attributed to the high bias and low variance nature of the linear model, making it better in extrapolation than the non-linear neural networks.³² This can be observed in Fig. 2, which shows that the gradient of the surrogate model CV also aligns better with the gradients of the underlying free energy surface and better distinguishes the transition state region compared to the original Deep-TDA CV. The Deep-TDA CV is only trained on configurations sampled within the metastable state minima (SI Fig. S5). Therefore, its accuracy and performance critically depend on its ability to extrapolate to unseen regions of the conformational landscape, particularly the transition states, which the systems need to visit during the enhanced sampling simulations. The improvement in efficiency also stems from the fact that the surrogate model CVs apply bias along a reduced set of descriptors, thereby directing the sampling only along relevant degrees of freedom.

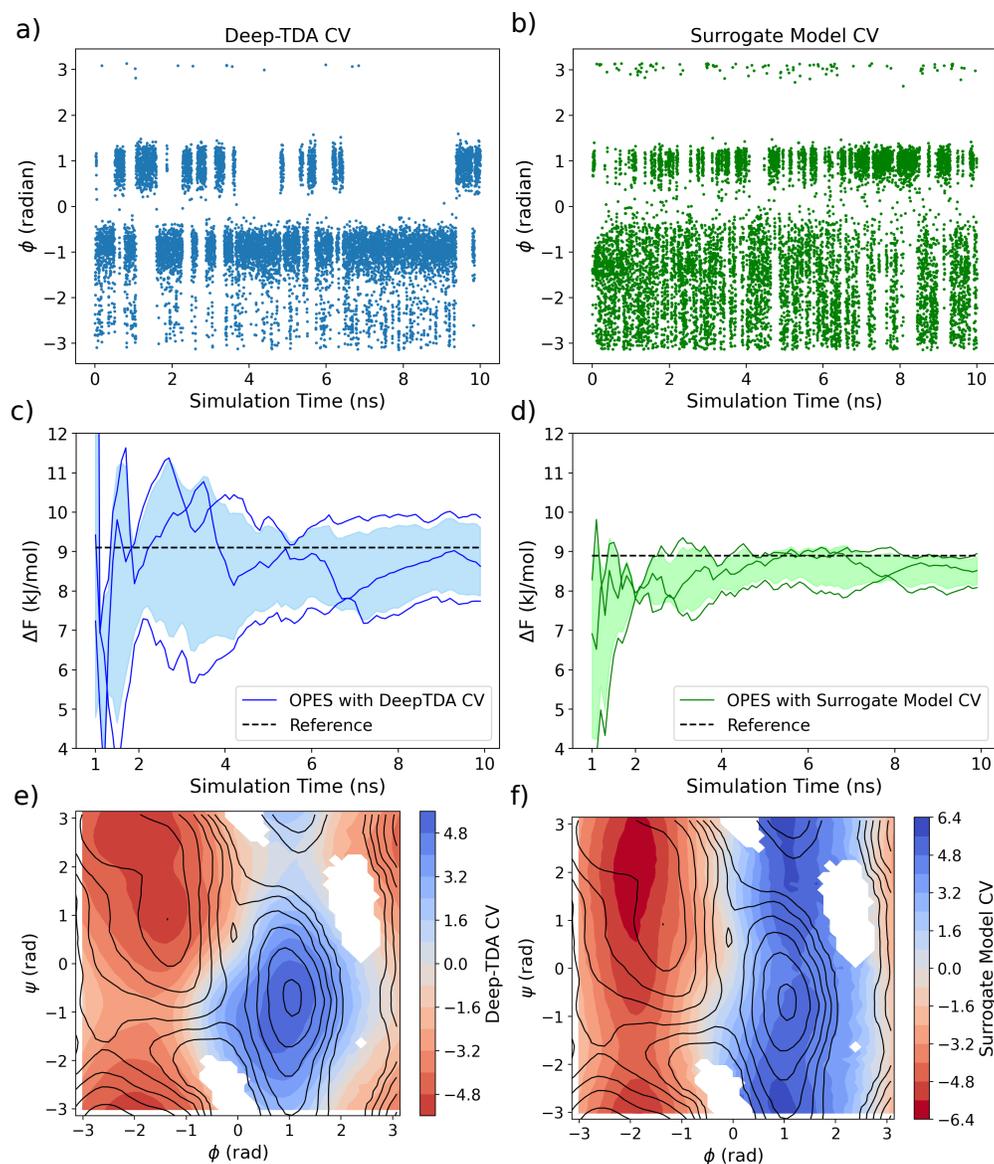


Figure 2: Upper panel: Time evolution of the ϕ torsion angle of alanine dipeptide in OPES simulation using (a) Deep-TDA CV and (b) its surrogate model CV. Middle Panel: The convergence of the free energy difference between the $C7_{eq}$ and $C7_{ax}$ states of alanine dipeptide using (a) the Deep-TDA CV and (b) its surrogate model CV for biasing. The reference free energy value is depicted in a black dashed line. Bottom Panel: Contour plots of the collective variables projected onto the two-dimensional Ramachandran angle (ϕ and ψ) space. The contours of the reference free energy surface are also shown as solid black lines.

While the original Deep-TDA CV applies bias along all the 45 distance descriptors (or 39 torsion-angle descriptors), the surrogate model CVs, trained using lasso regression, bias only along 15 distance descriptors (or 17 torsion descriptors).

Another edge of the surrogate model CV is that it is considerably less expensive to evaluate than its neural network counterpart due to the use of fewer parameters. It, therefore, significantly speeds up the enhanced sampling simulation, completing the same number of integration steps as the NN-CV in almost an order of magnitude less wall clock time.

In the case of the Deep-TICA CV, the convergence speed and the frequency of transitions between the two metastable states are comparable for the NN-CV and the surrogate model CV, despite the latter model using a linear combination of only 22 of the 45 descriptors. Unlike the Deep-TDA CV, we do not observe an improvement in convergence as the Deep-TICA CV is already trained on configurations coming from the entire conformational space to be sampled via enhanced sampling, including transition state regions (SI Fig. S5). Therefore, the extrapolation capability of the CV is not as critical as it is in the Deep-TDA CV. Nevertheless, the reduced CPU time makes the overall protocol significantly more efficient than directly biasing the NN-CV.

In addition to the accurate prediction of thermodynamic properties such as the free energy surface, when used in combination with methods like OPES-flooding,²⁴ surrogate model CVs of both Deep-TDA and Deep-TICA are equally capable of recovering the kinetics of the transition from the $C7_{eq}$ to the $C7_{ax}$ state. The surrogate model CVs led to mean first passage time (MFPT) estimates in agreement with the unbiased reference value of 1.28 μ s within the 95% confidence intervals. The reweighted first passage times follow an exponential distribution almost perfectly, which is the hallmark of reliable kinetics estimation for single barrier crossing events.⁵⁸

Another key advantage of the surrogate model CVs is their interpretability. The contribution of each descriptor toward the CV can be obtained directly from the regression coefficients. The same three distance descriptors contribute the most toward both Deep-

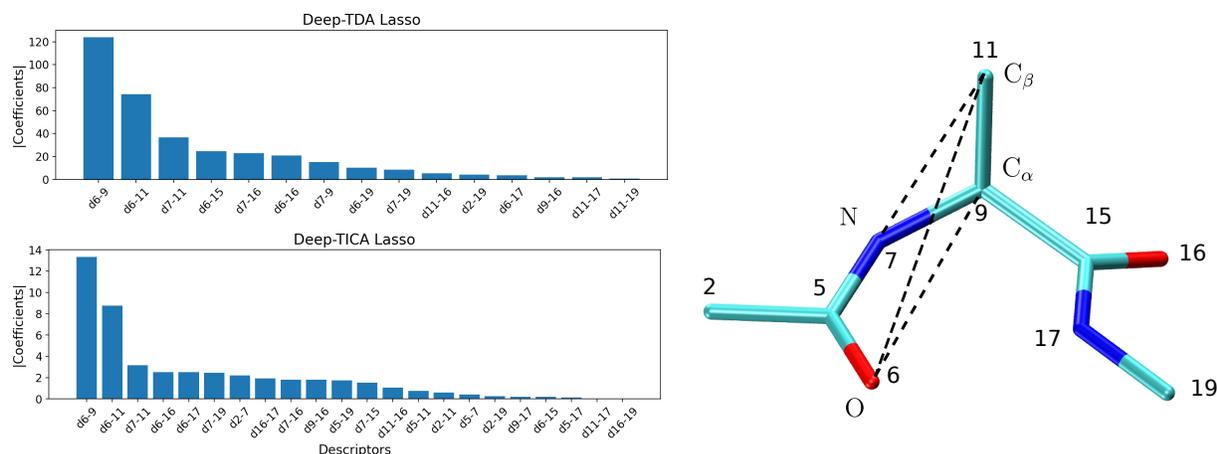


Figure 3: Left Panel: Magnitudes of the regression coefficients of the descriptors with non-zero contribution in the surrogate models of the Deep-TDA CV (upper) and Deep-TICA CV (lower) in alanine dipeptide. Right Panel: Molecular structure of the alanine dipeptide molecule. The atom numbers of the heavy atoms and the three interatomic distances with the strongest contributions toward the CVs are shown explicitly.

TDA and Deep-TICA CVs: the distances between atoms O and C_α , atoms O and C_β , and atoms N and C_β (Fig. 3). These three descriptors and their ordering are identical to the three most important descriptors identified in previous work through a more extensive committor analysis.⁵⁹ The regression coefficients of the torsion descriptors indicate the ϕ dihedral angle to have the most significant contribution towards the Deep-TDA CV (SI Fig. S6). This observation agrees with our preexisting knowledge that the ϕ angle constitutes an optimal CV for the alanine dipeptide conformational transition, reinforcing the ability of the surrogate model to correctly predict the most relevant descriptors.

4.2 Chignolin

In the case of chignolin, the deep neural network CVs apply bias along all 45 pairwise distances between the C_α atoms, whereas the surrogate models CVs trained through lasso regression of the NN output apply bias along a much smaller set of descriptors: 8 descriptors for Deep-TDA and 20 for Deep-TICA. This reduces the model complexity significantly, considering the original Deep-TDA and Deep-TICA CVs include ~ 2300 parameters. The

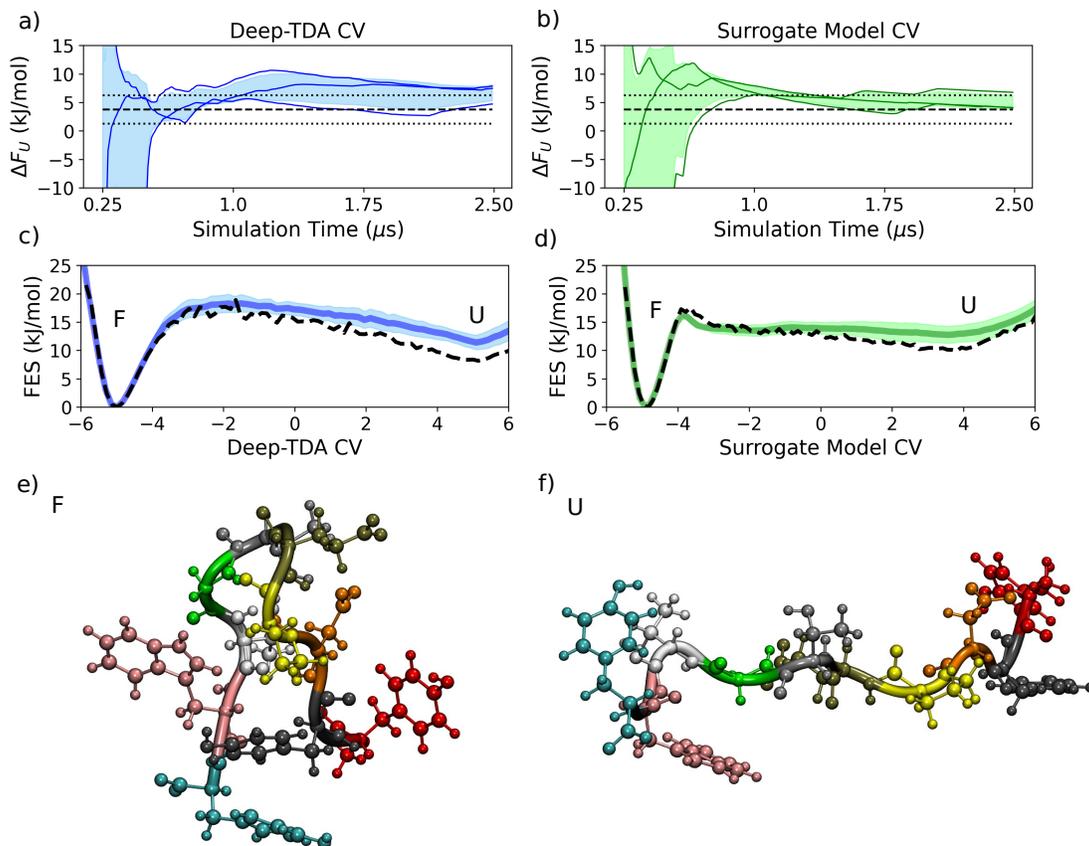


Figure 4: Top panel: Convergence of the free energy difference between the folded and unfolded states of chignolin using OPES simulations depositing bias along (a) Deep-TDA and (b) its corresponding surrogate model CV. The black dashed line denotes the reference unfolding free energy from Ref. 50. The black dotted line shows $\pm 1 k_B T$ variations from the reference value. The solid line and shaded region indicate the mean and standard deviation of unfolding free energy from three independent OPES simulations. Middle panel: Mean and uncertainties of the free energy profiles obtained from three independent OPES simulations using (c) Deep-TDA and (d) its corresponding surrogate model CV. The reference free energy profile from unbiased simulation⁵⁰ is shown in a black dashed line. The folded (F) and unfolded (U) states are labeled on the free energy profiles. Bottom panel: Representative structures of the folded (F) and unfolded (U) conformations of chignolin.

lasso regression CV showed nearly identical free energy convergence to that of the standard Deep-TDA (Fig. 4 a,b), demonstrating that the simplified surrogate model can perform as effectively as the neural network CV. The free energy surfaces (FES) for both CVs fall within one $k_B T$ of the reference value obtained from unbiased simulation of Ref. 50. In contrast, the surrogate model Deep-TICA CV shows a slightly slower convergence speed compared to its neural network counterpart (S9 a and b). This observation is consistent with the results of the alanine dipeptide system, where the surrogate model outperformed the Deep-TDA CV in terms of convergence speed but not the Deep-TICA CV. The Deep-TDA CV is trained only on the metastable state data and does not include any information about the saddle points of the free energy landscape, e.g., the transition state(s). Therefore, the superior extrapolation capability of the simplified regression model compensates for its lack of adaptability caused by the inclusion of a reduced number of descriptors and the strictly linear functional form. Contrarily, the Deep-TICA CV is trained on almost the entire conformational space of the protein, leaving little to no need for extrapolation during the enhanced sampling simulations. In such a data-rich environment, the high-variance neural network model can provide a better description of the slow modes through non-linear combinations of molecular descriptors.

It should be noted here that all four CVs led to a converged estimate of the free energy landscape within ~ 50 times less simulation time compared to the unbiased run. Therefore, the surrogate models for both Deep-TDA and Deep-TICA CVs can be used effectively in practical problems where the computational cost of evaluating a complex neural network and its gradients at every MD timestep can become a bottleneck. We observe that using identical computing hardware, the surrogate model CVs require almost half the wall clock time than that of the neural network CVs (SI Table S3). When we used a more complex model involving the 210 descriptors (introduced in Ref. 17) and ~ 44500 parameters, the NN CV required more than three times wall clock time than that of its corresponding surrogate model. In more realistic problems, therefore, the surrogate model CVs will have a clear edge over NN-CVs in terms of CPU time.

The second consideration is that in the case of more realistic systems, it is often not possible to sample the entire conformational landscape (either through unbiased simulation or through biased sampling with suboptimal CVs) to obtain the training configurations for the neural network CV. Therefore, the applicability of Deep-TICA CV in such systems can be rather limited. In contrast, the training of Deep-TDA CV only requires metastable state information. This is far easier to obtain in the case of large complex molecules by simply performing unbiased simulations in known metastable states, e.g., folded and unfolded states. As the surrogate model of Deep-TDA CV shows equal or better convergence speed and accuracy of the free energy landscapes, it can be used in general to study molecular systems of different sizes and complexity.

Our surrogate models also act as interpretable representations of their parent neural network CVs and provide molecular level details of the (un)folding process. Based on the magnitudes of the regression coefficients the distances between C_α atoms of the terminal residues such as $C_\alpha^9-C_\alpha^2$ [d9-2] contribute most strongly toward the Deep-TDA CV. In contrast, the distances $C_\alpha^6-C_\alpha^3$ [d6-3] and $C_\alpha^8-C_\alpha^6$ [d8-6] play a more significant role in the Deep-TICA CV (Fig. 5). Extensive committor analysis in previous work demonstrated that the hydrogen bond formations among Asp3, Thr6, and Thr8 residues play an important role in the transition state of chignolin folding.⁵⁹ The Deep-TICA CV is able to capture these interactions as its training data includes configurations from the whole conformational space, including the transition regions. In contrast, the Deep-TDA CV, which aims to distinguish the metastable states, emphasizes primarily the terminal residue interactions such as $C_\alpha^9-C_\alpha^2$ [d9-2] and $C_\alpha^{10}-C_\alpha^1$ [d10-1] as they alone are sufficient to discriminate between the compact folded hairpin and the extended unfolded structures. Despite the complete omission of the key hydrogen bonding interactions in the surrogate model of Deep-TDA, its ability to obtain the converged free energy landscape within a simulation time comparable to that of Deep-TICA CV reinforces the superior extrapolation capability of the lasso regression CV toward unseen regions of the conformational space.

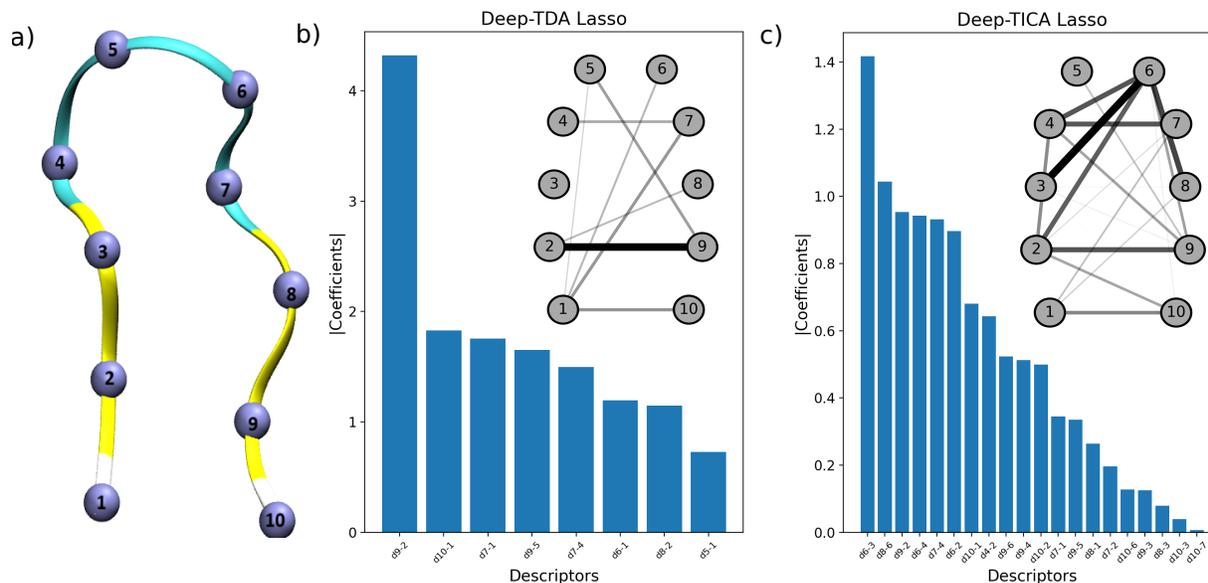


Figure 5: (a) Ribbon structure of the chignolin hairpin structure with C_{α} atoms shown explicitly and numbered according to their residue index. (b) and (c) Magnitudes of the non-zero regression coefficients of the surrogate models of the (b) Deep-TDA and (c) Deep-TICA CVs for the chignolin system. Networks depicting the pairwise distances with finite contributions toward the CV are provided as insets. The nodes denote the C_{α} atoms, and the thickness of the lines joining the nodes is proportional to the magnitude of the regression coefficient associated with the distance between those atoms.

Similar to alanine dipeptide, the surrogate model CVs are equally capable as their neural network counterparts in terms of recovering the unfolding kinetics from OPES-flooding simulations. All four CVs led to unfolding times within a factor of 3 from the unbiased estimate ($2.2 \pm 0.4 \mu\text{s}$) and acceptable exponential fit of the first passage time distributions ($p > 0.05$)⁵⁸ (Table 1). In fact, the unfolding kinetics from surrogate model CVs are in slightly better agreement with the reference value compared to their neural network counterparts. However, these differences are likely not statistically significant as they are within each other's confidence intervals. These results indicate that despite being only an approximation of the neural network CVs, the surrogate models perform equally well in OPES-flooding simulations for recovering the kinetic properties.

Table 1: Comparison of the unfolding kinetics of chignolin obtained from OPES flooding simulations using Deep-TDA, Deep-TICA, and their respective surrogate model CVs.

Method / CV	Number of runs	Mean first passage time (τ) in μs	p-value ^a	95% confidence interval in μs ^b
Deep-TDA	15	6.87	0.430	4.42 - 12.37
Deep-TDA (Surrogate)	15	4.66	0.397	3.01 - 8.42
Deep-TICA	15	3.58	0.992	2.81 - 7.87
Deep-TICA (Surrogate)	15	2.14	0.231	2.11 - 5.92

^aThe p-values were generated from 2 samples of the Kolmogorov–Smirnov test.⁵⁸ ^bWe computed the 95% confidence intervals using the method developed by Kaminsky.^{47,60}

5 Conclusions

Machine learning algorithms, including deep neural networks, have seen enormous success in the design of optimal collective variables for enhanced sampling simulation of molecular systems. However, the lack of interpretability and the high computational cost of evaluating large NN models in every integration timestep limits their applicability in studying large biomolecules. In this work, we present a surrogate model approach that reintroduces the interpretability of conventional intuitive CVs into the ones learned using a deep neural network. Surrogate models form part of the growing field of explainable artificial intelligence (XAI) that aims to make deep learning outcomes more interpretable. In our approach, we first train a deep-learning CV using an extensive set of molecular descriptors as input. Then, we use lasso regression to approximate the output of the neural network as a linear combination of a subset of the input descriptors, chosen by minimizing the L1 norm of regression coefficients and k -fold cross-validation. This approach automatically detects the highest contributing descriptors toward the neural network, removing the requirement of manual feature selection, commonly required in shallow learning algorithms, like HLDA.¹⁵ We show that despite this enormous simplification, the surrogate model CVs perform almost equally well in recovering the underlying free energy landscape and transition kinetics of molecular processes using enhanced sampling simulations. In fact, the surrogate models may outperform their neural network counterparts in terms of sampling and convergence

speed when no information about transition states and saddle points is included during the training of the neural network. This stems from the superior extrapolation capacity of the simpler regression approach over the more complex neural network models that do not generalize well beyond the training data. This is an important consideration for large and complex molecular systems, where it is difficult to sample the entire conformational space just to generate the training data for a machine learning CV. Particularly challenging is the sampling of transition states, although, in recent days, key advancements have been made in this direction.^{19,20,59,61}

In this proof-of-concept study, we limit ourselves to one of the simplest forms of surrogate model, linear regression. However, this work opens the possibility of incorporating other forms of surrogate models that include non-linearity and couplings between variables to improve the CVs without sacrificing interpretability and computational efficiency. We also envision a significantly higher performance gain from surrogate modeling of larger and significantly more complex neural network architectures, including Permutationally Invariant Networks for Enhanced Sampling (PINES)⁶² and Graph Neural Networks (GNN).^{35,63,64} In situations where the surrogate model CV is not as optimal as their neural network counterparts but provides noticeable gains in computational wall-clock time, they can be used in combination with the new variants of OPES specifically designed for quantitative sampling of free energy landscapes using suboptimal CVs.^{65,66} We aim to pursue these avenues in the future. This work, therefore, paves the way for advancing the field of enhanced sampling and CV discovery and facilitates mechanistic investigation of complex biological processes through molecular dynamics simulation.

Acknowledgement

This work has been supported by the new faculty start-up grant from the Office of the Vice President for Research and Innovation, University of Oregon. The authors thank Sudip Das,

Jintu Zhang, and Revanth Elangovan for stimulating discussions. The authors declare no competing financial interest.

Data Availability Statement

The input files for all simulations performed in this work and sample codes for training the Deep-TDA, Deep-TICA, and their surrogate model CVs are provided in the GitHub repository: https://github.com/sompriya/Surrogate_CV.git. The input files will also be made available through the PLUMED NEST repository.⁶⁷

Supporting Information Available

Additional results comparing two forms of deep NN CVs, Deep-TDA and Deep-TICA, with their surrogate models on the conformational transitions in alanine dipeptide and chignolin are shown in Figures S1–S13 and Tables S1–S4 in the Supporting Information.

References

- (1) Karplus, M.; McCammon, J. A. Molecular dynamics simulations of biomolecules. *Nature structural biology* **2002**, *9*, 646–652.
- (2) Torrie, G. M.; Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of computational physics* **1977**, *23*, 187–199.
- (3) Laio, A.; Parrinello, M. Escaping free-energy minima. *Proceedings of the National Academy of Sciences of the United States of America* **2002**, *99*, 12562–6.
- (4) Miao, Y.; Feher, V. A.; McCammon, J. A. Gaussian accelerated molecular dynam-

- ics: unconstrained enhanced sampling and free energy calculation. *Journal of chemical theory and computation* **2015**, *11*, 3584–3595.
- (5) Darve, E.; Pohorille, A. Calculating free energies using average force. *The Journal of chemical physics* **2001**, *115*, 9169–9183.
- (6) Hénin, J.; Lelièvre, T.; Shirts, M. R.; Valsson, O.; Delemotte, L. Enhanced Sampling Methods for Molecular Dynamics Simulations [Article v1.0]. *Living Journal of Computational Molecular Science* **2022**, *4*, 1583.
- (7) Tiwary, P.; Berne, B. Spectral gap optimization of order parameters for sampling complex molecular systems. *Proceedings of the National Academy of Sciences* **2016**, *113*, 2839–2844.
- (8) M. Sultan, M.; Pande, V. S. tICA-metadynamics: accelerating metadynamics by using kinetically selected collective variables. *Journal of chemical theory and computation* **2017**, *13*, 2440–2447.
- (9) Sultan, M. M.; Pande, V. S. Automated design of collective variables using supervised machine learning. *The Journal of chemical physics* **2018**, *149*.
- (10) Chen, W.; Ferguson, A. L. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *Journal of computational chemistry* **2018**, *39*, 2079–2102.
- (11) Ribeiro, J. M. L.; Bravo, P.; Wang, Y.; Tiwary, P. Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). *The Journal of chemical physics* **2018**, *149*.
- (12) Sidky, H.; Chen, W.; Ferguson, A. L. Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation. *Molecular Physics* **2020**, *118*, e1737742.

- (13) Wang, Y.; Ribeiro, J. M. L.; Tiwary, P. Machine learning approaches for analyzing and enhancing molecular dynamics simulations. *Current opinion in structural biology* **2020**, *61*, 139–145.
- (14) Fu, H.; Bian, H.; Shao, X.; Cai, W. Collective Variable-Based Enhanced Sampling: From Human Learning to Machine Learning. *The Journal of Physical Chemistry Letters* **2024**, *15*, 1774–1783.
- (15) Mendels, D.; Piccini, G.; Parrinello, M. Collective variables from local fluctuations. *The journal of physical chemistry letters* **2018**, *9*, 2776–2781.
- (16) Bonati, L.; Rizzi, V.; Parrinello, M. Data-driven collective variables for enhanced sampling. *The journal of physical chemistry letters* **2020**, *11*, 2998–3004.
- (17) Bonati, L.; Piccini, G.; Parrinello, M. Deep learning the slow modes for rare events sampling. *Proceedings of the National Academy of Sciences* **2021**, *118*, e2113533118.
- (18) Trizio, E.; Parrinello, M. From enhanced sampling to reaction profiles. *The Journal of Physical Chemistry Letters* **2021**, *12*, 8621–8626.
- (19) Ray, D.; Trizio, E.; Parrinello, M. Deep learning collective variables from transition path ensemble. *The Journal of Chemical Physics* **2023**, *158*.
- (20) Trizio, E.; Parrinello, M.; Kang, P. Everything everywhere all at once, a probability-based enhanced sampling approach to rare events. *arXiv preprint arXiv:2410.17029* **2024**,
- (21) David, C. C.; Jacobs, D. J. Principal component analysis: a method for determining the essential dynamics of proteins. *Protein dynamics: Methods and protocols* **2014**, 193–226.
- (22) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification

- of slow molecular order parameters for Markov model construction. *The Journal of chemical physics* **2013**, *139*.
- (23) Schwantes, C. R.; Pande, V. S. Improvements in Markov state model construction reveal many non-native interactions in the folding of NTL9. *Journal of chemical theory and computation* **2013**, *9*, 2000–2009.
- (24) Ray, D.; Ansari, N.; Rizzi, V.; Invernizzi, M.; Parrinello, M. Rare event kinetics from adaptive bias enhanced sampling. *Journal of Chemical Theory and Computation* **2022**, *18*, 6500–6509.
- (25) Rizzi, V.; Bonati, L.; Ansari, N.; Parrinello, M. The role of water in host-guest interaction. *Nature Communications* **2021**, *12*, 93.
- (26) Wilhelm, A.; Zweig, K. A. Hacking a surrogate model approach to XAI. *arXiv preprint arXiv:2406.16626* **2024**,
- (27) Luo, J.; Ma, X.; Ji, Y.; Li, X.; Song, Z.; Lu, W. Review of machine learning-based surrogate models of groundwater contaminant modeling. *Environmental Research* **2023**, 117268.
- (28) Nyshadham, C.; Rupp, M.; Bekker, B.; Shapeev, A. V.; Mueller, T.; Rosenbrock, C. W.; Csányi, G.; Wingate, D. W.; Hart, G. L. Machine-learned multi-system surrogate models for materials prediction. *npj Computational Materials* **2019**, *5*, 51.
- (29) Jakeman, J. D.; Kouri, D. P.; Huerta, J. G. Surrogate modeling for efficiently, accurately and conservatively estimating measures of risk. *Reliability Engineering & System Safety* **2022**, *221*, 108280.
- (30) McBride, K.; Sundmacher, K. Overview of surrogate modeling in chemical process engineering. *Chemie Ingenieur Technik* **2019**, *91*, 228–239.

- (31) Mehdi, S.; Tiwary, P. Thermodynamics of interpretation. *arXiv preprint arXiv:2206.13475* **2022**,
- (32) Edwards, R. E.; New, J.; Parker, L. E.; Cui, B.; Dong, J. Constructing large scale surrogate models from big data and artificial intelligence. *Applied energy* **2017**, *202*, 685–699.
- (33) Novelli, P.; Bonati, L.; Pontil, M.; Parrinello, M. Characterizing metastable states with the help of machine learning. *Journal of Chemical Theory and Computation* **2022**, *18*, 5195–5202.
- (34) Sparse linear models - mlcolvar documentation — mlcolvar.readthedocs.io. https://mlcolvar.readthedocs.io/en/latest/notebooks/tutorials/expl_lasso.html, [Accessed 30-07-2024].
- (35) Zhang, J.; Bonati, L.; Trizio, E.; Zhang, O.; Kang, Y.; Hou, T.; Parrinello, M. Descriptors-free collective variables from geometric graph neural networks. *arXiv preprint arXiv:2409.07339* **2024**,
- (36) Schwarz, G. Estimating the dimension of a model. *The annals of statistics* **1978**, 461–464.
- (37) Akaike, H. *Selected papers of hirotugu akaike*; Springer, 1998; pp 199–213.
- (38) Bozdogan, H.; Haughton, D. M. Informational complexity criteria for regression models. *Computational Statistics & Data Analysis* **1998**, *28*, 51–76.
- (39) Molgedey, L.; Schuster, H. G. Separation of a mixture of independent signals using time delayed correlations. *Physical review letters* **1994**, *72*, 3634.
- (40) McCarty, J.; Parrinello, M. A variational conformational dynamics approach to the selection of collective variables in metadynamics. *The Journal of chemical physics* **2017**, *147*.

- (41) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Structure, Function, and Bioinformatics* **2010**, *78*, 1950–1958.
- (42) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, *1*, 19–25.
- (43) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Computer physics communications* **2014**, *185*, 604–613.
- (44) Bonati, L.; Trizio, E.; Rizzi, A.; Parrinello, M. A unified framework for machine learning collective variables for enhanced sampling simulations: mlcolvar. *The Journal of Chemical Physics* **2023**, *159*.
- (45) Mendels, D.; de Pablo, J. J. Collective variables for free energy surface tailoring: Understanding and modifying functionality in systems dominated by rare events. *The Journal of Physical Chemistry Letters* **2022**, *13*, 2830–2837.
- (46) Invernizzi, M.; Parrinello, M. Rethinking metadynamics: from bias potentials to probability distributions. *The journal of physical chemistry letters* **2020**, *11*, 2731–2736.
- (47) Ray, D.; Parrinello, M. Kinetics from metadynamics: Principles, applications, and outlook. *Journal of Chemical Theory and Computation* **2023**, *19*, 5649–5670.
- (48) Honda, S.; Akiba, T.; Kato, Y. S.; Sawada, Y.; Sekijima, M.; Ishimura, M.; Ooishi, A.; Watanabe, H.; Odahara, T.; Harata, K. Crystal structure of a ten-amino acid protein. *Journal of the American Chemical Society* **2008**, *130*, 15327–15331.
- (49) Satoh, D.; Shimizu, K.; Nakamura, S.; Terada, T. Folding free-energy landscape of a 10-residue mini-protein, chignolin. *FEBS letters* **2006**, *580*, 3422–3426.

- (50) Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. How fast-folding proteins fold. *Science* **2011**, *334*, 517–520.
- (51) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. How robust are protein folding simulations with respect to force field parameterization? *Biophysical journal* **2011**, *100*, L47–L49.
- (52) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *The Journal of chemical physics* **1983**, *79*, 926–935.
- (53) Palazzesi, F.; Valsson, O.; Parrinello, M. Conformational entropy as collective variable for proteins. *The journal of physical chemistry letters* **2017**, *8*, 4752–4756.
- (54) Mendels, D.; Piccini, G.; Brotzakis, Z. F.; Yang, Y. I.; Parrinello, M. Folding a small protein using harmonic linear discriminant analysis. *The Journal of chemical physics* **2018**, *149*.
- (55) Belkacemi, Z.; Gkeka, P.; Lelièvre, T.; Stoltz, G. Chasing collective variables using autoencoders and biased trajectories. *Journal of chemical theory and computation* **2021**, *18*, 59–78.
- (56) Topel, M.; Ejaz, A.; Squires, A.; Ferguson, A. L. Learned Reconstruction of Protein Folding Trajectories from Noisy Single-Molecule Time Series. *Journal of chemical theory and computation* **2023**, *19*, 4654–4667.
- (57) Yang, S.; Nam, J.; Dietschreit, J. C.; Gómez-Bombarelli, R. Learning Collective Variables with Synthetic Data Augmentation through Physics-Inspired Geodesic Interpolation. *Journal of Chemical Theory and Computation* **2024**, *20*, 6559–6568.
- (58) Salvalaglio, M.; Tiwary, P.; Parrinello, M. Assessing the reliability of the dynamics reconstructed from metadynamics. *Journal of chemical theory and computation* **2014**, *10*, 1420–1425.

- (59) Kang, P.; Trizio, E.; Parrinello, M. Computing the committor with the committor to study the transition state ensemble. *Nature Computational Science* **2024**, 1–10.
- (60) Kaminsky, K. S. Confidence intervals for the exponential scale parameter using optimally selected order statistics. *Technometrics* **1972**, *14*, 371–383.
- (61) Debnath, J.; Invernizzi, M.; Parrinello, M. Enhanced sampling of transition states. *Journal of chemical theory and computation* **2019**, *15*, 2454–2459.
- (62) Herringer, N. S.; Dasetty, S.; Gandhi, D.; Lee, J.; Ferguson, A. L. Permutationally Invariant Networks for Enhanced Sampling (PINES): Discovery of Multimolecular and Solvent-Inclusive Collective Variables. *Journal of Chemical Theory and Computation* **2023**, *20*, 178–198.
- (63) Dietrich, F. M.; Advincula, X. R.; Gobbo, G.; Bellucci, M. A.; Salvalaglio, M. Machine learning nucleation collective variables with graph neural networks. *Journal of Chemical Theory and Computation* **2023**, *20*, 1600–1611.
- (64) Zou, Z.; Wang, D.; Tiwary, P. Graph Neural Network-State Predictive Information Bottleneck (GNN-SPIB) approach for learning molecular thermodynamics and kinetics. *arXiv preprint arXiv:2409.11843* **2024**,
- (65) Rizzi, V.; Aureli, S.; Ansari, N.; Gervasio, F. L. OneOPES, a combined enhanced sampling method to rule them all. *Journal of Chemical Theory and Computation* **2023**, *19*, 5731–5742.
- (66) Ray, D.; Rizzi, V. Enhanced Sampling with Sub-optimal Collective Variables: Reconciling Accuracy and Convergence Speed. **2024**,
- (67) Bonomi, M.; Bussi, G.; Camilloni, C.; Tribello, G. A.; Banáš, P.; Barducci, A.; Bernetti, M.; Bolhuis, P. G.; Bottaro, S.; Branduardi, D.; Capelli, R.; Carloni, P.; Ceriotti, M.; Cesari, A.; Chen, H.; Chen, W.; Colizzi, F.; De, S.; De La Pierre, M.;

Donadio, D.; Drobot, V.; Ensing, B.; Ferguson, A. L.; Filizola, M.; Fraser, J. S.; Fu, H.; Gasparotto, P.; Gervasio, F. L.; Giberti, F.; Gil-Ley, A.; Giorgino, T.; Heller, G. T.; Hocky, G. M.; Iannuzzi, M.; Invernizzi, M.; Jelfs, K. E.; Jussupow, A.; Kirilin, E.; Laio, A.; Limongelli, V.; Lindorff-Larsen, K.; Löhr, T.; Marinelli, F.; Martin-Samos, L.; Masetti, M.; Meyer, R.; Michaelides, A.; Molteni, C.; Morishita, T.; Nava, M.; Paissoni, C.; Papaleo, E.; Parrinello, M.; Pfaendtner, J.; Piaggi, P.; Piccini, G. M.; Pietropaolo, A.; Pietrucci, F.; Pipolo, S.; Provasi, D.; Quigley, D.; Raiteri, P.; Raniolo, S.; Rydzewski, J.; Salvalaglio, M.; Sosso, G. C.; Spiwok, V.; Šponer, J.; Swenson, D. W.; Tiwary, P.; Valsson, O.; Vendruscolo, M.; Voth, G. A.; White, A. Promoting transparency and reproducibility in enhanced molecular simulations. *Nature Methods* 2019 16:8 **2019**, 16, 670–673.

TOC Graphic

