



Cite this: *Digital Discovery*, 2025, 4, 211

A graph neural network-state predictive information bottleneck (GNN-SPIB) approach for learning molecular thermodynamics and kinetics†

Ziyue Zou,^a Dedi Wang^{bc} and Pratyush Tiwary  ^{acd}

Molecular dynamics simulations offer detailed insights into atomic motions but face timescale limitations. Enhanced sampling methods have addressed these challenges but even with machine learning, they often rely on pre-selected expert-based features. In this work, we present a Graph Neural Network-State Predictive Information Bottleneck (GNN-SPIB) framework, which combines graph neural networks and the state predictive information bottleneck to automatically learn low-dimensional representations directly from atomic coordinates. Tested on three benchmark systems, our approach predicts essential structural, thermodynamic and kinetic information for slow processes, demonstrating robustness across diverse systems. The method shows promise for complex systems, enabling effective enhanced sampling without requiring pre-defined reaction coordinates or input features.

Received 27th September 2024
Accepted 27th November 2024

DOI: 10.1039/d4dd00315b
rsc.li/digitaldiscovery

1 Introduction

Molecular dynamics (MD) simulations are widely used in computational research, offering detailed spatial and temporal resolution of atomic motions. However, standard MD faces a timescale challenge, as processes of practical interest can take months or years of computer time to simulate. To tackle this, enhanced sampling methods have been developed, but most of these approaches require collective variables (CVs) to effectively capture key system information.¹ These CVs are typically based on physical insights or experimental data, yet constructing them becomes challenging when transitions are unknown or difficult to sample.

In recent years, machine learning methods have been introduced for this purpose, enabling an automated representation learning framework that can enable CV discovery and enhanced sampling.^{1–15} However, many of these approaches still require hand-crafted expert-based features as input for the ML model. To overcome this limitation, in this work, we build upon the State Predictive Information Bottleneck (SPIB) method,^{16,17} enabling it to learn directly from atomic coordinates instead of relying on hand-crafted expert-based features. SPIB, a variant of the reweighted autoencoded variational Bayes for enhanced sampling (RAVE) method,¹⁸ is a machine learning technique for

dimensionality reduction under a semi-supervised framework. Its structure is based on a time-lagged variational autoencoder where the encoder learns a low-dimensional representation by approximating the posterior distribution of the latent variables given the input features at time t . Unlike in a traditional autoencoder where the decoder focuses on reconstructing the input, the SPIB decoder is trained to predict the state labels in the future $t + \Delta t$ from the latent variables.

The performance of the original SPIB approach is significantly influenced by two key factors. First is the quality of the sampled trajectory, which ideally should include back-and-forth transitions between target metastable states. This requirement is difficult to meet in complex systems with transitions occurring on long timescales, where alternative approaches like using collections of short MD trajectories initiated at initial and final states can be considered.^{19–21}

The second challenge, which we focus on in this manuscript, is that the input variables for the SPIB model are typically derived from prior knowledge of the system, such as expert-based metrics like root mean square deviation (RMSD),²² radius of gyration (R_g), coordination number,²³ and Steinhardt order parameters.²⁴ However, these hand-crafted variables often lack transferability across systems, necessitating parameter tuning to construct effective CVs.²⁵

To address this challenge, efforts have been made to construct ML-based CVs from elementary variables like pairwise distances,^{26–30} yielding valuable insights. However, as noted in ref. 28 and 29, the stability of these methods deteriorates when the input dimension exceeds 100 when biasing, making them less suitable for many-body systems or large biomolecules. Additionally, this approach does not resolve symmetry issues common, for instance, in materials science. Although pairwise

^aDepartment of Chemistry and Biochemistry, University of Maryland, College Park 20742, USA. E-mail: ptiwary@umd.edu

^bBiophysics Program, University of Maryland, College Park 20742, USA

^cInstitute for Physical Science and Technology, University of Maryland, College Park 20742, USA

^dUniversity of Maryland Institute for Health Computing, Bethesda 20852, USA

† Electronic supplementary information (ESI) available. See DOI: <https://doi.org/10.1039/d4dd00315b>



distances are invariant to translation and rotation, the learned latent variables in typical neural networks are not permutation-invariant, meaning reordering input features can alter the CV value for the same configuration. While symmetry functions can be introduced to enforce invariance, this is often time-consuming.^{25,31,32} Moreover, ML methods with multilayer perceptrons (MLPs) lack transferability to systems of different sizes, whereas GNNs can accommodate systems of varying sizes.

Graph neural networks (GNNs) have recently gained attention because of their effectiveness in constructing representations across various applications, particularly in materials science, due to their inherent permutation invariance.^{7,17,33-44} In this work, we extend the SPIB framework by incorporating a GNN head with different graph convolutional layers. This specifically addresses the limitations of the original SPIB algorithm which needed physical-inspired hand-crafted input variables, while here a meaningful representation is learned on the fly with invariant pairwise distance variables. This enables us to apply the same framework with similar architectures across diverse systems without relying on system-specific expert knowledge, making it more broadly applicable. We tested our enhanced method on three representative model systems using our machine-learned CVs, which we call GNN-SPIB CVs. The systems are the Lennard-Jones 7 cluster, where permutation symmetry is crucial, and alanine dipeptide and tetrapeptide, where high-order representations such as torsion angles are typically required. With straightforward graph construction and basic features, our approach successfully learns meaningful and useful representations across all systems, using three representative graph layers to show the flexibility of our proposed framework. Additionally, the latent variables derived from our method provided thermodynamic and kinetic estimates comparable to those obtained using metadynamics-based methods that bias physically inspired expert-based CVs. This demonstrates the robustness and adaptability of our approach in overcoming previous challenges.

2 Methods

In this section, we provide an overview of the different techniques we use to learn latent geometric representations and assess their quality through enhanced sampling. Details of the enhanced sampling methods,⁴⁵⁻⁴⁷ model system setups, neural network training protocols, and definition of the expert-based collective variables are provided in the ESI.†

2.1 State predictive information bottleneck

The state predictive information bottleneck (SPIB) developed by Wang and Tiwary¹⁶ is a variant of the Reweighted Autoencoded Variational Bayes for Enhanced sampling (RAVE) method.⁴⁸ RAVE allows one to learn a meaningful representation with a variational autoencoder (VAE) framework from biased or unbiased data in the form of a time-series Y_{t_n} , comprising generally of several expert-selected features as a function of time. Building on this theme, SPIB was introduced as a more interpretable and robust model within the RAVE family,

designed to learn a meaningful low-dimensional representation that accounts for the metastable states of most molecular systems, where the system spends extended periods undergoing fluctuations. Instead of predicting the details of these fluctuations Y within the states, it is then more crucial to predict which metastable state S the system will be in after a time delay Δt . To reflect this, the objective function, \mathcal{L} in SPIB, aims to predict future state labels $S_{t_n+\Delta t}$:

$$\mathcal{L} = I(z, S_{t_n+\Delta t}) - \beta I(Y_{t_n}, z), \quad (1)$$

where $I(x, y) = \int p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dx dy$ is the mutual information between variables x and y ; z is the low-dimensional latent representation; and $S_{t_n+\Delta t}$ is the state label after a time delay Δt . β is a tunable hyperparameter which controls regularization *versus* prediction, while tuning the time delay Δt controls the extent of temporal coarse-graining of the dynamics as learned by SPIB. After an initial trial assignment of states which can be very approximate, SPIB learns both the number of metastable states and their locations in the high-dimensional feature space Y . The number of metastable states generally reduces as a function of the time delay Δt . Introducing this metastability based prediction task makes SPIB latent space physically meaningful as now they correspond to the slow degrees of freedom, making the learned representation more interpretable. This approach has been successfully applied to enhancing sampling of transitions in complex systems^{28,49} and to approximating the 50% committor surface between metastable states.^{21,50}

2.2 Graph and graph neural network

Graph data, denoted as $G(V, E)$, consist of a set of nodes V and edges E , each carrying specific geometric information. Typically, three key components define this graph structure, which we summarize here for the sake of completeness:

(1) Node features (X_i): each node $i \in V$ is associated with a feature vector X_i , which encapsulates the intrinsic properties or characteristics of the node.

(2) Edge indices (i, j): the relationships between nodes are represented by edge indices (i, j) in an adjacency matrix $a_{ij} \in \{0, 1\}$, defining the neighborhood structure and indicating direct connections between nodes.

(3) Edge features (L_{ij}): these describe the nature of the connection between nodes i and j , capturing attributes like weight, distance, or type.

In order to make meaningful predictions with inputs in graph objects, a special type of neural network, known as a graph neural network (GNN), is needed. Given the complexity of typical graph data, layers in the GNN are designed and trained with care. Within each graph layer, message-passing can break into three sequential steps:

(1) At layer l , message m between node i and each neighbor $j \in \mathcal{N}_i$, defined by the adjacency matrix, is computed *via* a message function (eqn (2)). The embedding $h^0 = X$ at layer $l = 0$ will be updated *via* each message passing operation.

$$m_{ij}^l = \mathbf{msg}\{h_i^l, h_j^l, L_{ij}\} \quad (2)$$



(2) This information from neighbors is then combined *via* an aggregate function (eqn (3)), which can be as simple as summation and averaging.

$$m_i^l = \text{agg}\{m_{ij}^l\} \quad (3)$$

(3) Lastly, the collected information from step 2 is integrated into the node features through an update function (eqn (4)): ⁵¹

$$h_i^{l+1} = \text{upd}\{h_i^l, m_i^l\} \quad (4)$$

Note that preserving the permutation-invariant property in GNNs requires careful design of the above operations. In this work, we ensured that our models maintain this property by including invariant input features and allowing invariant operation during message passing in graph layers.

2.3 Graph-based SPIB

As introduced above, low-dimensional representations learned by an SPIB model are capable of learning the number of metastable states and their locations and capturing the slow processes that govern transitions between them. These joint capabilities differentiate SPIB from other dimensionality reduction schemes. However, SPIB still needs a dictionary of features Y which can collectively, in high dimensions, demarcate different metastable states. Defining these variables can be challenging when studying complex systems, as they often rely on prior system knowledge. Using poorly distinguishable input variables may degrade the performance of enhanced sampling

methods that use SPIB-trained variables. To fully automate the representation learning in SPIB and eliminate the need for expert-defined input variables, we integrate a graph neural network head into the existing SPIB model.

Integrated together, the architecture is summarized in Fig. 1, and we refer to it as GNN-SPIB. While it is overall akin to the original SPIB framework,¹⁶ the input trajectory is now a graph comprising geometric representations of the simulation cell at each time frame, whereas in SPIB, the input comprised expert-based features. Following the same notation, we denote the input graph as G_{t_n} at time frame t_n , and therefore the loss function in eqn (1) is rewritten as:

$$\mathcal{L} = I(z, S_{t_n+\Delta t}) - \beta I(G_{t_n}, z) \quad (5)$$

Additionally, the high-dimensional nature of graph data provides a better description of the simulation cell than expert-based features. In this work, SPIB and graph models were developed with Pytorch⁵² and Pytorch geometric⁵³ packages, respectively. Given the diversity of graph layers, we do not want to limit ourselves to certain graph layers or GNN architecture, and therefore, we selected three representative graph message passing layers when studying the three model systems to present the general applicability of our proposed framework. We believe that the selection of graph message passing layers can be rather flexible given the fact that the basic principle, message-passing operation is permutational-invariance, is enforced in most of these graph layers. The only design parameter here that requires care is that informative graph

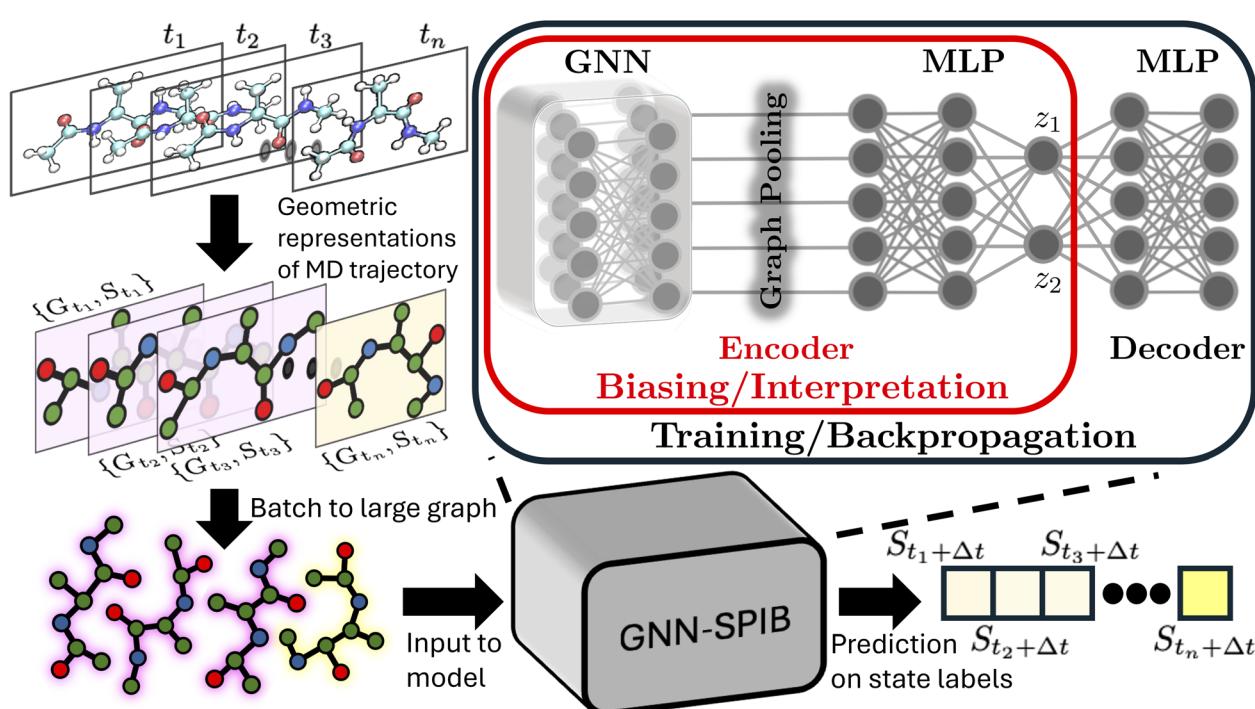


Fig. 1 Schematic of the workflow proposed in this work. Trajectories from unbiased/biased simulation are converted into timeseries graph data. The batched large graph is fed into graph neural networks. The GNN-SPIB model is then trained to predict the state labels of the time frame in lag time Δt as introduced in the original SPIB pipeline (box in black). The biasing variables (*i.e.*, z_1 and z_2) are then used in enhanced sampling methods (box in red).



layers are always at a higher computational cost which may largely slow down enhancing sampling methods.

Workarounds similar to ours have been used in other representation learning methods. The closest related method is the variational approach for Markov processes (VAMP) net, a machine learning architecture that constructs a Markov state model (MSM) and optimizes the VAMP-2 score derived from the single value decomposition of the corresponding Koopman operator.⁵⁴ In VAMPnet, MD configuration coordinates are used as input to the machine learning model. To address symmetry issues, alignment of configurations is necessary, which is common in computational studies of biological systems but less applicable to materials science. As an alternative, a graph representation was introduced into the standard VAMPnet architecture.^{44,55,56} All GNN layers chosen in this work are $E(3)$ -invariant GNNs, while a more data-efficient equivariant GNN representation learning scheme was introduced recently in ref. 44.

3 Results and discussion

We evaluate the ability of the GNN-SPIB low-dimensional latent representations to enhance sampling for three model systems. For all three systems, we perform well-tempered metadynamics (WTmetadD) (see the ESI† for a detailed introduction to this method) to quantify the quality of calculated free energy and infrequent metadynamics (imetaD) (see the ESI† for details about this method) to calculate kinetics. The three systems are

Lennard-Jones 7 (LJ7, in Sec. 3.1), alanine dipeptide (Sec. 3.2) and alanine tetrapeptide (Sec. 3.3). As a general pipeline for the three systems, we first train the models with data collected from short MD simulations at higher temperatures in which all targeted metastable states are visited but their sampling is incorrect/unconverged especially for the lower temperature of interest. We then perform WTmetaD simulations biasing along the GNN-SPIB latent variables, compute the free energy difference between states, and compare with results from much longer unbiased MD simulations and WTmetaD simulations biasing conventional expert-based CVs. After thermodynamic measurements with WTmetaD, we also collect kinetic information with imetaD simulations.

3.1 Lennard-Jones 7

As suggested by its name, Lennard-Jones 7 (LJ7) consists of a cluster of 7 Lennard-Jones particles in 2-d. It is considered one of the simplest model systems for colloidal rearrangements, where translational-, rotational-, and permutational-symmetry problems are encountered, with metastable states. The free energy landscape and state-to-state dynamics of LJ7 are well-studied, making it an excellent model system for benchmarking enhanced sampling methods of rare events.⁵⁷

Fig. 2 summarizes the results when biasing along the 2-dimensional GNN-SPIB latent variables, z_1 and z_2 . The input graphs come from snapshots of a trajectory of 1×10^7 steps at a temperature of $k_B T = 0.2\epsilon$. These are composed of identical

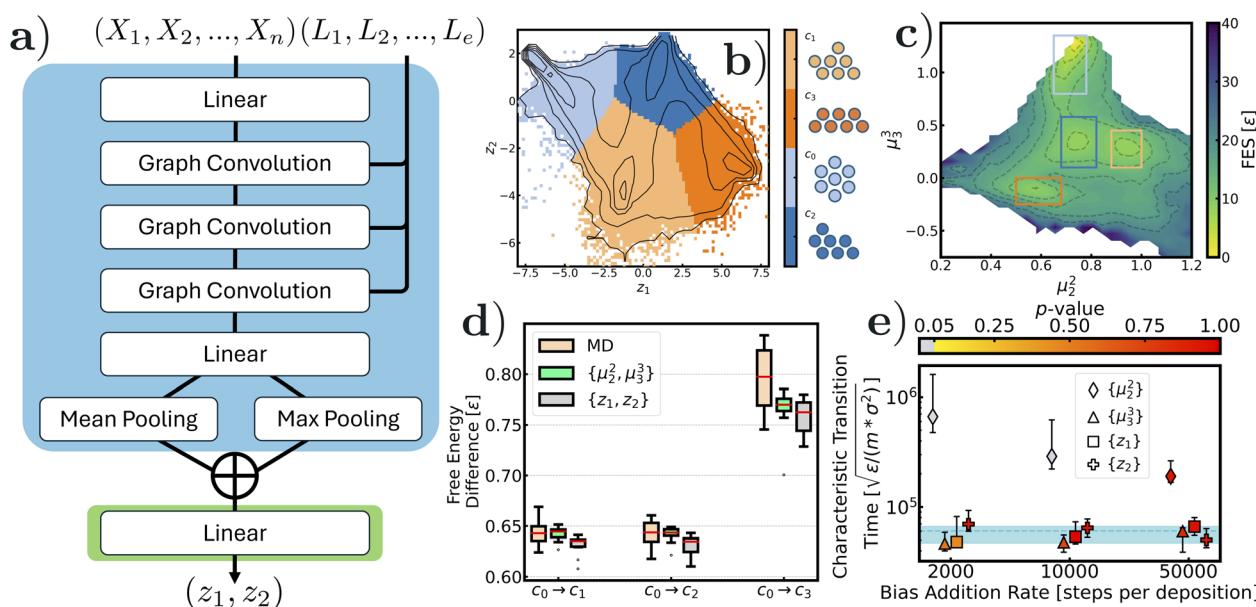


Fig. 2 Summary of WTmetaD simulation results biasing along machine learned reaction coordinates, z_1 and z_2 , in the Lennard-Jones 7 system. (a) A schematic of how the 2-d reaction coordinates, z_1 and z_2 , the output of the encoder, are computed with node features $\{V_n\}$ and edge features $\{L_e\}$ where \oplus denotes the concatenate operation. (b) State labels predicted by the model in RC space projected along the training data collected at $k_B T = 0.2\epsilon$. The highest contour line is at 10ϵ and each of the lines is separated by 2ϵ . (c) Reweighted free energy surface of WTmetaD using $\{z_1, z_2\}$ at $k_B T = 0.1\epsilon$ projected onto expert-based CV space, μ_2^2 and μ_3^3 , with state definitions in colored boxes. (d) Box plots of free energy differences from (c) between sampled metastable states comparing conventional long MD and WTmetaD biasing expert-based CVs. (e) Characteristic transition times of $c_0 \rightarrow c_3$ at $k_B T = 0.1\epsilon$ estimated by imetaD simulations using expert-based and machine learned RCs. Benchmark is drawn from standard MD simulation in cyan. The shaded region and error bar correspond to the 95% confidence interval. Colors in markers indicate the p -value from the K-S test, where a p -value less than 0.05 suggests that the result is unreliable.



node attributes $\{X_n\} = \{1, 1 \cdots 1\}$ and pair distance edge features $\{L_e\}$. The graph convolution layers, which are considered as one of the simplest graph layers, are directly adopted from ref. 58 and graph embeddings are pooled with mean and max operators (see Fig. 2a in blue). The hidden embedding of graph layers is then fed into the SPIB model (Fig. 2a in green; we have provided detailed information about the training process in the ESI†) for extrapolating dynamic information. Fig. 2b shows the SPIB converged state labels as colored regions where we observe four distinct, local minima corresponding to the four well known structures for the LJ-7 system. In particular, these 4 configurations are c_0 (hexagon), c_1 (capped parallelogram 1), c_2 (capped parallelogram 2), and c_3 (trapezoid) (see schematics next to Fig. 2b and the colorbar). Given the fact that the predicted state labels are correctly assigned to 4 distinct energy minima in GNN-SPIB latent space, we believe that the trained model is able to classify the configuration of the LJ7 cluster when providing the corresponding geometric representation in nodes and edges, and therefore, the encoded latent representation from GNN-SPIB can be used as a biasing variable in metadynamics simulations.

To verify this, WTmetaD simulations were performed at $k_B T = 0.1\epsilon$ biasing the 2-d reaction coordinates, z_1 and z_2 , and the resulting free energy surfaces are shown in (Fig. 2c). For better evaluation of the sampling quality, we projected the free energy surface onto the more meaningful space comprising the second and third moments of the coordination numbers, introduced previously in ref. 59. These expert-based CVs have been used previously to study this system⁶⁰ and thus provide a good way to test the quality of samples generated from biasing along the GNN-SPIB, which is physically less meaningful. Four distinct local minima were sampled, which is a strong indicator of good sampling quality for this system. To further evaluate the quality of the sampling, we tabulated the free energy differences between the sampled configurations and the initial state, c_0 . As a benchmark, we performed an additional 10 independent, unbiased long MD simulations, 10 WTmetaD simulations biasing the conventional CV set $\{\mu_2^2, \mu_3^3\}$, and 10 WTmetaD simulations biasing the machine-learned 2D CV $\{z_1, z_2\}$. The MD simulations lasted for 1×10^9 steps, and the WTmetaD simulations lasted for 1×10^8 steps. Inspection of Fig. 2d reveals that WTmetaD simulations using GNN-SPIB CV produce results comparable to those using the expert-crafted moments of coordination number CVs (refer to the ESI† for numerical values). Notably, no information about the coordination was directly provided as input to the model and only the pairwise distances of neighboring nodes were used. We further evaluate the GNN-SPIB against conventional CVs by computing their correlation coefficients. The results (see the ESI†) show that both z_1 and z_2 have a stronger correlation to μ_3^3 than to μ_2^2 .

As an even more demanding test, we ascertained the quality of the GNN-SPIB CV in obtaining accurate kinetics through imetaD calculations. In this task, we performed kinetic measurements by estimating the transition time of the slowest transition from initial state c_0 to c_3 using the imetaD method at $k_B T = 0.1\epsilon$. We performed these 1-d imetaD simulations separately biasing the GNN-SPIB z_1 and z_2 . We benchmarked on

characteristic transition time from reference unbiased MD simulations (dashed line with shaded errors in Fig. 3e). We also provide results of 1-d imetaD simulations using expert-based CVs μ_2^2 and μ_3^3 . In particular, since the 1-d projection along μ_3^3 suggests that the states c_0 and c_3 are well-separated, we expect μ_3^3 to be a better expert-based CV for biasing relative to μ_2^2 as it fails to distinguish c_0 from c_3 . Using the imetaD method, the characteristic transition time remained robust and accurate under various bias addition frequencies when μ_3^3, z_1 , and z_2 variables were used. Specifically, we obtained transition times of 53 829 (95% confidence interval (CI): [46 282, 73 303]) and 64 504 (95% CI: [52 986, 77 703]) $\sqrt{\epsilon/(m \times \sigma^2)}$ when biasing along the GNN-SPIB z_1 and z_2 , respectively, with bias added every 1×10^4 steps. For reference, the transition time from long unbiased MD for the $c_0 \rightarrow c_3$ transition was estimated to be 60 630.31 (95% CI [46 945, 66 598]) $\sqrt{\epsilon/(m \times \sigma^2)}$. The numerical values of 95% CI and the p -values of all performed simulations are provided in the ESI.† However, when μ_2^2 was used in imetaD simulations, the characteristic transition time was off by at least one order of magnitude compared to the benchmark value, with a p -value less than 0.05 (markers in grey in Fig. 2e). In summary, the thermodynamic and kinetic evidence suggests that either of our GNN-SPIB CVs shows results comparable to conventional expert-based CVs when used in enhanced sampling.

3.2 Alanine dipeptide

While the rearrangement of the LJ7 cluster is considered a simple model of colloidal system dynamics, the alanine dipeptide molecule serves as a popular toy model for biomolecular conformational changes. Here, we focus on transitions in a vacuum between three conformers: C5, C7_{eq}, and C7_{ax}. The set of expert-based CVs commonly used in enhanced sampling methods for this system includes the dihedral angles ϕ and ψ (Fig. 3a).^{27,62} In our model, unlike the previous example, the three elements carbon, nitrogen and oxygen in alanine dipeptide are one-hot encoded in their node features. Once again, we leverage the computational efficiency of using inter-atomic distances as edge features, similar to the LJ7 model. Although conformational changes in biomolecules like alanine dipeptide are generally described by high-order representations such as torsion angles, lower-dimensional descriptors for enhanced sampling methods can be learned through machine learning-based dimensionality reduction on inter-atomic distances.^{26,27,63,64} To increase the expressiveness of the model, we applied a basis function to slice the edge features (refer to the ESI† for details on this operation).⁴³

The overall architecture remains the same as in the study of the LJ7 cluster, but we replaced the graph convolution layers with more informative graph attention network (GAT) convolution layers, to show the robustness of our proposed framework. We have provided information about the training process and the hyperparameters used in the ESI.^{65†} However, this does not necessarily mean that the graph convolutional layer, which was adopted in the LJ7 system, is not applicable here. We trained another GNN-SPIB with the exact architecture shown in (Fig. 2a) on alanine dipeptide data and the results showing the



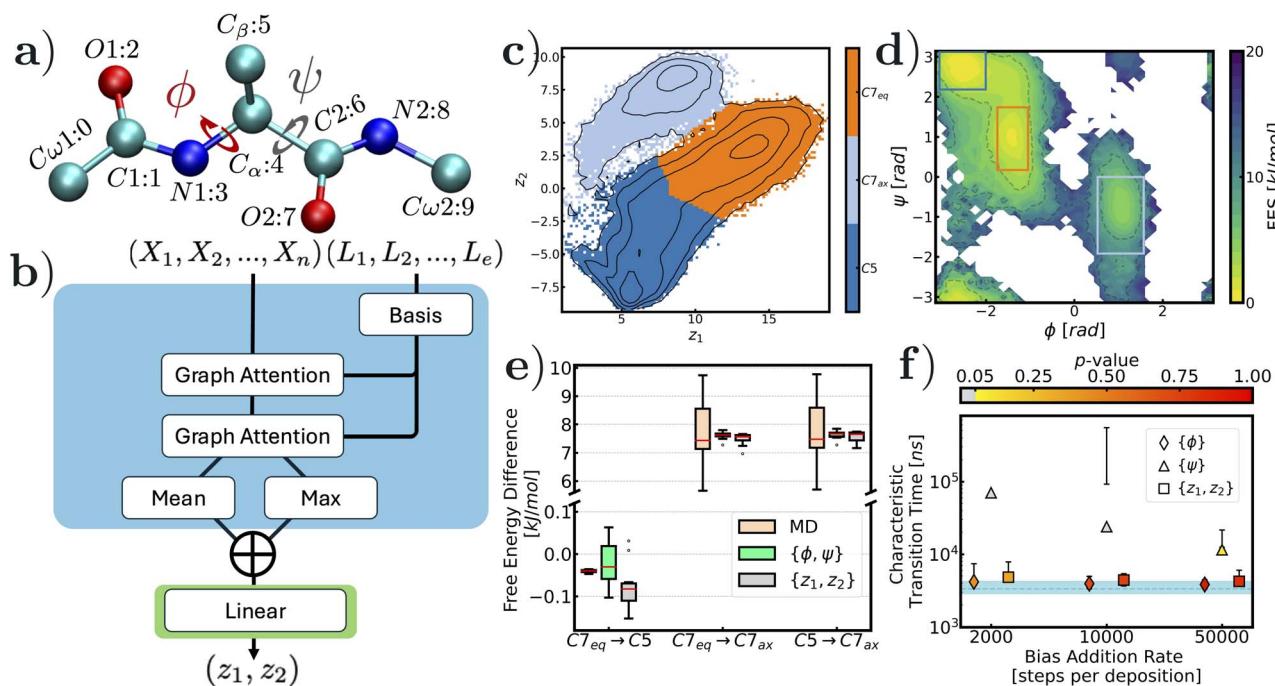


Fig. 3 Summary of WTmetaD simulation results for the alanine dipeptide system: (a) representation of the alanine dipeptide molecule⁶¹ with definition to expert-based CVs, ϕ and ψ . Graph representation is constructed with only heavy atoms and atomic labels are followed through the assigned node index to graphs; (b) a schematic of how the reaction coordinate, $\{z_1, z_2\}$, is computed with node $\{V_i\}$ and edge $\{L_e\}$ features; (c) state label predictions in different colors from the model decoder with contour lines separated by 3 kJ mol^{-1} ; (d) reweighted free energy surface biasing the machine learned RC at 300 K using $\{\phi, \psi\}$ projection with conformer definitions in boxes; (e) free energy differences with the state defined in (d) under different sampling schemes; and (f) kinetic measurements of the transition from $(C7_{eq}, C5)$ to $C7_{ax}$ at 300 K with imetaD simulation. Dashed line in cyan is the benchmark MD simulation and marker points are results from imetaD using different RCs. The shaded region and error bars are the 95% confidence intervals. p -values from the K-S test for imetaD simulations are reflected by the colors and when the p -value is less than 0.05 (in grey), the result is unreliable.

model's ability in distinguishing all metastable states are presented in the ESI.† Three distinct minima are shown when projected onto the learned latent space with input data sampled at $T = 400 \text{ K}$ (Fig. 3c), and the reweighted free energy surface in the ϕ, ψ space from WTmetaD simulations using the z_1 and z_2 variables is shown in Fig. 3d, where all targeted states are well-sampled. The free energy differences between individual conformers were tabulated and benchmarked with brute-force MD simulations and WTmetaD simulations using ϕ, ψ dihedrals (see Fig. 3e). The results are in good agreement, with discrepancies of less than 1 kJ mol^{-1} . The numerical values of the free energy difference are reported in the ESI.† In addition, we see a strong correlation of z_1 with ϕ and ψ variables and only a moderate correlation of z_2 with these two torsion angles (see the ESI† for the scatter plot).

We then move to the more challenging validation of kinetics through enhanced sampling. The evaluation metric of kinetics focused on the slowest transition, $C7_{eq} \rightarrow C7_{ax}$, in alanine dipeptide, and the results are summarized in Fig. 3f. The characteristic transition time was estimated to be 3.340 (95% CI [2.857, 4.259]) μs , which is in accordance with ref. 66. The values in Fig. 3f suggest that both the 1-d expert-based CV ϕ and our GNN-SPIB CV $\{z_1, z_2\}$ are effective at the obtained accurate reweighted kinetics (see the ESI† for complete reports on kinetics measurements). ψ alone is known to be a poor CV for

this system,⁶⁶ which is reflected in the inaccurate kinetics when biasing along ψ irrespective of the frequency of bias deposition. Notably, unlike the LJ7 case, we directly performed imetaD simulations on the 2-d CV set, as 1-d imetaD simulations biasing either z_1 or z_2 did not yield good estimations of transition times. This is unsurprising, as the GNN-SPIB latent variables were set and trained in 2-d, and thus the complete information of this complex conformational change was likely not captured by z_1 or z_2 alone. At a 50 ns^{-1} bias addition rate (i.e. 1 Gaussian deposition every 10 000 steps), the estimated transition time is 4.424 (95% CI [3.662, 5.384]) μs , which is in agreement with that estimated in MD simulations (see all estimated values at different deposition rates in the ESI†).

We further evaluated the attention coefficients from the GAT layers, which reflect how information is exchanged during message passing, providing insights into the system (see the ESI† for complete attention matrices). In the first graph attention convolution layer, long-range interactions describing global molecular orientations, such as $C_{\omega}1-C_{\omega}2$ and $N1-N2$ distances, drew the model's attention. In contrast, local arrangements (e.g., $O1-N1$ and $N1-C_{\beta}$ distances) had high attention weights in the second graph attention layer (see Fig. 3a for atom labels). This suggests that, without high-order representations of the conformation, transitions between conformers are decoded sequentially from far to near using



pairwise distances. Additionally, edges with large attention weights are those connected to the N1 atom, which is used to define the torsion angle ϕ , indicating the importance of the ϕ angle over the ψ angle.

3.3 Alanine tetrapeptide

For our third and final example, we study conformational changes in a much more complicated model system, namely alanine tetrapeptide in a vacuum. For this system, there exist at least 8 metastable states in total. For these states, we follow the notations from ref. 67, numbering these states as in ref. 67 as s_i , $i = 1 \dots 8$. In particular, to capture the intricacies of alanine tetrapeptide conformational changes, six dihedral angles are considered important: ϕ_1 , ϕ_2 , ϕ_3 , ψ_1 , ψ_2 , and ψ_3 (see Fig. 4b for their definitions). Using a similar workflow as before, the alanine tetrapeptide molecule is first converted into graph objects, *i.e.*, all hydrogens are removed and C, N, and O atoms are retained during the graph construction. The nodes in the graph are set to be fully connected, and edge features are defined as interatomic distances. We adopted a skip connection scheme that allows information from all graph layers to flow to the pooling operator, thereby improving the model's

expressiveness in capturing features from all metastable states. Additionally, as shown in Fig. 4a, we considered all three typical pooling operations: mean, max, and sum in this example.

Once again, the graph layers were switched and chosen to be an expressive Gaussian Mixture model (see the ESI† for details about the model).⁶⁸ Similarly, we trained a GNN-SPIB with the architecture shown in (Fig. 2a) to present the generalizability of our approach, and the results (see the ESI†) suggest that the model is able to identify metastable states in alanine tetrapeptide. Our input training data are from a 1 μ s-long MD simulation at 400 K, where s_6 is barely sampled. These input training data are projected onto the trained latent space $\{z_1, z_2\}$ in Fig. 4c. As shown by the color coding, the model successfully learned 7 out of 8 metastable states, though the states s_6 and s_7 remain indistinguishable from each other. This is due to the lack of samples from s_6 during training. Nevertheless, performing WTmetaD simulations using $\{z_1, z_2\}$ for 200 ns at a lower temperature of 350 K drives the system to visit all 8 metastable states for alanine tetrapeptide. To demonstrate the quality of our sampling, we project the trajectory onto the $\{\phi_1, \phi_2, \phi_3\}$ space and label all target states in Fig. 4d. The free energy differences between all 8 states are computed and shown in

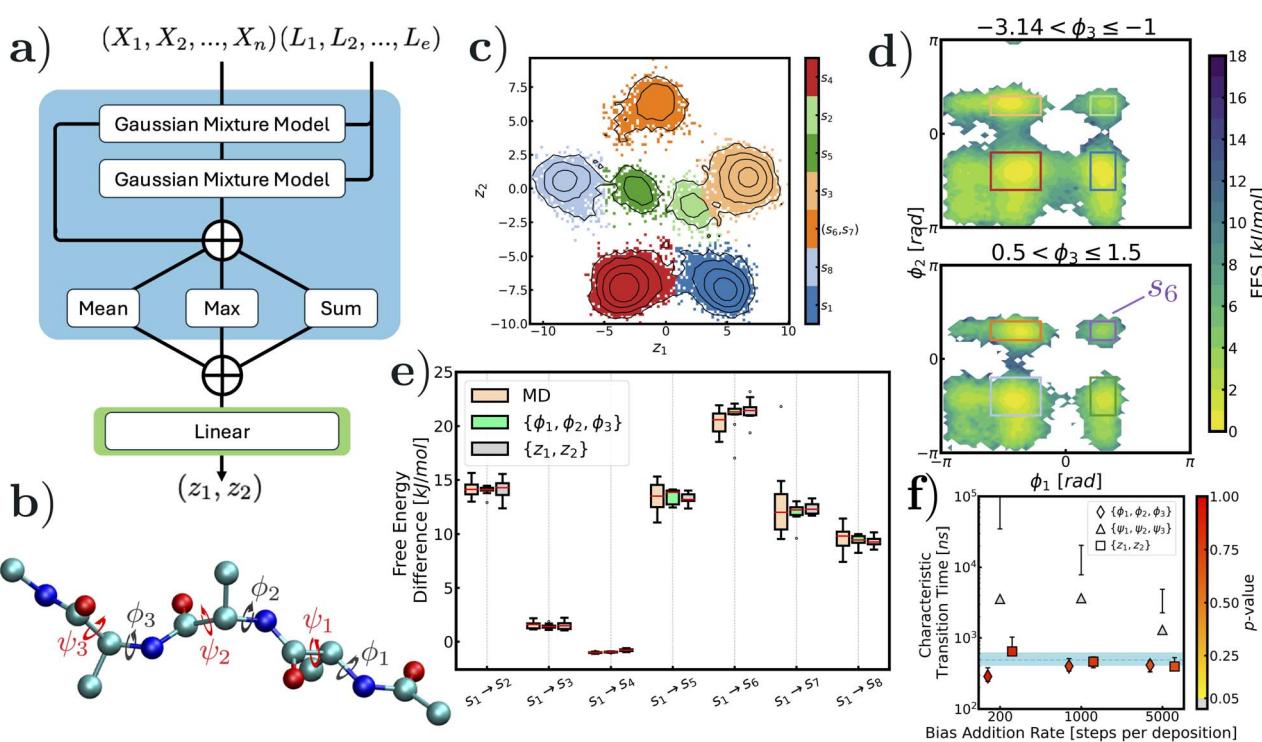


Fig. 4 Summary of WTmetaD simulation results in the alanine tetrapeptide system: (a) a schematic of reaction coordinate construction with a combination of embeddings of each graph convolution layers via skip connections before graph-level pooling operations; (b) representation of the alanine tetrapeptide molecule⁶¹ with definitions of characteristic dihedral angles, ϕ_1 , ϕ_2 , ϕ_3 , ψ_1 , ψ_2 , and ψ_3 and only heavy atoms are involved during graph construction; (c) the learned latent variable space $\{z_1, z_2\}$ with state labels predicted by the model on training data and the free energy surface with contours separated by 2 kJ mol^{-1} ; (d) reweighted free energy surface of WTmetaD simulations using 2-d $\{z_1, z_2\}$ variables at 350 K projected onto $\{\phi_1, \phi_2, \phi_3\}$ space; (e) tabulated free energy differences between all conformers from brute force MD simulations, WTmetaD simulations biasing $\{\phi_1, \phi_2, \phi_3\}$, and WTmetaD simulations biasing $\{z_1, z_2\}$; and (f) characteristic transition times of $s_1 \rightarrow s_7$ measured by imetaD using different RCs at 400 K. Dashed line in cyan is the benchmark MD simulation and marker points are results from imetaD using different RCs. The shaded region and error bars are the 95% confidence intervals. p -values from the K-S test for imetaD simulations are reflected by the colors and when the p -value is greater than 0.05, the estimation is reliable.



a box plot (Fig. 4e), along with two benchmark methods: standard long MD simulations and 3-d WTmetaD simulations using $\{\phi_1, \phi_2, \phi_3\}$ variables. The free energy differences converge to well-defined values, aligning with those from the benchmark methods (Fig. 4e) (see the ESI† for details). From evaluating the correlation between machine learned variables and conventional variables, we find that z_1 and z_2 show correlations with only a few torsion angles such as ϕ_3 and ψ_2 . This suggests that our machine learning model learned different descriptors of conformational changes in the system compared to conventional knowledge-based variable – torsion angles, such discrepancies have been detailed investigated in ref. 69.

While performing 3-d biasing along $\{\phi_1, \phi_2, \phi_3\}$ space reveals accurate kinetic measurements, imetaD simulations using $\{\psi_1, \psi_2, \psi_3\}$ performed poorly by largely overestimating the transition time under various biasing paces. Values estimated by imetaD simulations biasing the 2-d GNN-SPIB $\{z_1, z_2\}$ variables were in agreement with benchmark values and this variable remained robust to frequent biasing deposition rates. As summarized in the subplot of Fig. 4f, we estimated the characteristic transition time from the unfolded substate s_1 to the folded substate s_7 at 400 K. The timescale was estimated to be 489 (95% CI [409, 612]) ns using reference unbiased MD simulations. Values estimated by imetaD simulations biasing the 2-d GNN-SPIB $\{z_1, z_2\}$ variables were 460 (95% CI [378, 618]) ns with a 500 ns⁻¹ deposition rate (*i.e.* 1 Gaussian deposition every 1000 steps), in agreement with the benchmark values and remained robust to frequent bias deposition rates (see the ESI† for details). Finally, while performing 3-d biasing along the $\{\phi_1, \phi_2, \phi_3\}$ space reveals accurate kinetic measurements, imetaD simulations using $\{\psi_1, \psi_2, \psi_3\}$ performed poorly, significantly overestimating the transition time under various biasing paces. This reflects how one set of expert-based CVs can significantly differ from another in the quality of sampling.

4 Conclusion

While enhanced sampling methods have significantly extended the capabilities of molecular dynamics simulations, identifying optimal coordinates remains an ongoing challenge. Even methods using machine learning, while somewhat automating the process, still rely on expert-based features to be pre-selected. In this work, we introduced a hybrid framework combining graph neural networks and the state predictive information bottleneck, named the GNN-SPIB approach, to automatically learn low-dimensional representations of complex systems, further removing this limitation. This approach allows the model to learn configurations *via* graph layers while capturing system dynamics through the past-future information bottleneck.

We demonstrated the effectiveness of our method on three benchmark systems: the Lennard-Jones 7 cluster, alanine dipeptide, and alanine tetrapeptide. Each system presents distinct challenges in learning meaningful representations, such as the need for permutation invariance in the Lennard-Jones cluster or high-order variables like angles for peptide systems. By applying three representative graph message-

passing layers, we showcased the robustness and flexibility of the proposed framework. Importantly, our results are not confined to specific graph layers or architectures, underscoring the generalizability of this approach across diverse systems.

Biassing these GNN-SPIB in WTmetaD simulations yielded results comparable to those obtained using conventional CVs both for the calculation of free energy surfaces and kinetic transition times. Given the simplicity of the input features, specifically pairwise distances, we believe that this method holds promise for complex systems where optimal reaction coordinates for enhanced sampling methods are not known *a priori*. The major computational cost of the current method remains to be backpropagating low-dimensional representations to high-dimensional atomic coordinates, which can be properly accelerated with GPU-support⁴⁴ and/or by optimizing the design of current models with an optimal number of training parameters. Like other variants in the RAVE family, our learned reaction coordinates can be iteratively improved with better sampling, particularly for complex systems, as noted in ref. 29. Future work could expand this approach by incorporating high-order representations such as angles or spherical harmonics.⁵¹ Additionally, the input data for model training need not be limited to simulations; static data from metastable states, as demonstrated in previous studies, can also lead to meaningful latent representations.^{17,19–21}

Data availability

The codes used for reproducing the simulations are available at PLUMED-NEST <https://www.plumed-nest.org/eggs/24/020/>.

Author contributions

Z. Z., D. W., and P. T. designed research; Z. Z. and D. W. developed algorithms; Z. Z. performed research and analyzed data; Z. Z. wrote the initial draft; and Z. Z., D. W., and P. T. revised the paper.

Conflicts of interest

The authors declare no competing financial interest.

Acknowledgements

This research was entirely supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, CPIMS Program, under Award DE-SC0021009. Z. Z. thanks Zachary Smith and Suemin Lee for help in graph neural nets and infrequent metadynamics simulations. The authors thank Eric Beyerle, Bodhi P. Vani, Shams Mehdi, Akashnathan Arangathan, and Venkata Adury for fruitful discussions. The authors also thank Bodhi P. Vani, Zachary Smith, Vanessa J. Meraz, and Suemin Lee for proofreading the manuscript. P.T. is an investigator at the University of Maryland-Institute for Health Computing, which is supported by funding from Montgomery County, Maryland, and The University of Maryland Strategic Partnership: MPowering the State, a formal collaboration



between the University of Maryland, College Park, and the University of Maryland, Baltimore. We are also grateful to NSF ACCESS Bridges2 (project CHE180053) and the University of Maryland Zaratan High-Performance Computing cluster for enabling the work performed in this study.

References

- 1 S. Mehdi, Z. Smith, L. Herron, Z. Zou and P. Tiwary, Enhanced Sampling with Machine Learning, *Annu. Rev. Phys. Chem.*, 2024, **75**, 38382572.
- 2 F. Noé, A. Tkatchenko, K.-R. Müller and C. Clementi, Machine learning for molecular simulation, *Annu. Rev. Phys. Chem.*, 2020, **71**, 361–390.
- 3 M. Chen, Collective variable-based enhanced sampling and machine learning, *Eur. Phys. J. B*, 2021, **94**, 1–17.
- 4 S. Sarupria, S. W. Hall and J. Rogal, Machine learning for molecular simulations of crystal nucleation and growth, *MRS Bull.*, 2022, **47**, 949–957.
- 5 E. R. Beyerle, Z. Zou and P. Tiwary, Recent advances in describing and driving crystal nucleation using machine learning and artificial intelligence, *Curr. Opin. Solid State Mater. Sci.*, 2023, **27**, 101093.
- 6 H. Jung, R. Covino, A. Arjun, C. Leitold, C. Dellago, P. G. Bolhuis and G. Hummer, Machine-guided path sampling to discover mechanisms of molecular self-organization, *Nat. Comput. Sci.*, 2023, 1–12.
- 7 F. M. Dietrich, X. R. Advincula, G. Gobbo, M. A. Bellucci and M. Salvalaglio, Machine Learning Nucleation Collective Variables with Graph Neural Networks, *J. Chem. Theory Comput.*, 2024, **20**, 1600–1611.
- 8 N. S. M. Herringer, S. Dasetty, D. Gandhi, J. Lee and A. L. Ferguson, Permutationally Invariant Networks for Enhanced Sampling (PINES): Discovery of Multimolecular and Solvent-Inclusive Collective Variables, *J. Chem. Theory Comput.*, 2024, **20**, 178–198.
- 9 O. Elishav, R. Podgaetsky, O. Meikler and B. Hirshberg, Collective Variables for Conformational Polymorphism in Molecular Crystals, *J. Phys. Chem. Lett.*, 2023, **14**, 971–976.
- 10 W. Zhang and C. Schütte, Understanding recent deep-learning techniques for identifying collective variables of molecular dynamics, *PAMM*, 2023, **23**, e202300189.
- 11 E. Prašníkár, M. Ljubić, A. Perdih and J. Borišek, Machine learning heralding a new development phase in molecular dynamics simulations, *Artif. Intell. Rev.*, 2024, **57**, 102.
- 12 L. Müllender, A. Rizzi, M. Parrinello, P. Carloni and D. Mandelli, Effective data-driven collective variables for free energy calculations from metadynamics of paths, *PNAS Nexus*, 2024, **3**, pgae159.
- 13 A. Majumder and J. E. Straub, Machine Learning Derived Collective Variables for the Study of Protein Homodimerization in Membrane, *J. Chem. Theory Comput.*, 2024, **20**, 5774–5783.
- 14 A. France-Lanord, H. Vroylandt, M. Salanne, B. Rotenberg, A. M. Saitta and F. Pietrucci, Data-Driven Path Collective Variables, *J. Chem. Theory Comput.*, 2024, **20**, 3069–3084.
- 15 S.-K. Lee, S.-T. Tsai and S. C. Glotzer, Classification of complex local environments in systems of particle shapes through shape symmetry-encoded data augmentation, *J. Chem. Phys.*, 2024, **160**, 154102.
- 16 D. Wang and P. Tiwary, State predictive information bottleneck, *J. Chem. Phys.*, 2021, **154**, 134111.
- 17 Z. Zou and P. Tiwary, Enhanced Sampling of Crystal Nucleation with Graph Representation Learnt Variables, *J. Phys. Chem. B*, 2024, **128**, 3037–3045.
- 18 Y. Wang, J. M. L. Ribeiro and P. Tiwary, Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics, *Nat. Commun.*, 2019, **10**, 1–8.
- 19 B. P. Vani, A. Aranganathan, D. Wang and P. Tiwary, AlphaFold2-RAVE: From Sequence to Boltzmann Ranking, *J. Chem. Theory Comput.*, 2023, **19**, 4351–4354.
- 20 B. P. Vani, A. Aranganathan and P. Tiwary, Exploring Kinase Asp-Phe-Gly (DFG) Loop Conformational Stability with AlphaFold2-RAVE, *J. Chem. Inf. Model.*, 2024, **64**, 2789–2797.
- 21 X. Gu, A. Aranganathan and P. Tiwary, Empowering AlphaFold2 for protein conformation selective drug discovery with AlphaFold2-RAVE, *eLife*, 2024, **13**, RP99702.
- 22 S. K. Kearsley, On the orthogonal transformation used for structural comparisons, *Acta Crystallogr., Sect. A*, 1989, **45**, 208–210.
- 23 S.-T. Tsai, Z. Smith and P. Tiwary, Reaction coordinates and rate constants for liquid droplet nucleation: Quantifying the interplay between driving force and memory, *J. Chem. Phys.*, 2019, **151**, 154106.
- 24 P. J. Steinhardt, D. R. Nelson and M. Ronchetti, Bond-orientational order in liquids and glasses, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1983, **28**, 784.
- 25 P. Geiger and C. Dellago, Neural networks for local structure detection in polymorphic systems, *J. Chem. Phys.*, 2013, **139**, 164105.
- 26 J. Rydzewski and O. Valsson, Multiscale Reweighted Stochastic Embedding: Deep Learning of Collective Variables for Enhanced Sampling, *J. Phys. Chem. A*, 2021, **125**, 6286–6302.
- 27 L. Bonati, V. Rizzi and M. Parrinello, Data-Driven Collective Variables for Enhanced Sampling, *J. Phys. Chem. Lett.*, 2020, **11**, 2998–3004.
- 28 S. Mehdi, D. Wang, S. Pant and P. Tiwary, Accelerating All-Atom Simulations and Gaining Mechanistic Understanding of Biophysical Systems through State Predictive Information Bottleneck, *J. Chem. Theory Comput.*, 2022, **18**, 3231–3238.
- 29 S. Lee, D. Wang, M. A. Seeliger and P. Tiwary, Calculating Protein–Ligand Residence Times through State Predictive Information Bottleneck Based Enhanced Sampling, *J. Chem. Theory Comput.*, 2024, **20**, 6341–6349.
- 30 D. Wang and P. Tiwary, Augmenting Human Expertise in Weighted Ensemble Simulations through Deep Learning based Information Bottleneck, *J. Chem. Theory Comput.*, 2024, DOI: [10.1021/acs.jctc.4c00919](https://doi.org/10.1021/acs.jctc.4c00919).
- 31 J. Behler and M. Parrinello, Generalized neural-network representation of high-dimensional potential-energy surfaces, *Phys. Rev. Lett.*, 2007, **98**, 146401.



32 J. Rogal, E. Schneider and M. E. Tuckerman, Neural-Network-Based Path Collective Variables for Enhanced Sampling of Phase Transformations, *Phys. Rev. Lett.*, 2019, **123**, 245701.

33 D. Kuroshima, M. Kilgour, M. E. Tuckerman and J. Rogal, Machine Learning Classification of Local Environments in Molecular Crystals, *J. Chem. Theory Comput.*, 2024, **20**(14), 6197–6206.

34 M. Sipka, A. Erlebach and L. Grajciar, Constructing Collective Variables Using Invariant Learned Representations, *J. Chem. Theory Comput.*, 2023, **19**, 887–901.

35 R. S. DeFever, C. Targonski, S. W. Hall, M. C. Smith and S. Sarupria, A generalized deep learning approach for local structure identification in molecular simulations, *Chem. Sci.*, 2019, **10**, 7503–7515.

36 A. Moradzadeh, H. Oliaei and N. R. Aluru, Topology-Based Phase Identification of Bulk, Interface, and Confined Water Using an Edge-Conditioned Convolutional Graph Neural Network, *J. Phys. Chem. C*, 2023, **127**, 2612–2621.

37 S. Banik, D. Dhabal, H. Chan, S. Manna, M. Cherukara, V. Molinero and S. K. Sankaranarayanan, CEGANN: Crystal Edge Graph Attention Neural Network for multiscale classification of materials environment, *npj Comput. Mater.*, 2023, **9**, 23.

38 Q. Kim, J.-H. Ko, S. Kim and W. Jhe, GCIceNet: a graph convolutional network for accurate classification of water phases, *Phys. Chem. Chem. Phys.*, 2020, **22**, 26340–26350.

39 M. Fulford, M. Salvalaglio and C. Molteni, DeepIce: A deep neural network approach to identify ice and water molecules, *J. Chem. Inf. Model.*, 2019, **59**, 2141–2149.

40 T. Xie and J. C. Grossman, Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties, *Phys. Rev. Lett.*, 2018, **120**, 145301.

41 P. B. Jørgensen, K. W. Jacobsen and M. N. Schmidt, Neural message passing with edge updates for predicting properties of molecules and materials, *arXiv*, 2018, preprint, arXiv:1806.03146, DOI: [10.48550/arXiv.1806.03146](https://doi.org/10.48550/arXiv.1806.03146).

42 J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals and G. E. Dahl, Neural message passing for quantum chemistry, *International conference on machine learning*, 2017, pp. 1263–1272.

43 K. Schütt, P.-J. Kindermans, H. E. Sauceda Felix, S. Chmiela, A. Tkatchenko and K.-R. Müller, Schnet: A continuous-filter convolutional neural network for modeling quantum interactions, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, ed. I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan and R. Garnett, Long Beach, CA, USA, December 4–9, 2017, 30, pp. 991–1001.

44 J. Zhang, L. Bonati, E. Trizio, O. Zhang, Y. Kang, T. Hou and M. Parrinello, Descriptors-free Collective Variables From Geometric Graph Neural Networks, *arXiv*, 2024, preprint, arXiv:2409.07339, DOI: [10.48550/arXiv.2409.07339](https://doi.org/10.48550/arXiv.2409.07339).

45 O. Valsson, P. Tiwary and M. Parrinello, Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint, *Annu. Rev. Phys. Chem.*, 2016, **67**, 159–184.

46 A. Laio and M. Parrinello, Escaping free-energy minima, *Proc. Natl. Acad. Sci. U. S. A.*, 2002, **99**, 12562–12566.

47 P. Tiwary and M. Parrinello, From metadynamics to dynamics, *Phys. Rev. Lett.*, 2013, **111**, 230602.

48 J. M. L. Ribeiro, P. Bravo, Y. Wang and P. Tiwary, Reweighted autoencoded variational Bayes for enhanced sampling (RAVE), *J. Chem. Phys.*, 2018, **149**, 072301.

49 N. D. Pomarici, S. Mehdi, P. K. Quoika, S. Lee, J. R. Loeffler, K. R. Liedl, P. Tiwary and M. L. Fernández-Quintero, Learning high-dimensional reaction coordinates of fast-folding proteins using State Predictive information bottleneck and Bias Exchange Metadynamics, *bioRxiv*, 2023, preprint, DOI: [10.1101/2023.07.24.550401](https://doi.org/10.1101/2023.07.24.550401).

50 D. Wang, R. Zhao, J. D. Weeks and P. Tiwary, Influence of Long-Range Forces on the Transition States and Dynamics of NaCl Ion-Pair Dissociation in Water, *J. Phys. Chem. B*, 2022, **126**, 545–551.

51 A. Duval, S. V. Mathis, C. K. Joshi, V. Schmidt, S. Miret, F. D. Malliaros, T. Cohen, P. Liò, Y. Bengio and M. Bronstein, A Hitchhiker's Guide to Geometric GNNs for 3D Atomic Systems *arXiv*, 2024, preprint, arXiv:2312.07511, DOI: [10.48550/arXiv.2312.07511](https://doi.org/10.48550/arXiv.2312.07511).

52 A. Paszke, *et al.*, PyTorch: An Imperative Style, High-Performance Deep Learning Library, *arXiv*, 2019, preprint, arXiv:1912.01703, DOI: [10.48550/arXiv.1912.01703](https://doi.org/10.48550/arXiv.1912.01703).

53 M. Fey and J. E. Lenssen, Fast Graph Representation Learning with PyTorch Geometric, *arXiv*, 2019, preprint, arXiv:1903.02428, DOI: [10.48550/arXiv.1903.02428](https://doi.org/10.48550/arXiv.1903.02428).

54 A. Mardt, L. Pasquali, H. Wu and F. Noé, VAMPnets for deep learning of molecular kinetics, *Nat. Commun.*, 2018, **9**, 1–11.

55 M. Ghorbani, S. Prasad, J. B. Klauda and B. R. Brooks, GraphVAMPNet, using graph neural networks and variational approach to Markov processes for dynamical modeling of biomolecules, *J. Chem. Phys.*, 2022, **156**, 184103.

56 B. Liu, M. Xue, Y. Qiu, K. A. Konovalov, M. S. O'Connor and X. Huang, GraphVAMPnets for uncovering slow collective variables of self-assembly dynamics, *J. Chem. Phys.*, 2023, **159**, 094901.

57 P. Schwerdtfeger and D. J. Wales, 100 Years of the Lennard-Jones Potential, *J. Chem. Theory Comput.*, 2024, **20**, 3379–3405.

58 V. G. Satorras, E. Hoogeboom and M. Welling, E(n) Equivariant Graph Neural Networks, *arXiv*, 2021, preprint, arXiv:2102.09844, DOI: [10.48550/arXiv.2102.09844](https://doi.org/10.48550/arXiv.2102.09844).

59 G. A. Tribello, M. Ceriotti and M. Parrinello, A self-learning algorithm for biased molecular dynamics, *Proc. Natl. Acad. Sci. U. S. A.*, 2010, **107**, 17509–17514.

60 L. Evans, M. K. Cameron and P. Tiwary, Computing committors in collective variables *via* Mahalanobis diffusion maps, *Appl. Comput. Harmon. Anal.*, 2023, **64**, 62–101.

61 W. Humphrey, A. Dalke and K. Schulten, VMD: Visual molecular dynamics, *J. Mol. Graphics*, 1996, **14**, 33–38.

62 P. G. Bolhuis, C. Dellago and D. Chandler, Reaction coordinates of biomolecular isomerization, *Proc. Natl. Acad. Sci. U. S. A.*, 2000, **97**, 5877–5882.



63 T. Fröhling, L. Bonati, V. Rizzi and F. L. Gervasio, Deep learning path-like collective variable for enhanced sampling molecular dynamics, *J. Chem. Phys.*, 2024, **160**, 174109.

64 P. Kang, E. Trizio and M. Parrinello, Computing the committor with the committor to study the transition state ensemble, *Nat. Comput. Sci.*, 2024, **4**, 451–460.

65 S. Brody, U. Alon and E. Yahav, How Attentive are Graph Attention Networks?, *International Conference on Learning Representations*, 2022, <https://openreview.net/forum?id=F72ximsx7C1>.

66 O. Blumer, S. Reuveni and B. Hirshberg, Combining stochastic resetting with Metadynamics to speed-up molecular dynamics simulations, *Nat. Commun.*, 2024, **15**, 240.

67 S.-T. Tsai, Z. Smith and P. Tiwary, SGOOP-d: Estimating Kinetic Distances and Reaction Coordinate Dimensionality for Rare Event Systems from Biased/Unbiased Simulations, *J. Chem. Theory Comput.*, 2021, **17**, 6757–6765.

68 F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda and M. M. Bronstein, Geometric deep learning on graphs and manifolds using mixture model CNNs, *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5425–5434.

69 D. Errington, C. Schneider, C. Bouysset and F. A. Dreyer, Assessing interaction recovery of predicted protein-ligand poses, *arXiv*, 2024, preprint, arXiv:2409.20227, DOI: [10.48550/arXiv.2409.20227](https://doi.org/10.48550/arXiv.2409.20227).

