

Permutationally Invariant Networks for Enhanced Sampling (PINES): Discovery of Multi-Molecular and Solvent-Inclusive Collective Variables

Nicholas S.M. Herringer,[†] Siva Dasetty,[‡] Diya Gandhi,[‡] Junhee Lee,[‡] and Andrew L. Ferguson^{*,‡}

[†]*Department of Chemistry, University of Chicago, Chicago, Illinois 60637*

[‡]*Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637*

E-mail: andrewferguson@uchicago.edu

Abstract

The typically rugged nature of molecular free energy landscapes can frustrate efficient sampling of the thermodynamically relevant phase space due to the presence of high free energy barriers. Enhanced sampling techniques can improve phase space exploration by accelerating sampling along particular collective variables (CVs). A number of techniques exist for data-driven discovery of CVs parameterizing the important large scale motions of the system. A challenge to CV discovery is learning CVs invariant to symmetries of the molecular system, frequently rigid translation, rigid rotation, and permutational relabeling of identical particles. Of these, permutational invariance has proved a persistent challenge in frustrating the data-driven discovery of multi-molecular CVs in systems of self-assembling particles and solvent-inclusive CVs for solvated systems. In this work, we integrate Permutation Invariant Vector (PIV) featurizations with autoencoding neural networks to learn nonlinear CVs invariant to translation, rotation, and permutation, and perform interleaved rounds of CV discovery and enhanced sampling to iteratively expand sampling of configurational phase space and obtain converged CVs and free energy landscapes. We demonstrate the Permutationally Invariant Network for Enhanced Sampling (PINES) approach in applications to the self-assembly of a 13-atom Argon cluster, association/dissociation of a NaCl ion pair in water, and hydrophobic collapse of a $C_{45}H_{92}$ *n*-pentatetracontane polymer chain. We make the approach freely available as a new module within the PLUMED2 enhanced sampling libraries.

1 Introduction

Classical molecular dynamics (MD) simulations present a powerful technique to understand and engineer the molecular-level behaviors of chemical, materials, and biological systems.^{1,2} Rugged free energy landscapes are a generic feature of many molecular systems³ wherein high free energy barriers several multiples or more of the thermal energy can trap unbiased simulation trajectories and result in incomplete sampling of the thermally accessible phase space, failure to resolve the relevant metastable states, and un converged thermodynamic averages.^{4,5} Collective variable (CV)-based biasing techniques are a class of non-Boltzmann sampling techniques designed to ameliorate the sampling problem by accelerating barrier crossing along particular CVs.⁶ A panoply of CV biasing enhanced sampling techniques exist,^{4,5,7,8} including, for example, umbrella sampling,⁹ adaptive biasing force,¹⁰ and metadynamics.¹¹ The techniques differ in the numerical details of how the biasing forces are applied, but the success of all of these techniques relies upon the identification of appropriate CVs along which to drive sampling. Robert Zwanzig's observation regarding the projection operator formalism is apposite:¹² "Statistical mechanics does not tell us what the relevant variables are. This is our choice. If we choose well, the results may be useful; if we choose badly, the results (while still formally correct) will probably be useless." Generally, we wish to select CVs coincident with the large scale and/or slow motions of the system so that we can enhance sampling along the conformational degrees of freedom important to the system dynamics. We also wish to judiciously select a parsimonious number of CVs so as to enable efficient sampling and the opportunity for thermodynamically converged sampling within the CV subspace. CV discovery is fundamentally a problem of dimensionality reduction in the space of the Cartesian coordinates defining the configurational phase space of the system, and these objective functions – high-variance CVs for large scale motions or maximally autocorrelated CVs for slow motions – underpin the preponderance of systematic CV discovery techniques using variants of dimensionality reduction algorithms.^{7,8,13}

Autoencoders – artificial neural networks (ANNs) comprising an encoding of high-dimensional

observations into a low-dimensional latent bottleneck followed by decoding of the latent space back up to the high-dimensional observational space – have proven particularly well-suited to CV discovery applications.^{14,15} As universal function approximators capable of unsupervised learning and automatic differentiation, these architectures can be trained over unlabeled simulation trajectories to discover nonlinear CVs that are explicit and differentiable functions of the atomic coordinates.⁷ These favorable properties mean that it is not necessary to know the metastable states of the system *a priori* (i.e., unlabeled trajectory data is sufficient) and the differentiable CVs can be modularly incorporated into standard enhanced sampling algorithms by propagating biasing forces on the CVs to those on the atoms using the chain rule.^{8,16} In terms of high-variance CV discovery, we previously introduced Molecular Enhanced Sampling with Autoencoders (MESA) as an approach that interleaves successive rounds of CV learning and biased sampling to simultaneously discover CVs coincident with the large-scale dynamical motions of the system and enhance sampling along these coordinates.^{8,16} The process is iterated until we achieve convergence in phase space exploration and a set of CVs that globally parameterize the thermally relevant configurational subspace. Tiwary and co-workers developed reweighted autoencoded variational Bayes for enhanced sampling (RAVE)¹⁷ as an approach similar to MESA that focusses on the discovery of a single CV that is highly correlated with a physically interpretable order parameter. Belkacemi, Gkeka, Lelièvre, and Stoltz proposed free energy biasing and iterative learning with autoencoders (FEBILAE) as a technique very closely related to MESA but which reweights the biased simulation data prior to each new round of CV discovery to eliminate the effect of non-Boltzmann sampling in the learned CVs.¹⁸ Ensing and co-workers developed the FABULOUS approach that shares similarities to both MESA and RAVE as well as earlier work by Ma and Dinner.^{19,20} This approach essentially employs only the decoding block of an autoencoder to identify CVs capable of accurate reconstruction of atomic coordinates. Luber and co-workers developed deep autoencoder neural network (DAENN)²¹ and a framework called Deep Learning for Collective Variables (DeepCV)²² implementing it to iteratively learn CVs

of chemical reactions from limited training configurations covering only the reactant state using eXtended Social PeRmutation InvariaNT (xSPRINT) molecular representation. In terms of slow CV discovery, a number of unsupervised autoencoder-based architectures have also found success, including the time-lagged autoencoders (TAE) of Wehmeyer and Noé,²³ variational dynamics encoders (VDE) of Pande and co-workers,^{24–26} state predictive information bottleneck (SPIB) of Wang and Tiwary,²⁷ and our state-free reversible VAMPNets (SRVs)^{28–30} also known as Deep-TICA.³¹ In a supervised learning modality, Bolhuis and co-workers employed augmented/extended autoencoders that also take in labeled information on the molecular committor function to learn CVs capable of accurate reconstruction of both the atomic coordinates and progression along a reactive path.³²

An important consideration in CV discovery is the determination of symmetry-adapted CVs that respect the underlying symmetries of the molecular system.^{33,34} The most commonly encountered symmetries are translation, rotation, and permutation arising from the invariance of a molecular system to rigid translation and/or rotation and the permutational relabeling of identical and indistinguishable particles. Failure to respect these symmetries in the learned CVs is at best data inefficient and at worst can lead to fundamental failure modes associated with the artificial distinction of physically identical system configurations. In general, CVs for enhanced sampling should be invariant to these symmetries, although there are cases where covariance may be desirable.^{35–37} Symmetry adaptation is typically enforced by the choice of featurization of the system and/or within the neural network architecture. The learned CVs are relatively easily made invariant to translation and rotation by representing the system using internal coordinates (i.e., bonds, angles, torsions) that do not depend on the arbitrary location and orientation of the Cartesian coordinate system^{33,38,39} or by data augmentation.^{8,40}

Permutational invariance has proven more challenging to eliminate, and this has frustrated CV discovery and biased sampling in common and important applications such as self-assembling systems of identical particles or solutes immersed in a bath of identical sol-

vent molecules. Good methods do exist, however, to handle permutational invariance. At the level of the network architecture, Winter, Noé, and Clevert developed a permutation-invariant graph autoencoder (PIGAE) for graph reconstruction by employing an explicit permuter that learns an explicit permutation matrix for each input graph, although the quartic scaling in reordering operations during each training step makes this approach expensive.⁴¹ Similarly, Huang and co-workers⁴² developed a GraphVAMPnets based approach that utilizes graph neural networks to respect permutation and rotational symmetries of particles in self-assembling systems by enforcing identical node embeddings for such symmetric particles and learn the slow CVs from the resulting graph embeddings using VAMPnets. Prudente, Acioli, and Soares Neto⁴³ and Nguyen and Le⁴⁴ developed special purpose neural networks for the fitting of potential energy surfaces of small molecules in which the first layer of the network was modified to respect permutational invariance, although generalization and scaling to larger systems is challenging. A cheaper and more scalable approach is to represent the system to the network via permutationally-invariant featurizations.³³ Provided these featurizations can be expressed as explicit functions of the particle coordinates, the derivatives of the CVs learned by the neural network with respect to the Cartesian coordinates can be constructed from the chain rule via automatic differentiation and/or analytical derivatives. Permutational invariance is typically enforced by summing particle-centric basis functions over all identical particles or a permutationally-invariant ordering. Examples of the former class include permutation-invariant polynomials (PIPs) generated by summing monomial functions of internal coordinates over all like-particle orderings,^{45,46} atom-centered symmetry functions (ACSF) that achieve permutational invariance by a similar summation procedure,^{47,48} many-body tensor representations (MBTR) that define scalar geometry functions (e.g., atomic number, distance, angle) over small body-order numbers of atoms and arranges these into permutationally-invariant distribution functions for each class of interaction,⁴⁹ solvent signature vectors that use Gaussian kernel functions to quantify the distribution of solvent molecules around the solute,⁵⁰ solvent shell featurizations that

quantify the number of water molecules in nested shells around the solute,^{51,52} and smooth overlap of atomic orbitals (SOAP) that generates the partial power spectra of an expansion of spherical harmonic representations of the atomic-centered smoothed density fields summed over all identical atoms.⁵³ In general, the exponential scaling in the number of permutational orderings associated with the summations for each of these approaches must be controlled by defining a cutoff in the range and/or order of the interactions considered.³³ Ordering approaches typically offer a computationally cheaper means to achieve permutational symmetry adaptation.³³ Examples of this class include the bag-of-bonds (BoB) featurization based on a concatenation of the ordered (inverse) pairwise distances for each class of bond scaled by the product of the nuclear charges⁵⁴ and social permutation invariant (SPRINT) coordinates constructed from the leading eigenvector of a graph adjacency matrix.^{55,56} Of particular interest in this work is the Permutation Invariant Vectors (PIV) featurization developed by Pietrucci and co-workers.^{56–58} This representation can be conceived of as a concatenation of instantaneous radial distribution functions for classes of atom pairs in the system that, by virtue of distance-based ordering, is invariant to like-particle permutation and, through the use of exclusively pairwise distances, also translation and rotation. Physically motivated, bespoke order parameters engineered to preserve permutational invariance have also been developed, although these can be less generically extensible to arbitrary molecular systems. A trivial strategy to treat the permutational invariance of a solvent bath is to simply ignore the solvent coordinates and assume that the solvent effects are sufficiently strongly imprinted on the solute conformational dynamics that they can be driven implicitly by biasing solute-centric CVs.⁵⁹ Patel, Varilly, Chandler, and Garde developed the indirect umbrella sampling (INDUS) approach to drive water density fluctuations in user-defined control volumes and have used this technique to drive sampling of wetting/dewetting transitions.^{60,61} Zou *et al.* employed CVs based on coordination numbers, Steinhardt bond order parameters, interfacial water density, intermolecular angles, and pair orientational entropy, and combined these within the SPIB framework to drive sampling of crystal nucleation and polymorph

formation.⁶² Rizzi *et al.* combined CVs describing solvent coordination number and solvent structure around a binding pocket with the deep linear discriminant analysis (Deep-LDA) approach to incorporate water coordinates in driving sampling of host-guest interactions.⁶³

The primary contribution of this work is to present an approach for the integrated data-driven discovery of translationally, rotationally, and permutationally-invariant CVs in systems containing identical particles and their use within enhanced sampling calculations to accelerate phase space exploration. We term this approach Permutationally Invariant Networks for Enhanced Sampling (PINES). Our approach employs a PIV representation as input for an autoencoder to perform interleaved nonlinear CV discovery and CV biasing enhanced sampling using parallel bias metadynamics.⁶⁴ We draw inspiration for the synthesis of permutationally invariant featurizations with ANNs from the deep potential molecular dynamics (DeePMD) approach of Zhang *et al.* that employs an ordered distance featurization passed to an ANN to learn a potential energy function for molecular dynamics³⁹ and the PIP-NN approach of Jiang and Guo that performs a similar feat using a permutation-invariant polynomial (PIP) featurization.⁶⁵ We eschew the use of physically-motivated, hand-crafted descriptors in favor of a more generic, transferable, and system agnostic featurization strategy, but observe that the incorporation of system-specific knowledge may be straightforwardly achieved by augmenting the descriptor list with any additional features provided they are explicit and differentiable functions of the particle Cartesian coordinates. The PINES paradigm is also modular in the sense that different featurizations may be employed to adapt to different symmetries of the system (e.g., translation, rotation, permutation, inversion, chiral) and different ANNs may be used to perform different learning tasks (e.g., unsupervised CV discovery, supervised learning of potential energy surfaces). We demonstrate PINES in applications to three molecular systems in which proper accounting for particle permutational invariance is of paramount concern: (i) self-assembly of a 13-atom Ar cluster, (ii) assembly/disassembly of a NaCl ion pair in water, and (iii) hydrophobic collapse of a C₄₅H₉₂ *n*-pentatetracontane polymer chain. We make the approach freely available as a new user-

friendly module at <https://github.com/Ferg-Lab/pines.git> that can be patched with PLUMED2⁶⁶ enhanced sampling libraries.

The structure of the remainder of this paper is as follows. In Sec. 2, we introduce the PINES method for translationally, rotationally, and permutationally-invariant CV discovery and enhanced sampling using parallel bias metadynamics. This is followed by a description of the PIV featurization, PINES training procedure, and the enhanced sampling molecular dynamics simulations. In Sec. 3, we demonstrate applications of PINES in discovering CVs for enhanced sampling to self-assembly of 13-atom Ar cluster, association/disassociation of NaCl ion pair in water, and hydrophobic collapse of C₄₅H₉₂ polymer chain, and compare their free energy landscapes to prior work. Finally, in Sec. 4 we present our conclusions and perspectives for future work.

2 Methods

2.1 Permutationally Invariant Networks for Enhanced Sampling (PINES)

2.1.1 Molecular Enhanced Sampling with Autoencoders (MESA)

PINES can be conceived as a permutationally-invariant extension to the Molecular Enhanced Sampling with Autoencoders (MESA) approach that we previously developed for data-driven CV discovery and enhanced sampling.^{8,16} The key innovation of this approach is to use autoencoding ANNs to perform unsupervised, nonlinear dimensionality reduction of high-dimensional molecular simulation trajectories to discover low-dimensional CV parameterizations spanning the high-variance collective motions of the system. By virtue of the ANN architecture and automatic differentiation through the network, the functional dependence upon the first derivatives with respect to the Cartesian coordinates of the particles in the system are explicitly available, which enables propagation of biases in the CVs to forces

on atoms and permits the learned CVs to be seamlessly incorporated into any off-the-shelf CV biasing enhanced sampling approach. The “chicken-and-egg problem” of data-driven CV discovery observes that CV discovery requires training trajectories spanning the thermally relevant phase space, and trajectories spanning the thermally relevant phase space require good CVs to drive sampling.⁶⁷ MESA solves this problem by interleaving successive rounds of CV discovery and enhanced sampling until there is no additional round-to-round enhancement in exploration of phase space or change in number or character of the learned CVs. The first round of the CV discovery process may be initialized using short unbiased trajectories, simulated annealing, or biased trajectories in intuited or physically-motivated CVs. Subsequent rounds of the process are conducted by performing enhanced sampling in data-driven CVs discovered by the ANN. A schematic illustration of the PINES approach with parallel bias metadynamics as the CV biasing enhanced sampling approach is illustrated in Fig. 1.

2.1.2 Permutation Invariant Vector (PIV) featurization

PINES employs a Permutation Invariant Vector (PIV)^{56–58} representation of the molecular system to permit discovery of translationally, rotationally, and permutationally-invariant CVs for enhanced sampling. The computational graph underpinning PINES is simply a PIV featurization of the system followed by a regularized autoencoder for data-driven CV discovery (Fig. 1a,b). In this sense, it is very similar to MESA but employs a PIV representation of the molecular system. Importantly, gradients can be backpropagated end-to-end through the autoencoder encoder and PIV featurizer to permit calculation of the gradients in the CVs required for enhanced sampling calculations.

The PIV featurization can be conceived as a concatenation of instantaneous pair correlation functions specifying the ordered set of pairwise distances between the chemically identical and indistinguishable atoms of each class (e.g., all water O atoms, all Na^+ ions, etc.)^{56–58} Since the featurization is constructed from pairwise distances it is naturally invariant to rigid translation and rotation, and by virtue of the inherent ordering of pairwise

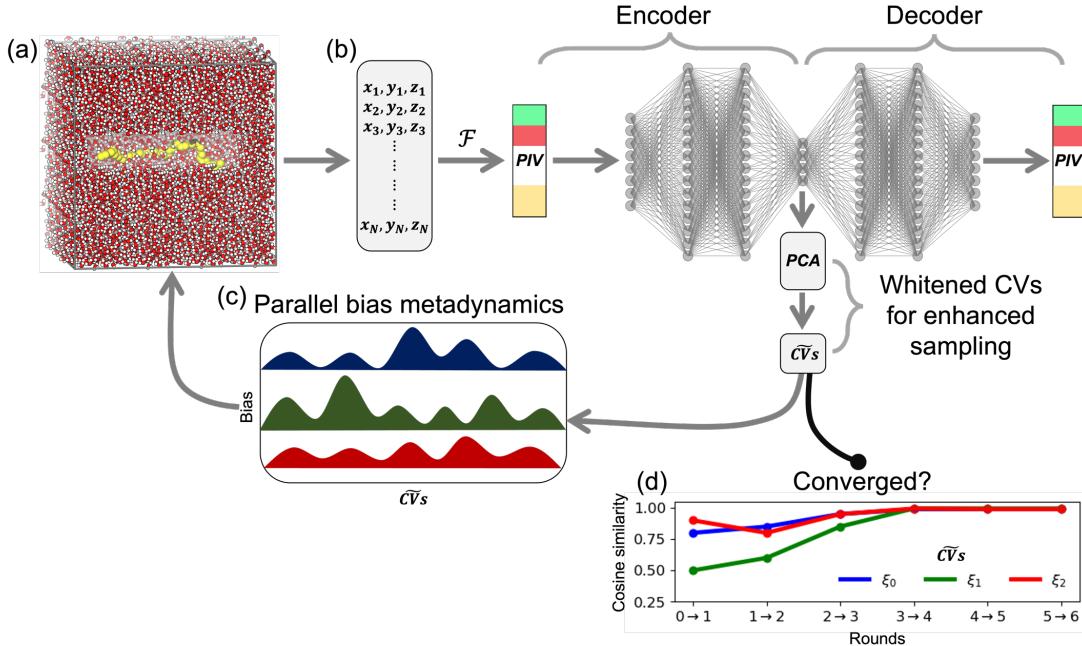


Figure 1: Schematic illustration of the Permutationally Invariant Network for Enhanced Sampling (PINES) approach for data-driven CV discovery and enhanced sampling. (a) Molecular simulation trajectories provide training data for CV discovery. The process is initialized with unbiased or biased simulations to provide an initial first sampling of phase space. Subsequent rounds are trained over biased simulations generated by enhanced sampling in the learned CVs from the previous round. (b) The trajectories are featurized and passed to a maximum mean discrepancy-Wasserstein autoencoder (MMD-WAE) to perform unsupervised nonlinear dimensionality reduction into a low-dimensional latent space spanned by the learned CVs. To stabilize the CV discovery process and help disentangle the latent space, we employ batch normalization between all layers with exception of the final layer. (c) After training the network, we extract the CVs from the bottleneck layer between the encoder and decoder. We find that it is useful to first pass the output of the bottleneck layer through a principal components analysis (PCA) whitening transformation to help align the discovered manifold with the highest variance directions for efficient parallel bias metadynamics (PBMetaD) sampling. The learned CVs are represented as the functional composition $\xi = h \circ f(\mathbf{r}) = h(\mathbf{z})$, where $\mathbf{z} = f(\mathbf{r})$ is the result of applying the featurization f – in this work a PIV featurization – to the Cartesian coordinates \mathbf{r} of the system, and $\xi = h(\mathbf{z})$ is the action of the MMD-WAE encoder plus PCA whitening transform upon the input featurization that generates a projection into the learned CVs. The derivative of the CVs with respect to the Cartesian coordinates of atom i is required for propagation of biasing forces in the CVs to biasing forces on the atoms, and comes from simple application of the chain rule to the function composition that is computed in practice by automatic differentiation, $\nabla_{\mathbf{r}_i} \xi = (\partial h / \partial \mathbf{z}) \nabla_{\mathbf{r}_i} \mathbf{z}$. (d) The learned CVs are implemented within PBMetaD calculations to drive sampling in the configurational phase space spanned by these CVs, completing one full iteration of PINES. The biased simulation trajectories from Round n are then passed forward as the training trajectories for CV discovery within Round $(n + 1)$ to determine new CVs and conduct new enhanced sampling calculations. The iterative process is terminated when the learned CVs stabilize, we observe no further exploration of phase space, and the reweighted free energy landscape converges.

distances within an instantaneous pair correlation function, it is also invariant to permutations of the labels attached to the indistinguishable atoms of each class. Since the PINES network is trained to reconstruct the PIV representation of the system, the learned CVs are also inherently invariant to translation, rotation, and permutation. Importantly, the PIV representation of a system is complete in the sense that if all pairwise distances are incorporated into the PIV featurization, rigid graph theory dictates that the Cartesian coordinates of the particles constituting the system can be unambiguously reconstructed up to translation, rotation, and permutation.^{68,69}

PIVs were introduced by Pietrucci and co-workers^{56–58} and have been used, among other applications, to efficiently drive sampling of different phases of water.⁵⁷ We illustrate in Fig. 2 the construction of a PIV featurization for a NaCl ion pair in water. The PIV featurization proceeds in three steps. First, we compute the distance matrix with elements a_{ij} recording the Euclidean pairwise distances between atoms i and j . We arrange the elements of the matrix into a block structure where each block contains the set of pairwise distances between two atom classes (i.e., groups of chemically indistinguishable atoms). Within each block the ensemble of pairwise distances are unchanged under permutational relabeling of the chemically indistinguishable atoms within each class. Second, we pass the pairwise distances through a rational switching function to produce the adjacency matrix,

$$v_{ij} = \frac{1 - \left(\frac{a_{ij} - d_0}{r_0}\right)^n}{1 - \left(\frac{a_{ij} - d_0}{r_0}\right)^m} \quad (1)$$

where the hyperparameter d_0 is minimum interparticle distance, r_0 is the interparticle distance which lies on the midpoint of the switching function curve (Fig. 2b), and n and m are non-negative numbers which control the steepness of the switching function curve and thus determine the resolution of the various pairwise interactions. This transformation has the two desirable effects of scaling the distances to the interval $(0, 1)$ and producing values close to unity for short distances, values close to zero for large distances, and a switching range

between these two limits over which the transformation has the most sensitivity to changes in a_{ij} . This ensures that the v_{ij} values passed as input to the encoder are appropriately scaled to $(0, 1)$ range and, by tuning the parameters of the switching function, permits us to focus the dynamic range of v_{ij} upon the physical range of distances over which we wish the network to pay the most attention. In principle, other switching functions⁷⁰ could be applied as well in place of the rational function (Eq. 1) for computing v_{ij} . Typically, we set the parameters in Eq. 1 based on a radial distribution function (RDF) collected over a short unbiased trial simulation of the system in its stable state (Fig. S1). We choose r_0 to be the mean of the first and second peaks of the RDF, set $d_0 = 0$ and $m = 2n$ and tune n such that $v_{ij} \approx 0.9$ at the a_{ij} value of the first peak and $v_{ij} \approx 0.1$ at the a_{ij} value of the second peak in RDF. This ensures that the v_{ij} values are most responsive to particles moving between the first two shells of their radial distribution where it is typically most important to highlight structural differences between various system configurations and become insensitive in the far field.⁵⁶ When appropriately tuned in this fashion, the PIV featurization show large inter-class differences between qualitatively different system configurations and small intra-class differences for configurations in the same structural class.⁵⁶ The important dynamical ranges for the systems studied in this work were relatively narrow such that we were able to use relatively steep switching functions that give good resolution over the region of interest. For larger dynamic ranges – such as one might encounter in large or complex self-assembling systems – one would typically use a longer-ranged switching function so as not to lose resolution in the flat regions of the switch with small gradients. Such a long-ranged switching function, however, offers reduced resolving power over its range due to the gradual nature of the switch. A superior, albeit more computationally expensive, solution would be to use multiple steep switching functions for each pairwise interaction tuned to be responsive over a range of concentric distance shells, wherein each switch provides high resolution over a local range and is arranged to partially overlap with its neighboring switching functions, thereby providing complete coverage of the relevant dynamical range. Third, we sort the v_{ij}

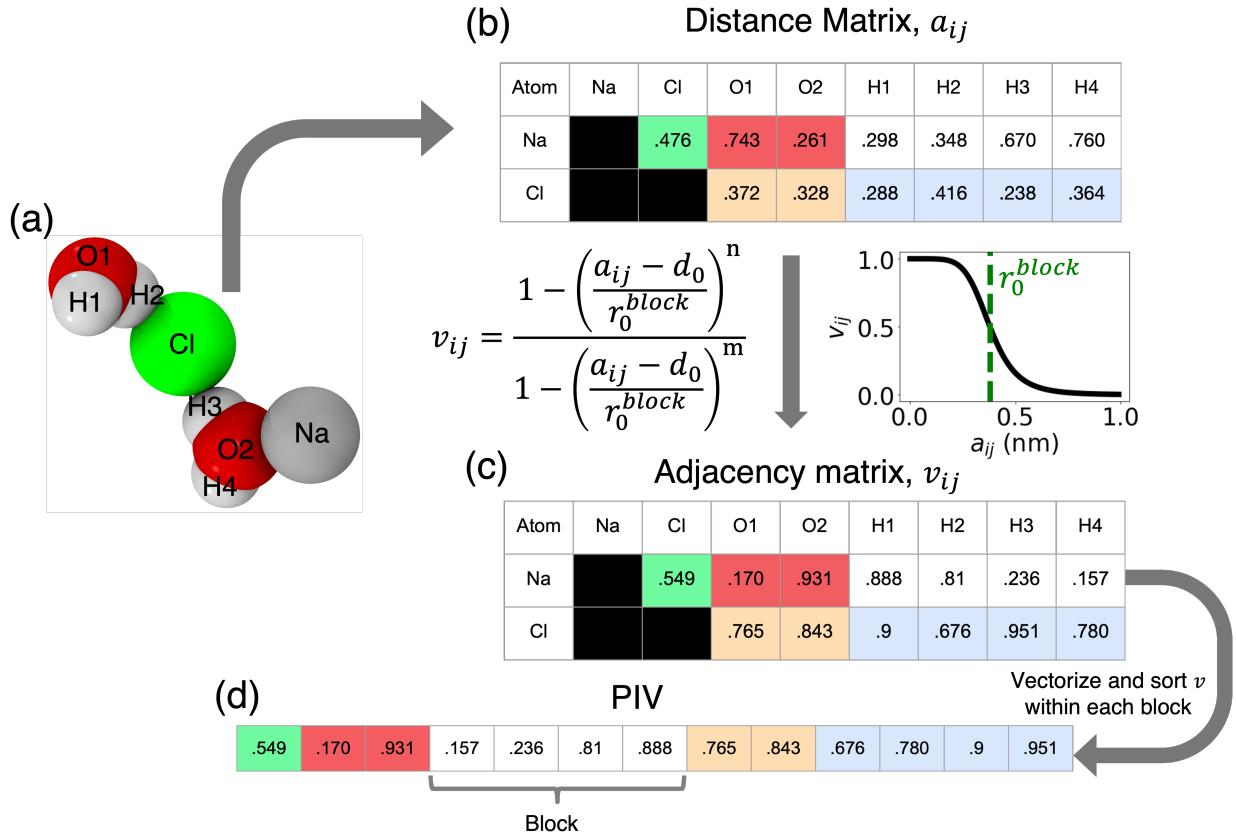
values within each block of the adjacency matrix in non-descending order and concatenate them into a single vector – the eponymous PIV. The sorting operation is the critical step in introducing permutational invariance.

Importantly, construction of the PIV from the Cartesian coordinates of the system is a deterministic and differentiable operation that permits backpropagation of gradients from the learned CVs to the PIV featurization and then the Cartesian coordinates by automatic differentiation. A well known deficiency of PIVs, and one that is shared by all featurizations that achieve permutational invariance by ordering, is that the featurization can exhibit kinks and its derivative therefore possess discontinuities at locations in the configurational phase space where the ordering of the particles changes.³³ Empirically, we encountered no difficulties or instabilities in either our CV learning process or our enhanced sampling calculations in the learned CVs. It is possible that the discontinuities are sufficiently gentle and/or rare that they do not present any difficulties, but this issue may warrant additional study to identify situations where these singularities may present difficulties.

2.1.3 Autoencoder network design and training

The second component of PINES is an autoencoding neural network that performs unsupervised nonlinear dimensionality reduction of molecular simulation training trajectories to discover low-dimensional embeddings into a small number of high-variance CVs.^{8,14–16} In this work, we employ a maximum mean discrepancy-Wasserstein autoencoder (MMD-WAE).⁷² The network is trained by minimizing the loss function comprising a mean squared error (MSE) term to encourage high fidelity reconstruction and a maximum mean discrepancy (MMD) loss to regularize the latent space,

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (PIV_i - \widehat{PIV}_i)^2 + E_P[k(x, x')] + E_Q[k(y, y')] - 2E_{P,Q}[k(x, y)], \quad (2)$$



where $k(z, z') = \exp(-|z - z'|^2)$ is a Gaussian kernel with unit bandwidth, $\frac{1}{N} \sum_{i=1}^N (PIV_i - \widehat{PIV}_i)^2$ is the element-wise mean squared error of the reconstruction, $E_P[k(x, x')] + E_Q[k(y, y')] - 2E_{P,Q}[k(x, y)]$ is the MMD loss, $P(x)$ is the probability distribution of the latent space, and $Q(y)$ is the target latent space distribution. The MMD loss is equal to zero only when the two distributions are identical. In this work, we choose the target distribution to be a normal distribution with a zero mean and unit variance.

We note that in principle it may be advantageous to sort the elements within each block of the reconstructed PIV into non-descending order prior to calculating the MSE loss. In this way the decoder need only learn the distribution of values within each block instead of also being required to learn the sorting rule. Empirically, we find that the decoder is able to readily learn to produce values in non-descending order and only rarely violates proper sorting within the reconstructed PIV. It appears, therefore, that pre-loss sorting of the reconstructed PIV is unnecessary. Additionally we draw attention to the possibility of using non-uniform weights over the various contributions to the MSE term in order to modulate the importance of reconstruction fidelity for particular interactions in the system,⁵⁶ although in our work we observed satisfactory performance using uniform weighting.

The encoder and decoder are both implemented as fully-connected feedforward neural networks with three hidden layers (Fig. 1b). For all systems considered in this work, the overall network architecture is $M-64-32-16-d-16-32-64-M$, where M is the dimensionality of the PIV featurization and d is the number of nodes in the bottleneck layer that defines the number of CVs to be discovered. An appropriate value for d is determined each round by training networks with d ranging from 1-8 and selecting the appropriate value of d from a knee in a plot of the fraction of variance explained curve using the L-method of Salvador and Chan^{8,16,73} (Fig. S2). For the systems studied in this work, we find a consistent latent space dimensionality of $d=3$ over the PINES iterations. For more complicated systems where the first round of sampling was insufficient to expose all of the important collective modes, these may emerge gradually over the PINES iterations and be manifest as an increase in the latent

space dimensionality exposed by this analysis. Biomolecular systems are archetypal complex systems in the sense that they tend to evince an emergent large-scale low-dimensionality due to nonlinear many-body couplings between the atomic degrees of freedom.^{59,74–79} Nevertheless, the enhanced sampling technique selected for this work, parallel bias metadynamics (cf. Section 2.1.4), scales well to quite large numbers of CVs^{64,80} thereby enabling the application of PINES to systems with high-dimensional latent spaces. We employ *tanh* activation functions between all hidden layers and a sigmoid activation in the final layer of the decoder. Batch normalization is employed between all hidden layers except the layer preceding the reconstructed PIV, which we find to be critical in avoiding saturation and stabilizing CV discovery. We construct and train the network using PyTorch⁸¹ with PyTorch Lightning⁸² and train it using backpropagation employing the Adam optimizer⁸³ with a batch size of 100 and learning rate of 10^{-4} . Training was terminated when the validation loss stopped decreasing employing a patience of 10 epochs. Training was performed on a NVIDIA Quadro RTX 6000 GPU card and, for the systems studied in this work, typically required hundreds of epochs and tens of GPU-minutes. After training is complete, we discard the decoder and pass the output of the bottleneck layer through a principal components analysis (PCA) whitening transformation (Fig. 1b).⁸⁴ We find this step useful in disentangling the learned CVs by imposing ordering by variance that helps to both stabilize the round-to-round discovery process and also improves enhanced sampling by orienting the biasing potentials along the orthogonalized PCA subspace. This transformation can be thought of as a simple change of basis in which the dimensionality of the latent space is conserved and the whitened CVs represent a rotation of the manifold to align the axes with the highest variance regions of the embedding.

2.1.4 Enhanced sampling along discovered CVs using parallel bias metadynamics

The learned CVs within each round are used to conduct enhanced sampling of configurational space using parallel bias metadynamics (PBMetaD).⁶⁴ The PINES approach is compatible with any off-the-shelf CV biasing enhanced sampling technique, but we elect to use PBMetaD due to its efficient scaling to large numbers of CVs by the application of multiple simultaneous one-dimensional biases (Fig. 1c). In this work, we never have need to drive sampling of more than three CVs, but the PBMetaD approach has been demonstrated up to several tens of CVs.⁸⁰ We couple all learned CVs to a well-tempered variant of PBMetaD enhanced sampling calculation implemented within the PLUMED2 library.^{64,66,85} The hyperparameters for the PBMetaD calculations are detailed for each of the three systems in Section 2.2.

Convergence of the PBMetaD calculations are assessed using two criteria.⁸⁶ First, we check that the biasing potential has become quasi-static by assuring that the height of the Gaussian biases deposited over the course of the simulation drop below a particular fraction of the initial height. In early iterations we adopt a less stringent criterion of 10% but in later rounds, we refine this to a more strict criterion of 1%. Second, we monitor the distribution of the CVs and check for diffusive sampling over the full range explored by each CV to verify that the simulation can freely diffuse over the barriers in all coupled coordinates.

2.1.5 Convergence assessment

Convergence of PINES is assessed by measuring the dimensionality d and cosine similarity ρ of the CVs learned between successive iterations and the expansion of the explored phase space. Stabilization of d of the MMD-WAE bottleneck layer indicates that no additional dimensions are being explored in successive biasing rounds, a cosine similarity of $\rho > 0.9$ between the CVs discovered in two consecutive rounds – computed by projecting the present round’s simulation data into the CVs learned in the present and previous rounds – is used to assess whether the CVs themselves are approximately constant between successive rounds,

and we verify that the boundary of explored phase space is no longer expanding (Fig. 1d).

Once all convergence criteria are satisfied, we declare the PINES procedure to be complete. At this point we construct a training data set by collating the biased trajectories from all prior rounds and use this concatenated data set to run a final round of CV discovery and PBMetaD enhanced sampling. For large datasets, we perform uniform sampling across the multidimensional CV space used for the latest round of biasing to generate a smaller, subsampled dataset that is representative of the whole CV space but more amenable to training a terminal network to identify the optimal embedding. These final CVs are used to perform a PBMetaD production run to demonstrate their impact on accelerating exploration of phase space. The unbiased free energy surface is recovered from the terminal PBMetaD biased trajectories by applying the standard Torrie-Valleau sampling weights to each configuration to eliminate the effect of the applied bias and enable projection into arbitrary CVs beyond those in which the biased sampling was conducted.^{9,64} Empirically, we do not observe any significant benefit from applying this reweighting to the biased trajectories in each individual CV discovery round, but this, of course, is possible.¹⁸

2.2 All-atom molecular dynamics simulations

All-atom molecular dynamics simulations of each of the three systems considered in this work were performed using the GROMACS 2021.1 simulation suite⁸⁷ patched with an in-house modified version of the PLUMED2 v2.8.0 enhanced sampling libraries.⁶⁶ This patch implements the PINES approach by integrating the newly developed module named PINES and the ANN module previously contributed by Chen and Ferguson.⁸ We developed the PINES module in PLUMED2 using the PIV module previously contributed by Pipolo and Pietrucci^{57,58} as a template. In the ANN module, we added support for batch normalization. Biasing of the translationally, rotationally, and permutationally-invariant CVs discovered by PINES can be performed by any off-the-shelf CV biasing enhanced sampling technique. In this work, we use the PBMetaD PLUMED2 module contributed by Bonomi and Pfaendt-

ner⁶⁴ in PLUMED2 for enhanced sampling calculations. The new PINES and modified ANN modules along with a helper bash script for patching with PLUMED2 are available at <https://github.com/Ferg-Lab/pines.git>, and will be submitted for incorporation into a future release of PLUMED2.⁶⁶ All the data and PLUMED input files required to perform the molecular dynamics simulations of the three systems have been hosted on PLUMED-NEST, the public repository of the PLUMED consortium,⁸⁸ under project ID **plumID:23.024** (<https://www.plumed-nest.org/eggs/23/024/>).

2.2.1 13-particle Ar cluster

We consider a system of 13 Ar atoms to demonstrate the utility of PINES in driving assembly and disassembly of a simple Lennard-Jones cluster of identical particles.⁸⁹ The Ar atoms were modeled in vacuum using Lennard-Jones non-bonded interactions with $\sigma = 0.34$ nm and $\epsilon = 1$ kJ mol⁻¹.^{90,91} Particles were randomly placed in a $2.5 \times 2.5 \times 2.5$ nm³ cubic box implementing periodic boundary conditions. The system was subjected to steepest descent energy minimization to eliminate forces larger than 1 kJ mol⁻¹ nm⁻¹. After energy minimization, a short equilibration simulation was performed in a NVT ensemble for 100 ps. The classical mechanical equations of motion were integrated using the leapfrog algorithm⁹² with a 2 fs time step. Particle velocities were initialized according to the Maxwell-Boltzmann distribution at 50 K. Simulations were conducted in the NVT ensemble in which temperature is maintained using a velocity-rescaling thermostat⁹³ employing a time constant of 0.1 ps. Lennard-Jones interactions were smoothly shifted to zero at a cutoff of 1.0 nm. The unbiased production run for the initial PINES round was conducted for 25 ns and frames were saved at a frequency of 1 ps for training the autoencoder in the subsequent PINES round. We conducted 25 ns of PBMetaD simulations in subsequent PINES iterations at a temperature of 50 K employing an initial Gaussian height of $W = 0.5$ kJ mol⁻¹, a Gaussian width of $\sigma = 0.1$ for each discovered CV, a bias factor of $\gamma = 10$, and deposition pace of 500 steps (1 ps). The PIV featurization contained one Ar-Ar interaction block with 78 elements

corresponding to all possible pairwise distances. Simulations were performed on 5×2.40 GHz Intel Xeon Gold 6148 CPU cores and $1 \times$ NVIDIA TITAN V GPU, achieving execution speeds of ~ 120 ns per day.

2.2.2 NaCl ion pair in water

An NaCl ion pair in water was used as a classic test system to assess the capacity of PINES to discover solvent-inclusive CVs governing the association/dissociation of the ion pair and use these to efficiently accelerate sampling⁹⁴⁻⁹⁸. We utilized the Amber99SB-ILDN forcefield⁹⁹ with Joung-Cheatham^{100,101} ion parameters. The Na^+ and Cl^- ions were solvated to a density of 1 g cm^{-3} with 508 TIP3P¹⁰² water molecules in a cubic simulation box of size $2.5 \times 2.5 \times 2.5 \text{ nm}^3$. A cutoff of 1 nm was applied for the Lennard-Jones interactions. Long-range electrostatics were treated by particle mesh Ewald summation^{103,104} employing a real-space cutoff of 1.0 and 0.16 nm Fourier grid spacing that were optimized during runtime. The neighbor list was generated using the Verlet method and was updated every 10 steps (20 fs). An energy minimization procedure with steepest-descent algorithm was applied to relax the system after solvation by adjusting the configuration to have forces between particles with less than $1 \text{ kJ mol}^{-1} \text{ nm}^{-1}$. After energy minimization, a short equilibration simulation was performed in a NVT ensemble for 100 ps. The leapfrog algorithm⁹² with a timestep of 2 fs was applied for propagating the equations of motion in the simulations. The initial velocities for the equilibration run were generated using the Maxwell-Boltzmann distribution at 300 K. A velocity-rescale thermostat⁹³ was applied for controlling the temperature during the equilibration run. For the production run, the final configuration and velocities from the equilibration run were used to initiate the simulation. Frames at a frequency of 2 ps were saved from the production run for training the autoencoder in the subsequent PINES rounds. Our PIV vectors consider the distance between the ion pair and the ion distances to the O and H atoms of the 10 most proximate water molecules such that the PIV featurization contained five interaction blocks with a total of 61 elements distributed across the blocks as

1 Na-Cl, 10 Na-O, 20 Na-H, 10 Cl-O, and 20 Cl-H elements.

The initial production unbiased simulation was performed for 50 ns. For PBMetaD in subsequent rounds, we set the initial gaussian height as $W = 1.2 \text{ kJ/mol}$, a Gaussian width of $\sigma = 0.1$ for each of the discovered CVs, a bias factor of $\gamma = 20$, and a deposition pace of 500 timesteps (1 ps). Each simulation in the PINES rounds with enhanced sampling along discovered CVs using PBMetaD was initially performed for 15 ns in a NVT ensemble at 300 K but was extended until all the convergence metrics for PBMetaD were satisfied. Simulations were performed on $5 \times 2.40 \text{ GHz}$ Intel Xeon Gold 6148 CPU cores and $1 \times$ NVIDIA TITAN V GPU, achieving execution speeds of $\sim 50 \text{ ns per day}$.

2.2.3 Aqueously solvated *n*-pentatetracontane ($\text{C}_{45}\text{H}_{92}$) polymer chain

An aqueously solvated *n*-pentatetracontane ($\text{C}_{45}\text{H}_{92}$) polymer chain was used to demonstrate the ability of PINES to learn solvent-inclusive CVs for hydrophobic collapse and use these CVs to enhance sampling of the elongation/collapse transition.⁶¹ The initial structure of the *n*-pentatetracontane alkane chain was modeled using the TraPPE-UA^{105,106} potential that models the chain as a string of CH_3/CH_2 united atoms. The system was solvated with 11,188 SPCE¹⁰⁷ water molecules in a cubic simulation box of size $7.0 \times 7.0 \times 7.0 \text{ nm}^3$. A cutoff of 1 nm was applied for the Lennard-Jones interactions. Long-range electrostatic interactions are handled by particle mesh Ewald method^{103,104} employing a real-space cutoff of 1.0 and 0.16 nm Fourier grid spacing that were optimized during runtime. The neighbor list was generated using the Verlet method^{108,109} and is updated every 20 steps (40 fs). An energy minimization procedure with steepest-descent algorithm was applied to relax the system after solvation by adjusting the configuration to have forces between particles with less than $1 \text{ kJ mol}^{-1} \text{ nm}^{-1}$. After energy minimization, a short equilibration simulation was performed in a NVT ensemble for 100 ps. The leapfrog algorithm⁹² with a timestep of 2 fs was applied for propagating the equations of motion. The initial velocities for the equilibration run were generated using the Maxwell-Boltzmann distribution at 300 K. A velocity-rescale

thermostat⁹³ was applied for controlling the temperature during the equilibration run. For the production run, the final configuration and velocities from the equilibration run were used to initiate the simulation. The unbiased production run in the initial PINES round was conducted for 100 ns with a timestep of 2 fs. Frames at a frequency of 2 ps were saved from the production run for training the autoencoder in the subsequent PINES rounds. The PIV featurization contained 45 interaction blocks with a total of 108 elements distributed across the blocks. The blocks are characterized by the intramolecular interactions between nine evenly spaced CH_2 united atoms from position 3 to 43 in the chain, yielding 36 blocks each containing only a single element, as well as the intermolecular interactions with their respective eight closest water oxygens, yielding nine blocks each containing eight elements. We note that the $\text{CH}_2\text{-CH}_2$ united atom blocks could be represented as 18 blocks of two elements to mod out the symmetry associated with the head-to-tail inversion symmetry of the alkane chain, however we choose not to do this here.

For the PBMetaD calculations, we set the initial gaussian height as $W = 1.2 \text{ kJ mol}^{-1}$, a Gaussian width of $\sigma = 0.1$ for each discovered CV, a bias factor of $\gamma = 20$, and a deposition pace of 500 timesteps (1 ps). Each simulation in the PINES rounds with enhanced sampling along discovered CVs using PBMetaD was initially performed for 30 ns in a NVT ensemble at 300 K but was extended until all the convergence metrics for PBMetaD were satisfied. Simulations were performed on $5 \times 2.40 \text{ GHz}$ Intel Xeon Gold 6148 CPU cores and $1 \times$ NVIDIA TITAN V GPU, achieving execution speeds of $\sim 10 \text{ ns per day}$.

3 Results and Discussion

We demonstrate the general utility and applicability of PINES to discover translationally, rotationally, and permutationally invariant CVs for enhanced sampling in applications to three molecular systems: (i) self-assembly of a 13-atom Ar cluster, (ii) assembly/disassembly of a NaCl ion pair in water, and (iii) hydrophobic collapse of a $\text{C}_{45}\text{H}_{92}$ *n*-pentatetracontane

polymer chain.

3.1 Self-assembly of a 13-particle Ar cluster

We first demonstrate the application of PINES to the self-assembly of a 13-particle Ar cluster in vacuum. We hypothesize that PINES can learn permutationally-invariant multi-body collective variables associated with the assembly and disassembly of the identical Ar atoms between solid-like, liquid-like, and vapor-like phases. After some trial-and-improvement exploration, we chose to perform the simulations at a temperature of 50 K, which we observed to provide good unbiased sampling of the solid, liquid, and vapor-like phases and therefore access to an unbiased ground truth against which to compare the biased sampling calculations under PINES. As in our previous work with MESA, we determine the dimensionality of the latent space by training independent networks with 1-8 nodes in the bottleneck layer of autoencoder and utilizing the L-method by Salvador and Chan⁷³ to identify a knee in the fraction of variance explained (FVE) curve (Fig. S2). With this method, we identified a latent space of dimensionality $d=3$.

In Fig. 3a we show the round-by-round sampling as 2D projections along the three CVs ξ_0 , ξ_1 , and ξ_2 learned from the terminal PINES network. The initial unbiased simulation used to seed PINES (gray) provided a good initial coverage of the accessible configurational space and each subsequent biased round (different colors) extends the frontier into new unsampled regions of phase space around the periphery of the explored region and probes higher free energy configurations. We present in Fig. 3b the reweighted (i.e., unbiased) free energy surfaces (FES) associated with each sampling round that better expose the variations in sampling density over the three CVs. In particular, we observe that the biased sampling identifies a low- ξ_0 metastable minimum that was not resolved by the unbiased calculations in Round 0. As we discuss below, this minimum at $(\xi_0 \approx -2, \xi_1 \approx 0, \xi_2 \approx 1)$ corresponds to a rarefied vapor-like state that is incompletely explored in the unbiased calculations.

We assess convergence of the iterative cycles of CV learning and biased sampling by com-

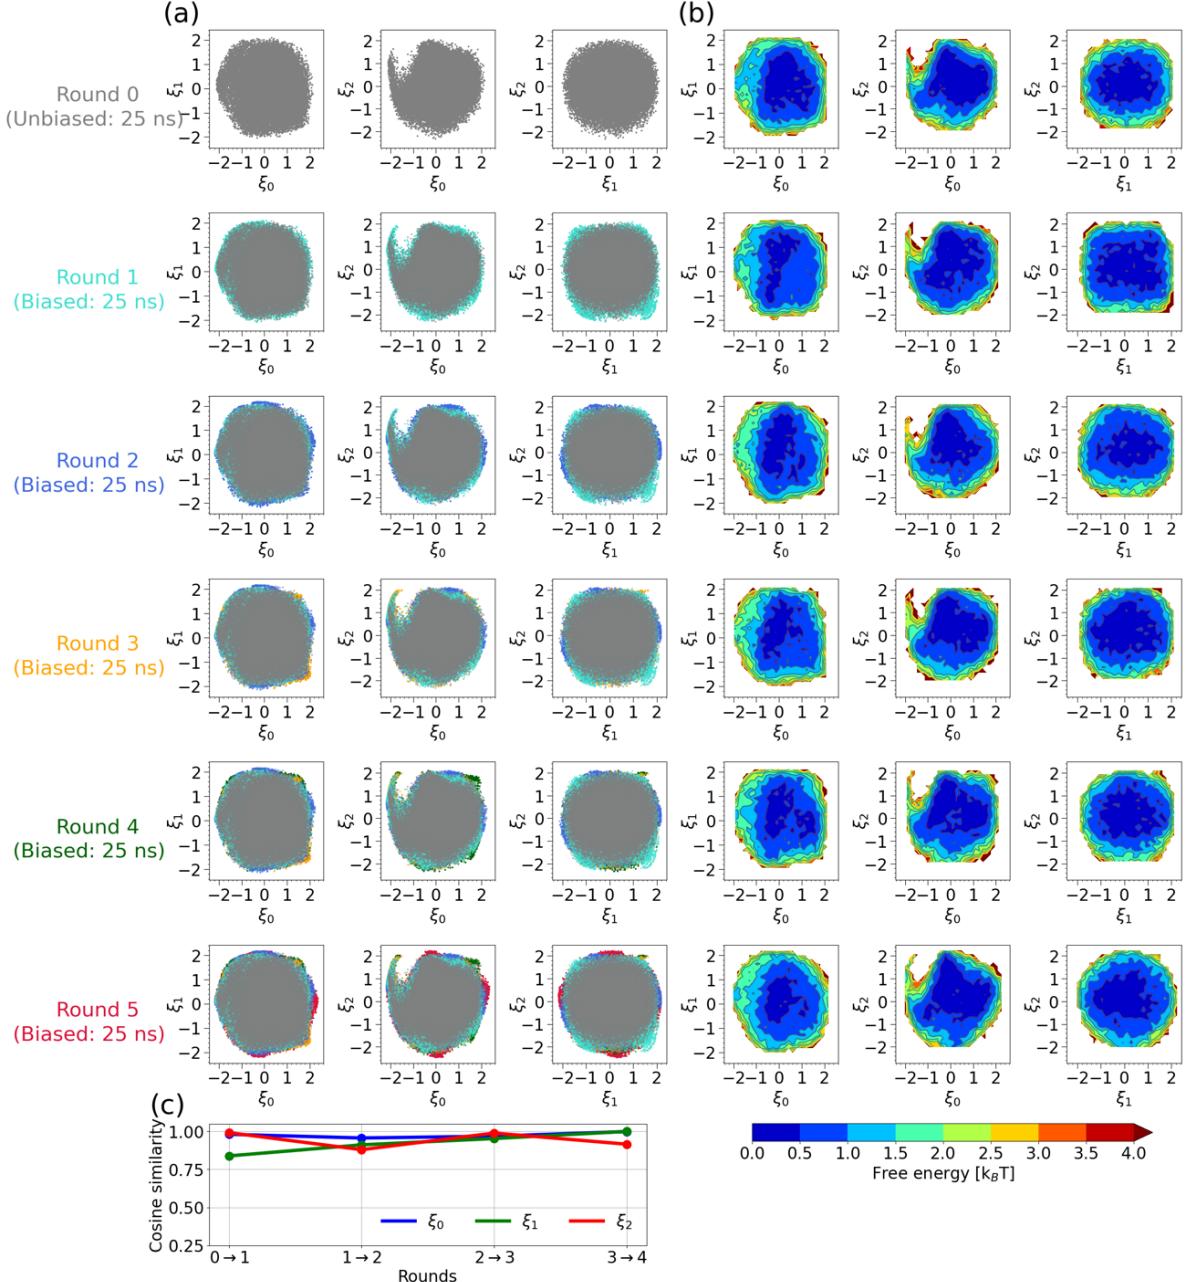


Figure 3: Application of PINES to a 13-particle Ar system in vacuum at 50 K. (a) Each column shows a 2D embedding of the round-by-round accrued sampling projected into the converged CVs ξ_0 , ξ_1 , and ξ_2 at the terminal Round 5. The unbiased simulation quite thoroughly samples the thermally accessible phase space, yet each biased round advances the frontier into unsampled regions on the periphery of the explored region. (b) Corresponding free energy surfaces (FES) computed by reweighting of the biased simulation data collected in the current round of sampling. (c) Cosine similarity between corresponding CVs in successive rounds. The high cosine similarity value indicates rapid convergence of the iterative CV discovery and enhanced sampling cycles within the five rounds of PINES.

puting the cosine similarity between corresponding CVs in successive rounds. As illustrated in Fig. 3c, we find each of the three learned CVs to quickly stabilize with round-to-round cosine similarity scores in excess of $\rho = 0.9$ for Round 1 and above. Together with the relatively modest expansion of the configurational phase space (Fig. 3a) and stabilization of the FES (Fig. 3b), we declare the PINES procedure to be converged within the five rounds.

It is informative to analyze the terminal learned CVs to provide insight about the high variance emergent collective modes within the system. By extracting the weights of the trained encoder, we perform layerwise matrix multiplication to compute a single scalar weight encapsulating the extent to which each PIV element contributes to each of the learned CVs integrated over all possible pathways through the network. Although this analysis does not account for correlations between the various PIV elements, it is useful in providing a gross assessment of how strongly each permutationally invariant pairwise distances influences each of the learned CVs. We present in Fig. 4a a graphical illustration of the contribution of each PIV element v_k to each of the the terminal learned CVs $\{\xi_0, \xi_1, \xi_2\}$, and also a cumulative sum over these weights indicating the net contribution and cooperative impact over multiple PIV elements at increasing pairwise distances. The absence of strong outliers in the bar charts indicates that none of the $\binom{13}{2} = 78$ PIV elements substantially dominate the values of the three learned CVs. A similar trend is manifested in the cumulative sums that indicate an approximately linear trend indicative of similar magnitude and similar signed contributions out to PIV element 35 or so. Recalling that the pairwise distances are passed through a switching function (Fig. 2) the larger pairwise distances appear earlier in the PIV, and for the shorter range interactions beyond element 35, we see a more interesting and challenging to interpret structure in the sign and magnitude of the PIV weights. The identical nature of all particles in this system and necessarily cooperative nature of assembly and disassembly is consistent with the absence of any single dominant contribution from the PIV.

We complement this bottom-up analysis of the learned CVs with a top-down interrogation in which we color projections of the Round 5 simulation data into the terminal CVs by

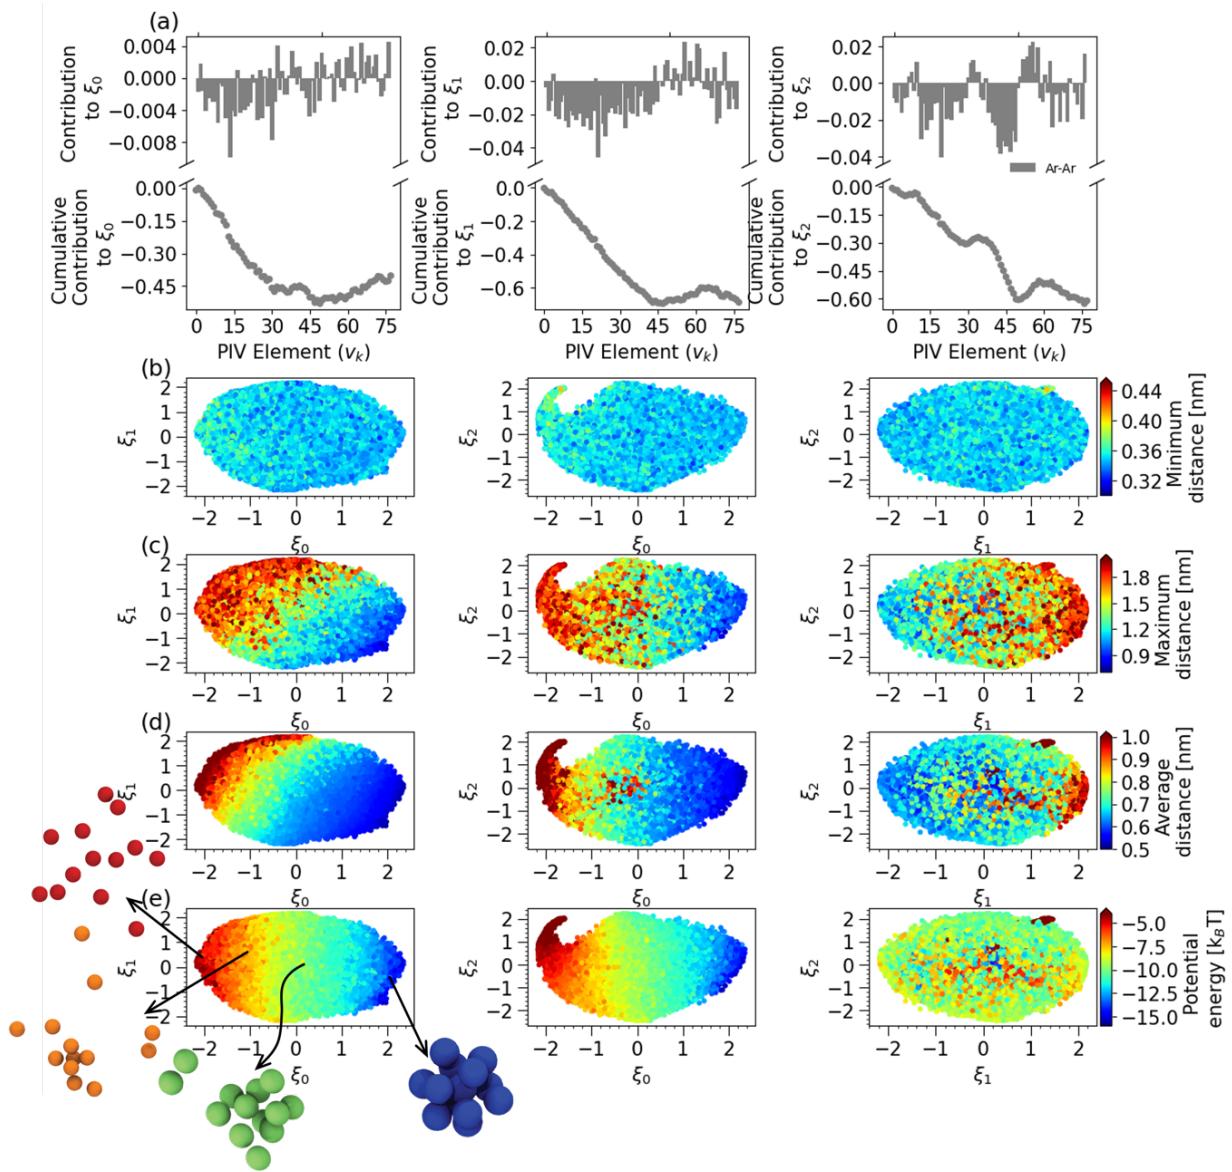


Figure 4: Bottom-up and top-down analysis of terminal learned CVs for the 13-particle Ar system. (a) Split plots show the signed contribution from each PIV element to each CV $\{\xi_0, \xi_1, \xi_2\}$ as a bar chart on top and as a cumulative sum of the signed contributions on the bottom. The magnitude of contributions on the bar chart emphasize the importance of individual PIV elements while the slopes in different regions of the scatter plot illuminate the concerted effect of atom groups on each CV. (b-e) Projections of the terminal Round 5 sampling data into the terminal CVs and colored by physically interpretable candidate variables. The potential energy, which is known a priori to be the putative CV for a Lennard-Jones fluid, correlates well with ξ_0 . Selected snapshots are presented to show the smooth transition across ξ_0 from the vapor-like phase (red, orange) to the condensed phase (green), and finally to the highly ordered solid-like phase (blue). The gradients in ξ_0 and ξ_1 with respect to the average and maximum pairwise distances indicates that PINES has learned these as important variables in driving assembly/disassembly.

candidate physical order parameters (Fig. 4b-e). In doing so, we find ξ_0 to be strongly correlated with the potential energy of the system (Fig. 4e). Since the potential energy is a sum over pairwise interactions, this makes ξ_0 a good correlate for the degree of association of the 13-particle cluster, with the compact, low-potential energy solid-like state residing at high ξ_0 , the dissociated, high-potential energy vapor-like state at low ξ_0 , and the liquid-vapor-like, partially associated / partially dissociated states at intermediate ξ_0 . The average pairwise distance (Fig. 4d) and maximum pairwise distance (Fig. 4c) are both well-correlated with a linear combination of ξ_0 and ξ_1 , with the “red hook” apparent in the ξ_0 - ξ_2 projections containing the most highly dissociated and rarefied vapor-like configurations. The minimum pairwise distance fails to show strong correlations with any of the three CVs (Fig. 4b). In sum, the three permutationally-invariant CVs learned by PINES as important collective modes to drive sampling of the configurational exploration of the system reflect the potential energy, compactness, and average/maximum pairwise interatomic distance.

Finally, in Fig. 5 we present our reweighted FES collected from our final round of enhanced sampling in the converged terminal learned CVs. We emphasize that these data were collected from enhanced sampling calculations conducted in the learned CVs (cf. Fig. 3), but that after sampling is complete, we are at liberty to reweight and project these data into arbitrary order parameters of our choosing for the purposes of expository illustration. The 1D FES constructed in the average pairwise distance clearly exposes the solid-like, filled icosahedron as the global free energy minimum of the system at an average distance of 0.62 nm (Fig. 5a). At higher average distances we enter into a liquid/vapor-like phase in which some fraction of the particles in the system are associated, and then terminate in a vapor-like phase with all particles dispersed lying within a weak free energy minimum at 1.17 nm approximately $4 k_B T$ higher in free energy than the solid-like phase. The 2D FES constructed in the potential energy and maximum pairwise distance better illustrates the expansion of accessible phase space for the more disordered liquid-like and vapor-like states compared to the solid-like filled icosahedron (Fig. 5b).

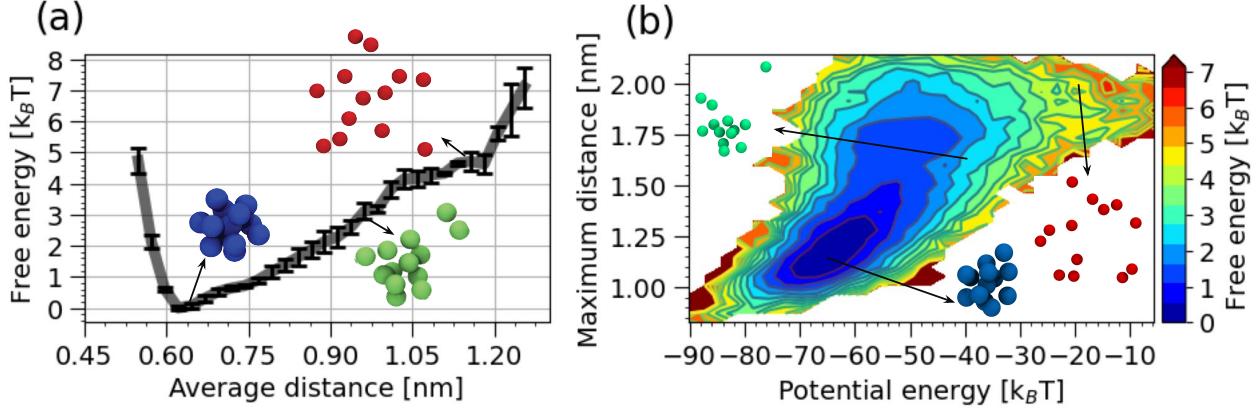


Figure 5: Reweighted free energy surfaces from the terminal round of enhanced sampling for the 13-particle Ar system projected into physically interpretable order parameters. (a) A 1D FES projected into the average pairwise interatomic distance. Representative snapshots show that the global free energy minimum to be occupied by a solid-like, filled icosahedron with an average pairwise distance of 0.62 nm. (b) A 2D FES projected into the potential energy and maximum pairwise interatomic distance. Representative snapshots illustrate the solid-like global free energy minimum and liquid-like and vapor-like phases.

The location and relative stability of the solid, liquid, and vapor-like phases is a strong function of the box size and temperature. In particular, the $2.5 \times 2.5 \times 2.5 \text{ nm}^3$ box places strict finite size effects on the maximum possible pairwise distances available to the particles in the vapor-like phase and limits the phase space volume accessible to it. The temperature also changes the relative balance of energetic and entropic contributions and modulating this away from the 50 K temperature at which our simulations were conducted can strongly modify the underlying free energy surface. To test the transferability of the converged CVs learned by PINES in a $2.5 \times 2.5 \times 2.5 \text{ nm}^3$ box at 50 K, we conducted enhanced sampling calculations in these CVs to recover the FES in an 8-fold larger $5 \times 5 \times 5 \text{ nm}^3$ simulation box (Fig. S3) and at temperatures of 25 K and 100 K (Fig. S4). Pleasingly, the CVs learned by PINES at $2.5 \times 2.5 \times 2.5 \text{ nm}^3$ and 50 K were capable of efficient exploration of the thermally-accessible configurational phase under these new conditions without any additional retraining. In particular, the enhanced sampling calculations had no difficulty in sampling the thermodynamically unfavorable solid-like phase in the large $5 \times 5 \times 5 \text{ nm}^3$ box (Fig. S4) or in traversing the $\sim 40 \text{ } k_B T$ difference between the solid-like and vapor-like phases in the 25 K

low temperature calculations (Fig. S4).

Taken together, this application of PINES to a simple 13-particle Ar system demonstrates that the method is capable of learning translationally, rotationally, and permutationally invariant CVs, using these CVs to efficiently drive PBMetaD enhanced sampling calculations, and furnishing transferable order parameters capable of driving sampling at other thermodynamic state points. We now proceed to test the method in applications to multi-component systems with more complex, rugged, and challenging-to-sample free energy surfaces.

3.2 Association/dissociation of a NaCl ion pair in water

We now apply PINES to discover CVs to drive the association/dissociation of a NaCl ion pair in water at 300 K. It is well known that solvent degrees of freedom play an important role in mediating this process via many-body collective rearrangements of solvation shell water molecules.^{94–98} We hypothesize that PINES should be able to identify permutationally-invariant ion and solvent-inclusive CVs associated with the indistinguishable solvent molecules and use these to drive sampling of the solvated ion pair dynamics.

We identified a $d=3$ dimensional latent space from an initial 50 ns unbiased simulation trajectory and confirm that this dimensionality remained unchanged in each of the five successive rounds of PINES (Fig. S2). We illustrate in Fig. 6a the expansion of the sampled phase space under each PINES iteration. The initial unbiased simulation provides good initial sampling of the central portion of the latent space (gray) but each successive biased round (different colors) provide additional sampling of the periphery. In particular, Rounds 1-3 (blue, yellow, green) show increased exploration of the CV space in half or less of the simulation time compared to the unbiased run and indicate that the CVs are accelerating sampling of novel states. Rounds 4-5 (red, black), despite being longer than the previous runs, offer diminishing returns in additional configurational sampling. (We recall that the run length in each round is dictated by convergence criteria of the PBMetaD simulations (cf. Sec. 2.1.4).)

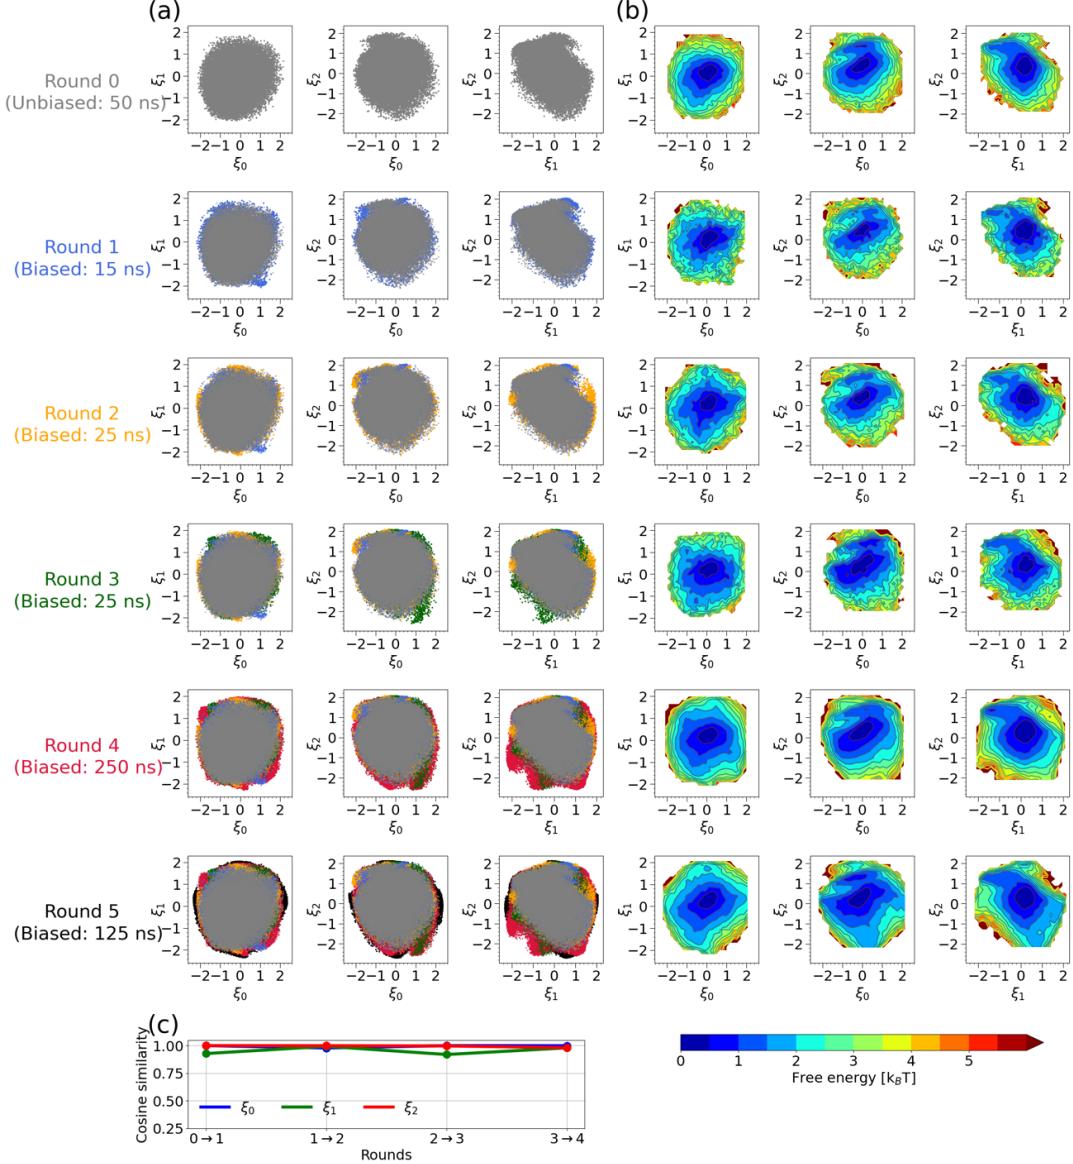


Figure 6: Application of PINES to a NaCl ion pair in water at 300 K. (a) Each column shows a 2D embedding of the round-by-round accrued sampling projected into the converged CVs ξ_0 , ξ_1 , and ξ_2 at the terminal Round 5. The unbiased simulation quite thoroughly samples the thermally accessible phase space, yet each biased round advances the frontier into unsampled regions on the periphery of the explored region. Rounds 1-3 show broader exploration of the CV space in half or less of the simulation time compared to the unbiased run and indicate that the CVs are accelerating sampling of novel states. Rounds 4-5, despite being much longer than the previous biased runs, provide diminishing returns in the exploration of new areas of CV space. (b) Corresponding free energy surfaces (FES) computed by reweighting of the biased simulation data collected in the current round of sampling. (c) Cosine similarity between corresponding CVs in successive rounds. The high cosine similarity values between all neighboring rounds indicate that the CVs discovered in the initial unbiased round are not substantially altered by learning over the biased data and the PINES protocol rapidly converges.

In Fig. 6b, we present the corresponding reweighted FES for each round. We observe visually apparent expansion of the FES into high-free energy peripheral regions and changes in the free energy distributions through Rounds 0-4. The FES between Rounds 4-5, however, are very similar in terms of the explored regions and free energy contours. As we discuss below, the core of the latent space that was readily sampled by the unbiased simulation corresponds to dissociated configurations with average levels of hydration. The peripheral states that require biased simulations for adequate exploration are associated with higher free energy states corresponding to metastable configurations such as the solvent-shared pair.

We present in Fig. 6c the cosine similarity between CVs in successive PINES rounds. In this case we find that the CVs almost immediately stabilize, with those learned from the unbiased simulation data changing relatively little over the course of the iterative CV discovery and enhanced sampling protocol and always remaining above $\rho = 0.9$. Taken together with the observed trends in the latent space and FES, we surmise that the initial CVs provide a good global parameterization of the system configurational space, but require approximately three rounds of enhanced sampling and minor CV refinement to attain convergence in sampling of the thermally accessible phase space.

We elected to use PIVs that considered the 10 closest water molecules to each ion (Section 2.2.2), with this number selected to be sufficiently large to track all water molecules within approximately the first hydration shell. To assess the dependence of the number of water molecules included within the PIV, we reran our analysis using the 15, 25, and 35 most proximate water molecules to explore the effect of including water molecules in higher hydration shells. We find that both the CVs discovered by the PINES protocol (Fig. S5) and that the exploration of phase space within the CVs (Fig. S6) are quite insensitive to the number of proximal water molecules included in the PIV. This suggests that the approximately 10 proximal water molecules necessary to capture the important solvation effects within the first hydration shell expose sufficient information for the discovery of CVs capable of characterizing and driving good global sampling of the relevant configurational motions

of the system. In general, the appropriate number of closest solvent molecules to consider is anticipated to be context and system dependent, and a parametric analysis such as this can be useful in identifying an appropriate threshold.

In Fig. 7a, we expose the integrated weights linking each of the 61 PIV elements to each of the three CVs. Inspection of the individual (upper) and cumulative (lower) weights indicates that particular interactions appear to contribute more significantly to each of the three CVs. In particular, PIV elements corresponding to Na-O and Cl-O distances make large magnitude contributions to ξ_0 , whereas Na-H and Cl-H distances seemingly contribute more strongly to ξ_1 , and Na-O makes the largest contribution to ξ_2 . Interestingly, the Na-Cl interaction is not a top contributor to any CVs, suggesting that the information about the Na-Cl state may be delocalized among the information on solvation shell water distances.

We further explore these trends in Fig. 7b-f where we color the latent space by candidate physical variables. These plots are illuminating in showing that ξ_0 is well-correlated with the average Cl-O and Cl-H distance, ξ_1 with the average Na-O and Na-H distance, and ξ_2 with the Na-Cl distance. This analysis offers a straightforward interpretation of the three learned CVs as corresponding to Cl ion hydration (ξ_0), Na ion hydration (ξ_1), and Na-Cl proximity (ξ_2).

We present in Fig. 8 the reweighted FES collected from our final round of PINES enhanced sampling projected into physical order parameters and annotated with representative molecular configurations. A 1D FES constructed in the Na-Cl distance clearly shows the contact pair ($r_{\text{Na-Cl}} = 0.26$ nm), solvent-shared ($r_{\text{Na-Cl}} = 0.51$ nm), and disassociated ($r_{\text{Na-Cl}} \gtrsim 0.6$ nm) states (Fig. 8a). The free energy barrier between the contact pair and solvent-shared state lies at $r_{\text{Na-Cl}} = 0.36$ nm with a height relative to the contact pair of $(4.9 \pm 0.1) k_B T$. This stands in good agreement with reported literature values of $4.8\text{-}5.3 k_B T$ and provides good support for the accurate sampling of the configurational space and FES generated by PINES.^{110,111} We present in Fig. 8b a 2D FES constructed in the Na-Cl distance and the solvent coordination number of the Na ion n_{Na} . This projection is illuminating in showing a

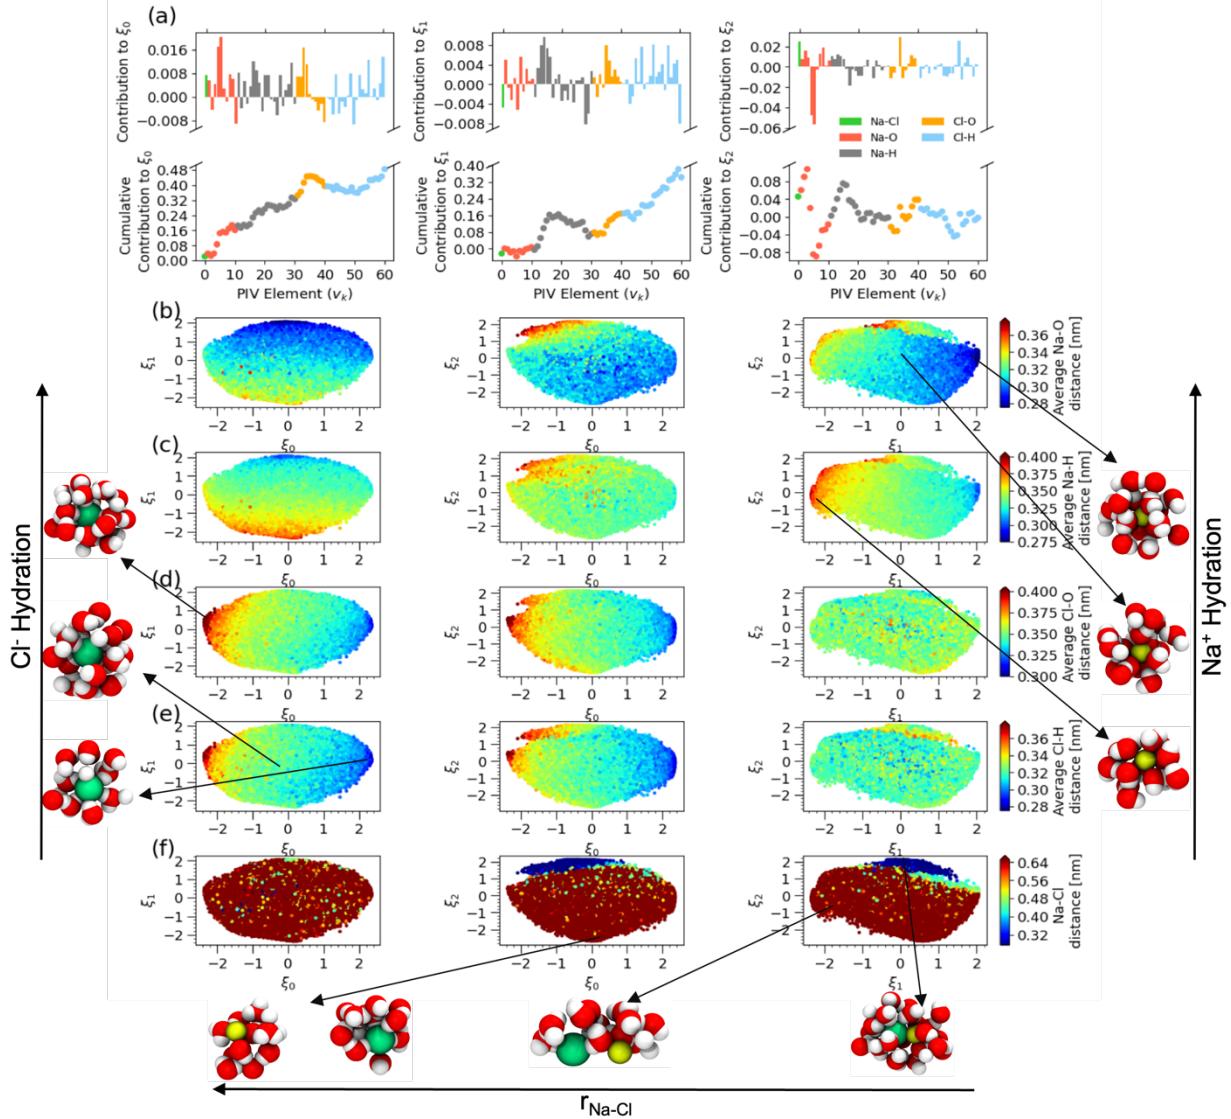


Figure 7: Bottom-up and top-down analysis of terminal learned CVs for the NaCl ion pair in water. (a) Split plots show the signed contribution from each PIV element to each CV $\{\xi_0, \xi_1, \xi_2\}$ as a bar chart on top and as a cumulative sum of the signed contributions on the bottom. The magnitude of contributions on the bar chart emphasize the importance of individual PIV elements while the slopes in different regions of the scatter plot illuminate the concerted effect of atom groups on each CV. (b-f) Projections of the terminal Round 5 sampling data into the terminal CVs and colored by physically interpretable candidate variables. ξ_0 is well-correlated with the average Cl-O and Cl-H distance, ξ_1 with the average Na-O and Na-H distance, and ξ_2 with the Na-Cl distance. Three distinct color bands are apparent in panel (f) corresponding to the contact pair (blue), the solvent-shared (green), and disassociated (red) states (cf. Fig. 8c).

splitting of the contact pair at $r_{\text{Na-Cl}} = 0.26$ nm into two basins with coordination numbers of $n_{\text{Na}} = 4$ and $n_{\text{Na}} = 5$. The FES suggests that the transition to the solvent-shared state at

$r_{\text{Na-Cl}} = 0.51 \text{ nm}$ appears to proceed from the $n_{\text{Na}} = 5$ state and then drops into the lower free energy $n_{\text{Na}} = 6$ state after separation has completed, although dynamical analyses such as transition path sampling would be required to test this conjecture.

To assess the importance of the solvent features learned by PINES in sampling the thermally relevant phase space, we present in Fig. S7 a comparison of the 1D FES in the Na-Cl distance and the 2D FES in the Na-Cl distance and the solvent coordination number of the Na ion n_{Na} computed by PINES and by well-tempered metadynamics enhanced sampling in the Na-Cl distance alone. This analysis shows that while metadynamics in the Na-Cl distance alone accurately recovers the 1D FES in the Na-Cl distance in good accord with PINES, it fails to expose low-free energy regions associated with highly solvated states of the Na ion that are only exposed by driving sampling in solvent-based collective variables learned by PINES.

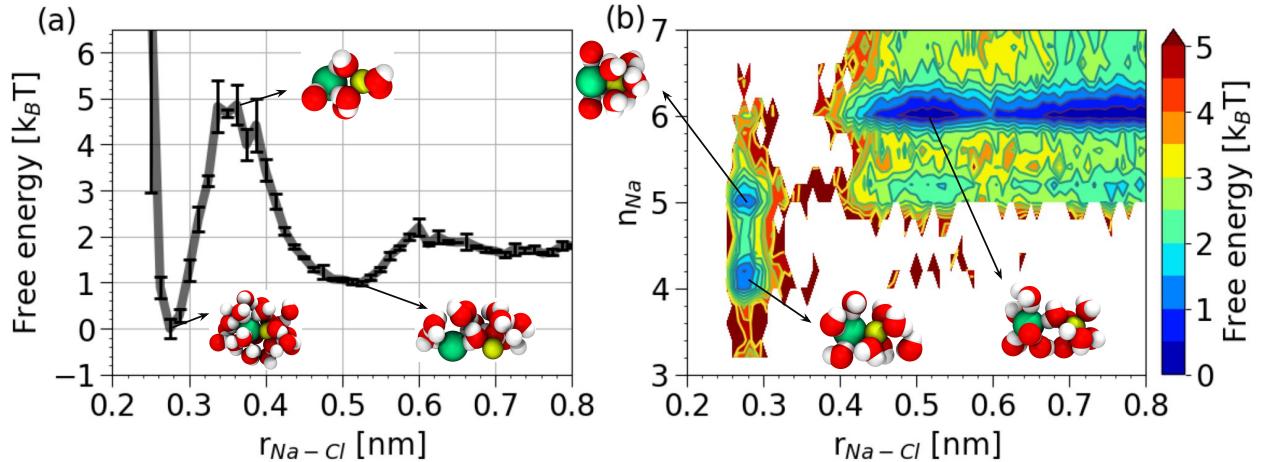


Figure 8: Reweighted free energy surfaces from the terminal round of enhanced sampling for the NaCl ion pair in water projected into physically interpretable order parameters. (a) A 1D FES projected into the Na-Cl distance. The landscape clearly illustrates the contact pair ($r_{\text{Na-Cl}} = 0.26 \text{ nm}$), solvent-shared ($r_{\text{Na-Cl}} = 0.51 \text{ nm}$), and disassociated ($r_{\text{Na-Cl}} \gtrsim 0.6 \text{ nm}$) states. The free energy barrier lying between the contact pair and solvent-shared states is $(4.9 \pm 0.1) k_B T$ as measured from the contact pair, and involves a collective rearrangement of solvent shell water molecules to accommodate the separating ion pair.^{97,112} Uncertainties in the free energy correspond to standard errors estimated by five-fold block averaging. (b) A 2D FES projected into the Na-Cl distance and the solvent coordination number of the Na ion n_{Na} . The contact pair comprises multiple solvation states of the Na ion. Transitions to the solvent separated pair appear to proceed most readily through the $n_{\text{Na}} = 5$ contact pair state.

To assess the transferability of the converged CVs learned by PINES for the NaCl ion pair association/dissociation process, we conducted enhanced sampling calculations in these CVs for a hypothetical divalent ion system generated by doubling the charge on each ion (Fig. S8). Whereas an unbiased simulation of the divalent ion pair is completely localized near the initial state corresponding to the contact pair because of the strong electrostatic attractions (Fig. S8a,b), PBMetaD enhanced sampling simulations in the monovalent CVs are capable of surmounting a $\sim 45 k_B T$ free energy barrier to drive sampling of the dissociated state (Fig. S8c,d). This analysis indicates the transferability of CVs learned over one system applied to another in which it is more challenging to drive sampling.

In sum, this application to a NaCl ion pair in water demonstrates the capacity of PINES to identify permutationally-invariant CVs incorporating information on the ion pair and local solvation environment and use these variables to efficiently drive enhanced sampling of ion pair association/dissociation. This illustrates the capacity of the method to learn permutationally-adapted CVs in a solvated system of chemically indistinguishable water molecules. The unbiased simulations used to seed the PINES campaign provided a relatively comprehensive sampling of the phase space and the learned CVs remained relatively unchanged over the course of the iterative search. In the next application, we demonstrate a case where the PINES enhanced sampling is instrumental in good sampling of phase space.

3.3 Hydrophobic collapse of an *n*-pentatetracontane ($C_{45}H_{92}$) polymer chain in water

In our final application, we consider hydrophobic collapse of an *n*-pentatetracontane ($C_{45}H_{92}$) polymer chain in water to demonstrate an application of PINES to a solvated polymer system. The dynamical collapse of hydrophobic polymers in water is known to be intimately related to drying transitions and the formation of water cavities.^{59,61,113,114} Solvent-centric CVs are therefore expected to play an important role in the collapse/extension transitions of the $C_{45}H_{92}$ polymer chain and, for such a long chain, it can also be challenging to comprehensively

explore the configurational phase space using unbiased simulations alone. This system is therefore anticipated to be a good test of the PINES methodology.

Once again, we identified a $d=3$ dimensional latent space from an initial 50 ns unbiased trajectory that remained unchanged over the course of the seven subsequent rounds of PINES (Fig. S2). We present in Fig. 9a the expansion of the sampled latent space over the PINES iterations. We see significant expansion into new, unexplored regions by the enhanced sampling calculations in Rounds 1-4, before sampling approximately converges and achieves only minor additional exploration in Rounds 5-7. We draw particular attention to the interior “hole” in the latent space apparent at $(\xi_0=-1, \xi_1=0)$ in Round 0 (gray) and which is filled in by the enhanced sampling trajectories in Round 1 (cyan) and Round 2 (blue). As we discuss below, this hole corresponds to a high-free energy transition region between collapsed and extended states of the chain that is only adequately sampled under PBMetaD biasing in the learned CVs. Fig. 9b presents analogous FES plots for each PINES round that expose the density distribution of points over the latent space and convergence of the landscapes over the PINES rounds. Fig. 9c presents the cosine similarity between CVs in successive PINES rounds. We observe significant changes in the learned CVs over Rounds 1-4 as they are exposed to additional configurational data supplied by the enhanced sampling calculations. The CVs stabilize in Rounds 5 and 6 with all CVs maintaining cosine similarity values of $\rho > 0.9$.

In Fig. 10, we gain insight into the learned CVs by presenting the network weights linking each of the 108 PIV elements to the three CVs and coloring the latent space embeddings by candidate physical variables. Interpreting the CVs for more complex systems can be challenging, but we find ξ_0 and ξ_1 to possess significant correlation with both the chain radius of gyration R_g and the number of waters N in the first solvation shell of the chain, defined by counting the number of water molecule O atoms within 0.6 nm of any united atom CH_2 or CH_3 bead of the alkane chain. ξ_2 is correlated, albeit non-monotonically, with the fraction of helical pitch f_{helix} . This observable, defined as $f_{\text{helix}} = r_{\text{CH}3-\text{CH}43} - (r_{\text{CH}3-\text{CH}23} +$

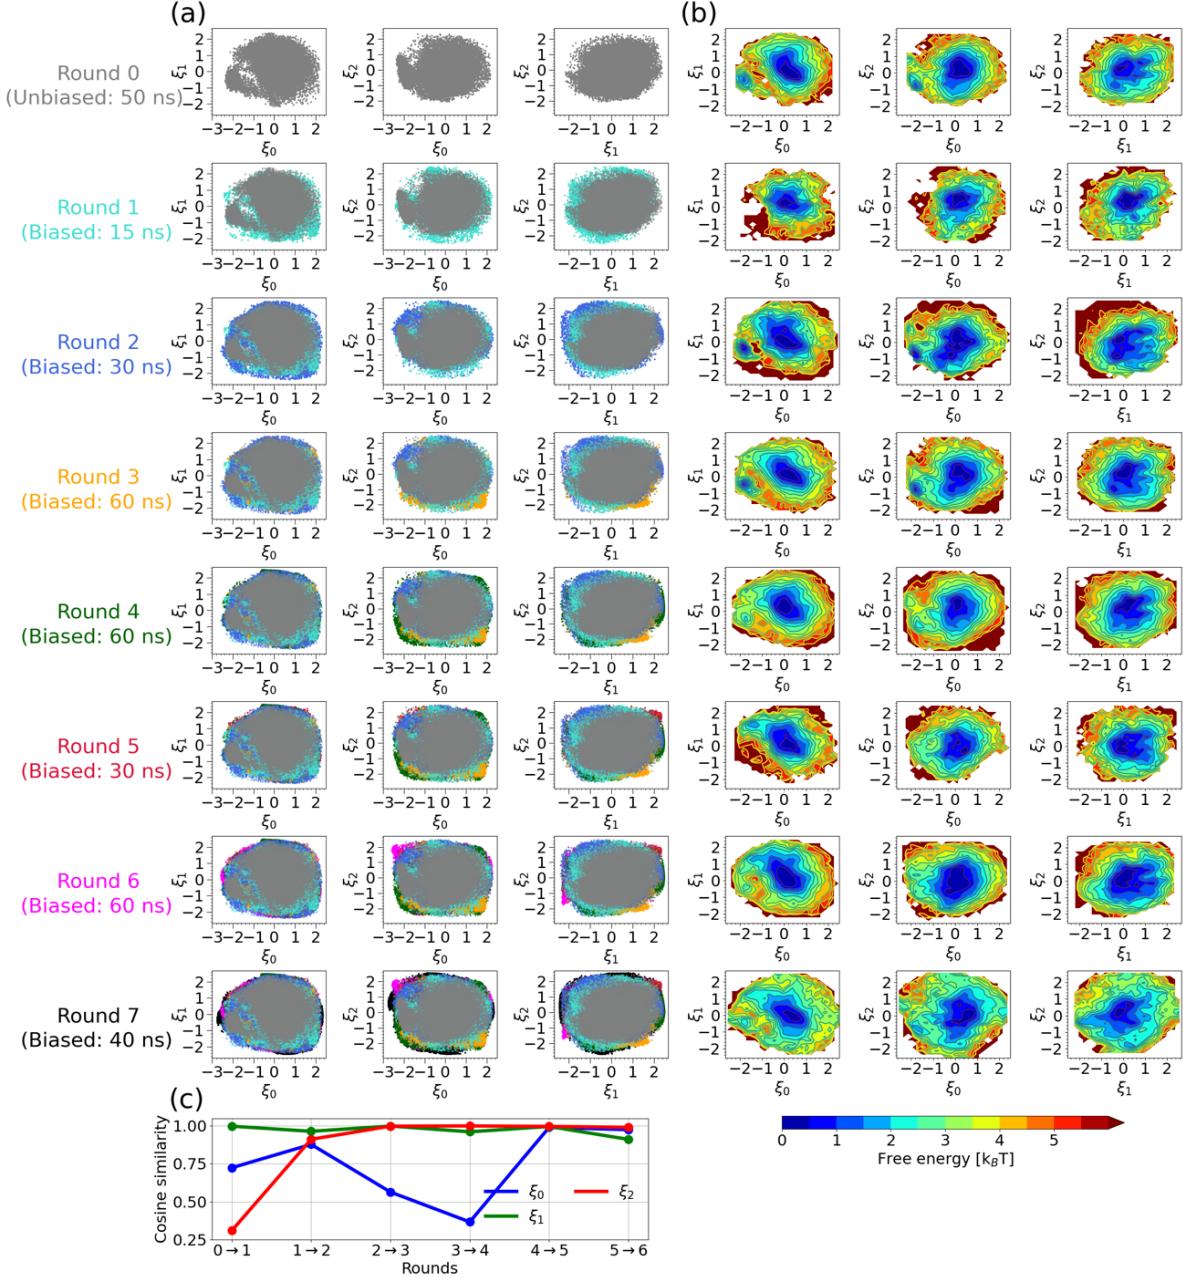


Figure 9: Application of PINES to an *n*-pentatetracontane ($C_{45}H_{92}$) chain in water at 300 K. (a) Each column shows a 2D embedding of the round-by-round accrued sampling projected into the converged CVs ξ_0 , ξ_1 , and ξ_2 at the terminal Round 7. Rounds 1-4 show significantly increased exploration of novel states compared to the unbiased run and indicate that the CVs are accelerating sampling of unexplored regions of phase space. In Rounds 5-6 the sampling of the CV space begins to saturate and suggests that the CVs have converged. (b) Corresponding free energy surfaces (FES) computed by reweighting of the biased simulation data collected in the current round of sampling. (c) Cosine similarity between corresponding CVs in successive rounds. The CVs change significantly round-to-round over Rounds 1-4 before converging in Rounds 5-7.

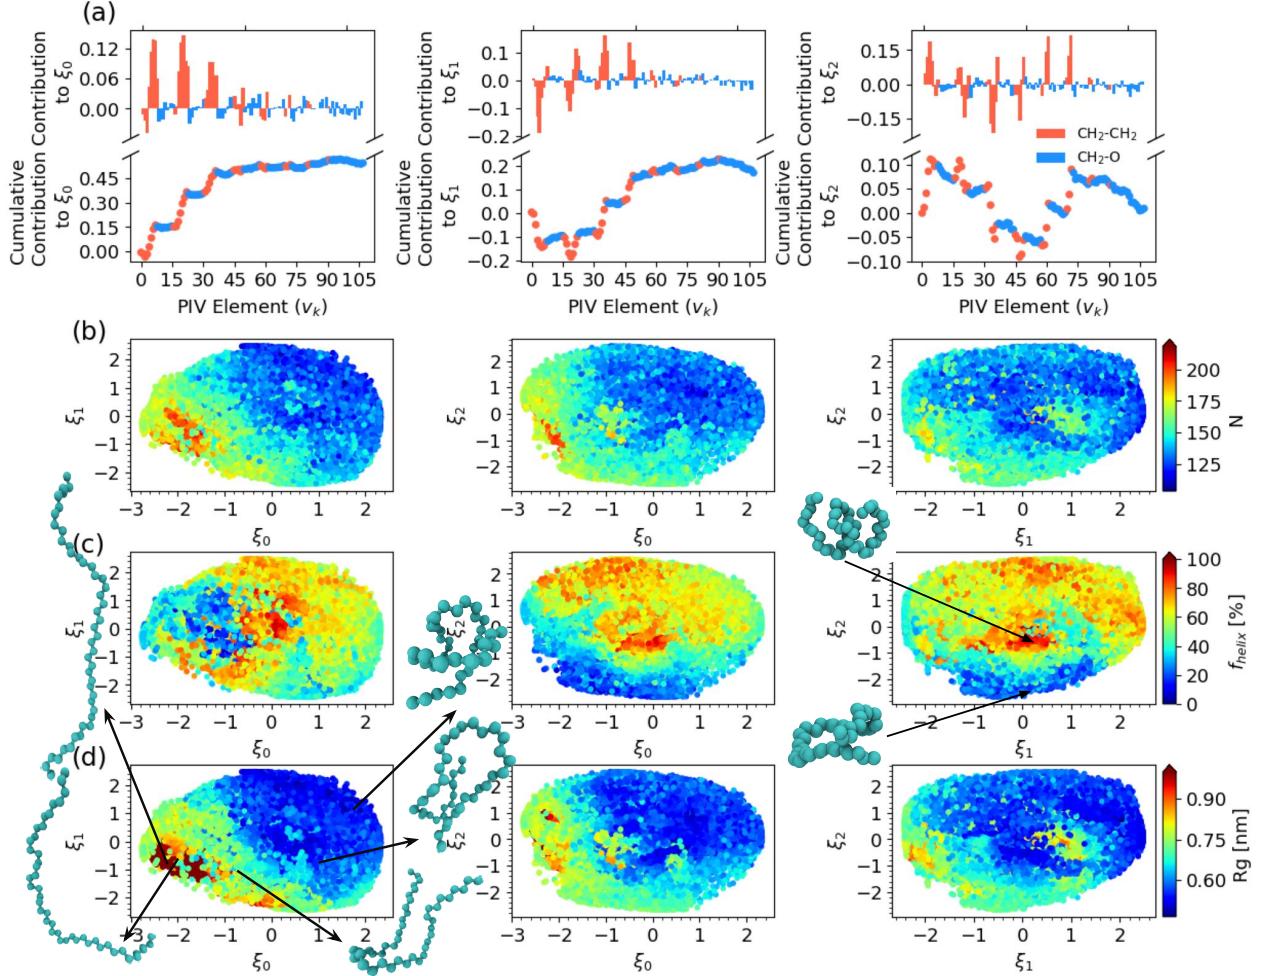


Figure 10: Bottom-up and top-down analysis of terminal learned CVs for the *n*-pentatetracontane chain in water. (a) Split plots show the signed contribution from each PIV element to each CV $\{\xi_0, \xi_1, \xi_2\}$ as a bar chart on top and as a cumulative sum of the signed contributions on the bottom. The magnitude of contributions on the bar chart emphasize the importance of individual PIV elements while the slopes in different regions of the scatter plot illuminate the concerted effect of atom groups on each CV. (b-d) Projections of the terminal Round 7 sampling data into the terminal CVs and colored by physically interpretable candidate variables. ξ_0 and ξ_1 are correlated with the chain radius of gyration, R_g , and the number of waters, N , in the first solvation shell of the chain. We highlight representative configurations on the ξ_0 - ξ_1 surface corresponding to the extended (red), collapsed (blue), and partially collapsed (green) structures. ξ_2 is correlated with the helical pitch f_{helix} . We highlight representative configurations on the ξ_1 - ξ_2 surface corresponding to planar rings (blue) and corkscrews (red).

$r_{CH13-CH33} + r_{CH23-CH43}$) and min-max normalized, measures deviations from a common helical chain structure with a 20 united atom turn, where $r_{CH3-CH43}$ is the distance between the 3rd and 43rd united atoms in the chain and $(r_{CH3-CH23} + r_{CH13-CH33} + r_{CH23-CH43})$ are the distances between 20 united atom separations in the chain. Large values of f_{helix}

correspond to an extended chain with regular turns of 20 united atoms while low values correspond to wider helices and ring-like structures. It also appears that solute-solvent PIV elements are most strongly represented within ξ_2 . Taken together, the CVs appear to contain signatures of the compactness, global hydration, and helicity of the polymer chain within the water solvent.

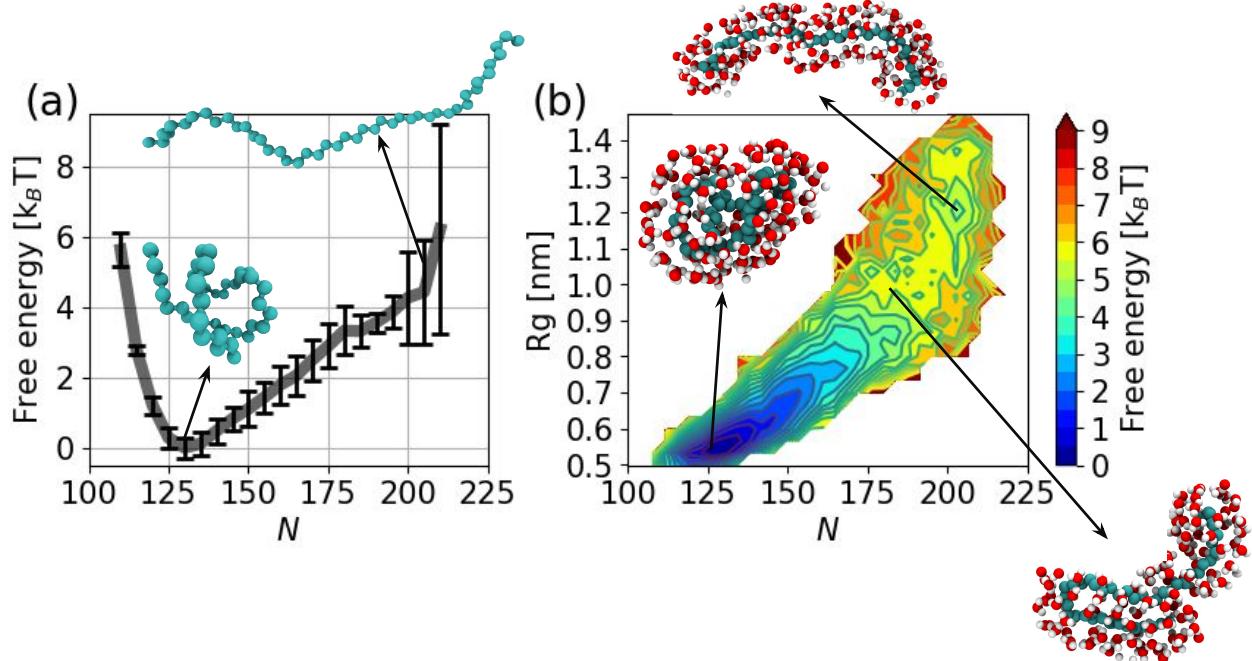


Figure 11: Reweighted free energy surfaces from the terminal round of enhanced sampling for the *n*-pentatetracontane chain in water. (a) A 1D FES projected into the number of waters N in the first solvation shell of the chain. A collapsed state with $N \approx 130$ is more stable than the fully extended state with $N \approx 200$ by approximately $4 k_BT$. Uncertainties in the free energy correspond to standard errors estimated by five-fold block averaging. (b) A 2D FES projected into the number of waters N in the first solvation shell of the chain and the radius of gyration R_g . The global free energy minimum at $(N \approx 130, R_g \approx 0.57 \text{ nm})$ contains the hydrophobically collapsed chain.

In Fig. 11 we show our converged and reweighted FES from the terminal round of PINES projected into physical order parameters. A 1D FES projected into the number of waters N in the first solvation shell of the chain exposes the collapsed state of the chain to be more stable than the extended by approximately $(4 \pm 1) k_BT$, in good agreement with the value of $3.3 k_BT$ reported by Dhabal *et al.*⁶¹ (Fig. 11a). A 2D FES constructed in N and R_g clearly exposes the relationship between hydration and chain extent, with the collapsed state of

the chain better shielding the hydrophobic beads from solvent and occupying the global free energy minimum at this thermodynamic state point (Fig. 11b).

4 Conclusions

In this work we have introduced Permutationally Invariant Network for Enhanced Sampling (PINES) as an approach for data-driven discovery and enhanced sampling in translationally, rotationally, and permutationally invariant collective variables for molecular systems. The approach essentially integrates Permutation Invariant Vector (PIV) featurizations with the Molecular Enhanced Sampling with Autoencoders (MESA) approach to perform interleaved rounds of symmetry-adapted collective variables and parallel bias metadynamics enhanced sampling to perform comprehensive sampling of the thermally accessible phase space and obtain converged thermodynamic averages. Convergence of the iterative sampling process is assessed by stabilization of the learned collective variables and saturation of the explored configurational phase space. PINES is well-suited to data-driven collective variable discovery and enhanced sampling in challenging-to-sample systems of indistinguishable particles frequently encountered in self-assembling and solvated systems. We have demonstrated PINES in applications to the assembly/disassembly of a 13-particle Ar cluster, association/dissociation of a NaCl ion pair in water, and hydrophobic collapse of a *n*-pentatetracontane ($C_{45}H_{92}$) chain in water. In all instances, PINES efficiently discovers symmetry-adapted multi-body collective variables that are used to efficiently drive sampling of configurational space. We anticipate that PINES may be of value in applications to diverse molecular systems comprising indistinguishable particles and we have made the approach freely available as a new user-friendly module at <https://github.com/Ferg-Lab/pines.git> that can be patched with PLUMED2 enhanced sampling libraries and will be submitted for incorporation into a future release of PLUMED2.⁶⁶

In future work, we would like to apply PINES to engage more complex processes such as

protein-folding, DNA self-assembly, and ligand-binding. Additionally, we envision a number of future technical improvements and modifications to PINES. One of the key advantages of PINES is its flexibility through modularity. Any generic CV-based biasing method is immediately compatible with PINES, but with small modifications to the open-source code the featurization choice and the network architecture may also be adapted to employ alternative CV discovery techniques such as, for example, SRVs to discover slow, as opposed to high-variance, CVs. We also note that a PyTorch PLUMED2 module^{115,116} enables extension of the PINES approach to use arbitrarily complicated ANN architectures within the PINES pipeline. In addition, we are also implementing PINES in the flexible GPU accelerated enhanced sampling library PySAGES.¹¹⁷ A principal use case of PINES is its application to solvated systems. The current implementation of PINES employs only solute-solvent distances within the PIV and does not consider the solvent-solvent interaction block. While this has the benefit of reducing the size of the PIV, the inclusion of such interactions would provide additional information on the structure of the near-field solvent that may be important for particular systems such as proteins with highly structured solvation shells. Lastly, we note that the current method makes no use of the decoder portion of the network other than its necessity in training the autoencoder network. The decoder, however, could possibly be utilized to generate novel conformations by sampling from sparse regions in the latent space to seed enhanced sampling calculations and accelerate convergence of the PINES protocol.

Conflict of Interest Disclosure

A.L.F. is a co-founder and consultant of Evozyne, Inc. and a co-author of US Patent Applications 16/887,710 and 17/642,582, US Provisional Patent Applications 62/853,919, 62/900,420, 63/314,898, 63/479,378, and 63/521,617, and International Patent Applications PCT/US2020/035206 and PCT/US2020/050466.

Acknowledgement

Work by N.S.M.H. was supported by the National Science Foundation under Grant No. DMR-1841807. Work by S.D. was supported by MICCoM (Midwest Center for Computational Materials), as part of the Computational Materials Science Program funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division, through Argonne National Laboratory, under Contract No. DE-AC02-06CH11357. This work was completed in part with resources provided by the University of Chicago Research Computing Center. We gratefully acknowledge computing time on the University of Chicago high-performance GPU-based cyberinfrastructure supported by the National Science Foundation under Grant No. DMR-1828629.

Supporting Information Available

Supplementary information comprising switching function parameterization process from pairwise RDFs, FVE plots used to determine the appropriate dimensionality of the latent space, analysis of convergence and converged FES for 13-atom Argon system in an entropically dominated $5 \times 5 \times 5$ nm³ box, analysis of transferability of learned CVs for the 13 atom Argon system to out-of-training temperature regimes, cosine similarity of discovered CVs for NaCl system with various numbers of proximal water considered, projection of the trajectories for various NaCl water systems into the CVs identified for the 35 water system, comparison of the free energy surfaces for the NaCl system obtained by PINES to enhanced sampling in the Na-Cl distance alone, and analysis of transferability of learned CVs from the NaCl system to a divalent ion pair by scaling the Na⁺ and Cl⁻ valency.

References

- (1) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press: New York, NY, 2017.
- (2) Frenkel, D.; Smit, B. *Understanding Molecular Simulation: From Algorithms to Applications*; Elsevier: Amsterdam, Netherlands, 2001.
- (3) Dill, K. A.; Chan, H. S. From Levinthal to pathways to funnels. *Nat. Struct. Biol.* **1997**, *4*, 10–19.
- (4) Hénin, J.; Lelièvre, T.; Shirts, M.; Valsson, O.; Delemonette, L. Enhanced sampling methods for molecular dynamics simulations [Article v1. 0]. *Living J. Comput. Mol. Sci.* **2022**, *4*, 1583–1583.
- (5) Yang, Y. I.; Shao, Q.; Zhang, J.; Yang, L.; Gao, Y. Q. Enhanced sampling in molecular dynamics. *J. Chem. Phys.* **2019**, *151*, 070902.
- (6) Abrams, C.; Bussi, G. Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. *Entropy* **2013**, *16*, 163–199.
- (7) Sidky, H.; Chen, W.; Ferguson, A. L. Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation. *Mol. Phys.* **2020**, *118*, e1737742.
- (8) Chen, W.; Ferguson, A. L. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *J. Comput. Chem.* **2018**, *39*, 2079–2102.
- (9) Torrie, G. M.; Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *J. Comput. Phys.* **1977**, *23*, 187–199.
- (10) Darve, E.; Pohorille, A. Calculating free energies using average force. *J. Chem. Phys.* **2001**, *115*, 9169–9183.

(11) Laio, A.; Parrinello, M. Escaping free-energy minima. *Proc. Natl. Acad. Sci. U. S. A.* **2002**, *99*, 12562–12566.

(12) Zwanzig, R. *Nonequilibrium Statistical Mechanics*; Oxford University Press: New York, NY, 2001.

(13) Wang, J.; Ferguson, A. Nonlinear machine learning in simulations of soft and biological materials. *Mol. Simul.* **2018**, *44*, 1090–1107.

(14) Kingma, D. P.; Welling, M. An introduction to variational autoencoders. *Found. Trends Mach. Learn.* **2019**, *12*, 307–392.

(15) Kingma, D. P.; Welling, M. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114* **2022**,

(16) Chen, W.; Tan, A. R.; Ferguson, A. L. Collective variable discovery and enhanced sampling using autoencoders: Innovations in network architecture and error function design. *J. Chem. Phys.* **2018**, *149*, 072312.

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

(18) Belkacemi, Z.; Gkekka, P.; Lelièvre, T.; Stoltz, G. Chasing Collective Variables using Autoencoders and biased trajectories. *J. Chem. Theory Comput.* **2021**, *18*, 59–78.

(19) Hooft, F.; Pérez de Alba Ortíz, A.; Ensing, B. Discovering collective variables of molecular transitions via genetic algorithms and neural networks. *J. Chem. Theory Comput.* **2021**, *17*, 2294–2306.

(20) Ma, A.; Dinner, A. R. Automatic method for identifying reaction coordinates in complex systems. *J. Phys. Chem. B* **2005**, *109*, 6769–6779.

(21) Ketkaew, R.; Creazzo, F.; Luber, S. Machine learning-assisted discovery of hidden states in expanded free energy space. *J. Phys. Chem. Lett.* **2022**, *13*, 1797–1805.

(22) Ketkaew, R.; Luber, S. DeepCV: A Deep learning framework for blind search of collective variables in expanded configurational space. *J. Chem. Inf. Model.* **2022**, *62*, 6352–6364.

(23) Wehmeyer, C.; Noé, F. Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *J. Chem. Phys.* **2018**, *148*, 241703.

(24) Hernández, C. X.; Wayment-Steele, H. K.; Sultan, M. M.; Husic, B. E.; Pande, V. S. Variational encoding of complex dynamics. *Phys. Rev. E* **2018**, *97*, 062412.

(25) Sultan, M. M.; Wayment-Steele, H. K.; Pande, V. S. Transferable neural networks for enhanced sampling of protein dynamics. *J. Chem. Theory Comput.* **2018**, *14*, 1887–1894.

(26) Wayment-Steele, H. K.; Pande, V. S. Note: Variational encoding of protein dynamics benefits from maximizing latent autocorrelation. *J. Chem. Phys.* **2018**, *149*, 216101.

(27) Wang, D.; Tiwary, P. State predictive information bottleneck. *J. Chem. Phys.* **2021**, *154*, 134111.

(28) Chen, W.; Sidky, H.; Ferguson, A. L. Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets. *J. Chem. Phys.* **2019**, *150*, 214114.

(29) Chen, W.; Sidky, H.; Ferguson, A. L. Capabilities and limitations of time-lagged autoencoders for slow mode discovery in dynamical systems. *J. Chem. Phys.* **2019**, *151*, 064123.

(30) Shmilovich, K.; Ferguson, A. L. Girsanov Reweighting Enhanced Sampling Technique (GREST): On-the-fly data-driven discovery of and enhanced sampling in slow collective variables. *J. Phys. Chem. A* **2023**, *127*, 3497–3517.

(31) Bonati, L.; Piccini, G.; Parrinello, M. Deep learning the slow modes for rare events sampling. *Proc. Natl. Acad. Sci. U.S.A.* **2021**, *118*, e2113533118.

(32) Frassek, M.; Arjun, A.; Bolhuis, P. An extended autoencoder model for reaction coordinate discovery in rare event molecular dynamics datasets. *J. Chem. Phys.* **2021**, *155*, 064103.

(33) Musil, F.; Grisafi, A.; Bartók, A. P.; Ortner, C.; Csányi, G.; Ceriotti, M. Physics-inspired structural representations for molecules and materials. *Chem. Rev.* **2021**, *121*, 9759–9815.

(34) Himanen, L.; Jäger, M. O.; Morooka, E. V.; Canova, F. F.; Ranawat, Y. S.; Gao, D. Z.; Rinke, P.; Foster, A. S. Dscribe: Library of descriptors for machine learning in materials science. *Comput. Phys. Commun.* **2020**, *247*, 106949.

(35) Anderson, B.; Hy, T. S.; Kondor, R. Cormorant: Covariant molecular neural networks. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 14537–14546.

(36) Smidt, T. E. Euclidean symmetry and equivariance in machine learning. *Trends Chem.* **2021**, *3*, 82–85.

(37) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **2022**, *13*, 2453.

(38) Sittel, F.; Jain, A.; Stock, G. Principal component analysis of molecular dynamics: On the use of Cartesian vs. internal coordinates. *J. Chem. Phys.* **2014**, *141*, 07B605_1.

(39) Zhang, L.; Han, J.; Wang, H.; Car, R.; E, W. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.* **2018**, *120*, 143001.

(40) Shorten, C.; Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *J. Big Data* **2019**, *6*, 1–48.

(41) Winter, R.; Noé, F.; Clevert, D.-A. Permutation-invariant variational autoencoder for graph-level representation learning. *Adv. Neural. Inf. Process. Syst.* **2021**, *34*, 9559–9573.

(42) Liu, B.; Xue, M.; Qiu, Y.; Konovalov, K.; O'Connor, M.; Huang, X. GraphVAMPnets for uncovering slow collective variables of self-assembly dynamics. *ChemRxiv preprint 10.26434/chemrxiv-2023-v0dwk* **2023**,

(43) Prudente, F. V.; Acioli, P. H.; Neto, J. S. The fitting of potential energy surfaces using neural networks: Application to the study of vibrational levels of H 3+. *J. Chem. Phys.* **1998**, *109*, 8801–8808.

(44) Nguyen, H. T.; Le, H. M. Modified feed-forward neural network structures and combined-function-derivative approximations incorporating exchange symmetry for potential energy surface fitting. *J. Phys. Chem. A* **2012**, *116*, 4629–4638.

(45) Braams, B. J.; Bowman, J. M. Permutationally invariant potential energy surfaces in high dimensionality. *Int. Rev. Phys. Chem.* **2009**, *28*, 577–606.

(46) Xie, Z.; Bowman, J. M. Permutationally invariant polynomial basis for molecular energy surface fitting via monomial symmetrization. *J. Chem. Theory Comput.* **2010**, *6*, 26–34.

(47) Behler, J.; Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **2007**, *98*, 146401.

(48) Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **2011**, *134*, 074106.

(49) Huo, H.; Rupp, M. Unified representation of molecules and crystals for machine learning. *Mach. Learn.: Sci. Technol.* **2022**, 045017.

(50) Gu, C.; Chang, H.-W.; Maibaum, L.; Pande, V. S.; Carlsson, G. E.; Guibas, L. J. Building Markov state models with solvent dynamics. *BMC Bioinf.* **2013**, *14*, 1–9.

(51) Harrigan, M. P.; Shukla, D.; Pande, V. S. Conserve water: A method for the analysis of solvent in molecular dynamics. *J. Chem. Theory Comput.* **2015**, *11*, 1094–1101.

(52) Raddi, R.; Voelz, V. A Markov state Model of solvent features reveals water dynamics in protein-peptide binding. *ChemRxiv preprint 10.26434/chemrxiv-2023-r6njk* **2023**,

(53) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Phys. Rev. B* **2013**, *87*, 184115.

(54) Hansen, K.; Biegler, F.; Ramakrishnan, R.; Pronobis, W.; Von Lilienfeld, O. A.; Müller, K.-R.; Tkatchenko, A. Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *J. Phys. Chem. Lett.* **2015**, *6*, 2326–2331.

(55) Pietrucci, F.; Andreoni, W. Graph theory meets ab initio molecular dynamics: atomic structures and transformations at the nanoscale. *Phys. Rev. Lett.* **2011**, *107*, 085504.

(56) Pietrucci, F. In *Handbook of Materials Modeling: Methods: Theory and Modeling*; Andreoni, W., Yip, S., Eds.; Springer Nature: Switzerland AG, 2020; Chapter 28: Novel enhanced sampling strategies for transitions between ordered and disordered structures, pp 597–619.

(57) Pipolo, S.; Salanne, M.; Ferlat, G.; Klotz, S.; Saitta, A. M.; Pietrucci, F. Navigating at will on the water phase diagram. *Phys. Rev. Lett.* **2017**, *119*, 245701.

(58) Gallet, G. A.; Pietrucci, F. Structural cluster analysis of chemical reactions in solution. *J. Chem. Phys.* **2013**, *139*, 074101.

(59) Ferguson, A. L.; Panagiotopoulos, A. Z.; Debenedetti, P. G.; Kevrekidis, I. G. Systematic determination of order parameters for chain dynamics using diffusion maps. *Proc. Natl. Acad. Sci. U. S. A.* **2010**, *107*, 13597–13602.

(60) Patel, A. J.; Varilly, P.; Chandler, D.; Garde, S. Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. *J. Stat. Phys.* **2011**, *145*, 265–275.

(61) Dhabal, D.; Jiang, Z.; Pallath, A.; Patel, A. J. Characterizing the interplay between polymer solvation and conformation. *J. Phys. Chem. B* **2021**, *125*, 5434–5442.

(62) Zou, Z.; Beyerle, E. R.; Tsai, S.-T.; Tiwary, P. Driving and characterizing nucleation of urea and glycine polymorphs in water. *Proc. Natl. Acad. Sci. U. S. A.* **2023**, *120*, e2216099120.

(63) Rizzi, V.; Bonati, L.; Ansari, N.; Parrinello, M. The role of water in host-guest interaction. *Nat. Commun.* **2021**, *12*, 1–7.

(64) Pfaendtner, J.; Bonomi, M. Efficient sampling of high-dimensional free-energy landscapes with parallel bias metadynamics. *J. Chem. Theory Comput.* **2015**, *11*, 5062–5067.

(65) Jiang, B.; Guo, H. Permutation invariant polynomial neural network approach to fitting potential energy surfaces. *J. Chem. Phys.* **2013**, *139*, 054112.

(66) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Comput. Phys. Commun.* **2014**, *185*, 604–613.

(67) Rohrdanz, M. A.; Zheng, W.; Clementi, C. Discovering mountain passes via torchlight: Methods for the definition of reaction coordinates and pathways in complex macromolecular reactions. *Annu. Rev. Phys. Chem.* **2013**, *64*, 295–316.

(68) Singer, A. A remark on global positioning from local distances. *Proc. Natl. Acad. Sci. U. S. A.* **2008**, *105*, 9507–9511.

(69) Dokmanic, I.; Parhizkar, R.; Ranieri, J.; Vetterli, M. Euclidean distance matrices: essential theory, algorithms, and applications. *IEEE Signal Process. Mag.* **2015**, *32*, 12–30.

(70) PLUMED2, Switching Functions. <https://www.plumed.org/doc-v2.9/user-doc/html/switchingfunction.html>, (Accessed October 2023).

(71) Humphrey, W.; Dalke, A.; Schulten, K. VMD: Visual molecular dynamics. *J. Mol. Graphics* **1996**, *14*, 33–38.

(72) Zhao, S.; Song, J.; Ermon, S. InfoVAE: Information Maximizing Variational Autoencoders. *arXiv preprint arXiv:1706.02262* **2018**,

(73) Salvador, S.; Chan, P. Determining the number of clusters/segments in hierarchical clustering/segmentation algorithms. 16th IEEE International Conference on Tools with Artificial Intelligence. 2004; pp 576–584.

(74) Hegger, R.; Altis, A.; Nguyen, P. H.; Stock, G. How complex is the dynamics of peptide folding? *Phys. Rev. Lett.* **2007**, *98*, 028102.

(75) Lange, O. F.; Grubmüller, H. Can principal components yield a dimension reduced description of protein dynamics on long time scales? *J. Phys. Chem. B* **2006**, *110*, 22842–22852.

(76) García, A. E. Large-amplitude nonlinear motions in proteins. *Phys. Rev. Lett.* **1992**, *68*, 2696–2699.

(77) Amadei, A.; Linssen, A. B.; Berendsen, H. J. Essential dynamics of proteins. *Proteins: Struct., Funct., Bioinf.* **1993**, *17*, 412–425.

(78) Zhuravlev, P. I.; Materese, C. K.; Papoian, G. A. Deconstructing the native state: energy landscapes, function, and dynamics of globular proteins. *J. Phys. Chem. B* **2009**, *113*, 8800–8812.

(79) Das, P.; Moll, M.; Stamati, H.; Kavraki, L. E.; Clementi, C. Low-dimensional, free-energy landscapes of protein-folding reactions by nonlinear dimensionality reduction. *Proc. Natl. Acad. Sci. U. S. A.* **2006**, *103*, 9885–9890.

(80) Prakash, A.; Fu, C. D.; Bonomi, M.; Pfaendtner, J. Biasing smarter, not harder, by partitioning collective variables into families in parallel bias metadynamics. *J. Chem. Theory Comput.* **2018**, *14*, 4985–4990.

(81) Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. In *Adv. Neural Inf. Process. Syst.*; Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc., 2019; Vol. 32; Chapter PyTorch: An Imperative Style, High-Performance Deep Learning Library.

(82) Falcon, William, PyTorch Lightning. <https://github.com/PyTorchLightning/pytorch-lightning>, (Accessed April 2022).

(83) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* **2017**,

(84) Shlens, J. A Tutorial on Principal Component Analysis. *arXiv preprint arXiv:1404.1100* **2014**,

(85) Barducci, A.; Bussi, G.; Parrinello, M. Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Phys. Rev. Lett.* **2008**, *100*, 020603.

(86) Bussi, G.; Laio, A. Using metadynamics to explore complex free-energy landscapes. *Nat. Rev. Phys.* **2020**, *2*, 200–212.

(87) 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.

(88) 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.; 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. H.; Tiwary, P.; Valsson, O.; Vendruscolo, M.; Voth, G. A.; White, A. Promoting transparency and reproducibility in enhanced molecular simulations. *Nat. Methods* **2019**, *16*, 670–673.

(89) Doye, Jonathan, Lennard-Jones Clusters. <http://doye.chem.ox.ac.uk/research/forest/LJ.html>, (Accessed September 2022).

(90) Dawid, A.; Gburski, Z. Interaction-induced light scattering in Lennard-Jones Argon clusters: Computer simulations. *Phys. Rev. A* **1997**, *56*, 3294–3296.

(91) Briant, C.; Burton, J. Molecular dynamics study of the structure and thermodynamic properties of Argon microclusters. *J. Chem. Phys.* **1975**, *63*, 2045–2058.

(92) Hockney, R.; Goel, S.; Eastwood, J. Quiet high-resolution computer models of a plasma. *J. Comput. Phys.* **1974**, *14*, 148–158.

(93) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **2007**, *126*, 014101.

(94) Jung, H.; Covino, R.; Hummer, G. Artificial intelligence assists discovery of reaction coordinates and mechanisms from molecular dynamics simulations. *arXiv preprint arXiv:1901.04595* **2019**,

(95) Mullen, R. G.; Shea, J.-E.; Peters, B. Transmission coefficients, committors, and solvent coordinates in ion-pair dissociation. *J. Chem. Theory Comput.* **2014**, *10*, 659–667.

(96) Kellermeier, M.; Raiteri, P.; Berg, J. K.; Kempter, A.; Gale, J. D.; Gebauer, D. Entropy drives calcium carbonate ion association. *ChemPhysChem* **2016**, *17*, 3535–3541.

(97) Geissler, P. L.; Dellago, C.; Chandler, D. Kinetic pathways of ion pair dissociation in water. *J. Phys. Chem. B* **1999**, *103*, 3706–3710.

(98) Wang, D.; Zhao, R.; Weeks, J. D.; Tiwary, P. 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.

(99) 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: Struct., Funct., Bioinf.* **2010**, *78*, 1950–1958.

(100) Joung, I. S.; Cheatham III, T. E. Molecular dynamics simulations of the dynamic and energetic properties of alkali and halide ions using water-model-specific ion parameters. *J. Phys. Chem. B* **2009**, *113*, 13279–13290.

(101) Joung, I. S.; Cheatham III, T. E. Determination of alkali and halide monovalent ion parameters for use in explicitly solvated biomolecular simulations. *J. Phys. Chem. B* **2008**, *112*, 9020–9041.

(102) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.* **1983**, *79*, 926–935.

(103) Darden, T.; York, D.; Pedersen, L. Particle mesh Ewald: An $N \log(N)$ method for Ewald sums in large systems. *J. Chem. Phys.* **1993**, *98*, 10089–10092.

(104) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A smooth particle mesh Ewald method. *J. Chem. Phys.* **1995**, *103*, 8577–8593.

(105) Martin, M. G.; Siepmann, J. I. Transferable potentials for phase equilibria. 1. United-atom description of n-alkanes. *J. Phys. Chem. B* **1998**, *102*, 2569–2577.

(106) Eggimann, B. L.; Sunnarborg, A. J.; Stern, H. D.; Bliss, A. P.; Siepmann, J. I. An online parameter and property database for the TraPPE force field. *Mol. Simul.* **2014**, *40*, 101–105.

(107) Berendsen, H.; Grigera, J.; Straatsma, T. The missing term in effective pair potentials. *J. Phys. Chem.* **1987**, *91*, 6269–6271.

(108) Verlet, L. Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules. *Phys. Rev.* **1967**, *159*, 98–103.

(109) Páll, S.; Hess, B. A flexible algorithm for calculating pair interactions on SIMD architectures. *Comput. Phys. Commun.* **2013**, *184*, 2641–2650.

(110) Sidky, H.; Colón, Y. J.; Helfferich, J.; Sikora, B. J.; Bezik, C.; Chu, W.; Giberti, F.; Guo, A. Z.; Jiang, X.; Lequieu, J.; Li, J.; Moller, J.; Quevillon, M. J.;

Rahimi, M.; Ramezani-Dakhel, H.; Rathee, V. S.; Reid, D. R.; Sevgen, E.; Thapar, V.; Webb, M. A.; Whitmer, J. K.; de Pablo, J. J. SSAGES: Software Suite for Advanced General Ensemble Simulations. *J. Chem. Phys.* **2018**, *148*, 044104.

(111) Timko, J.; Bucher, D.; Kuyucak, S. Dissociation of NaCl in water from ab initio molecular dynamics simulations. *J. Chem. Phys.* **2010**, *132*, 114510.

(112) Ballard, A. J.; Dellago, C. Toward the mechanism of ionic dissociation in water. *J. Phys. Chem. B* **2012**, *116*, 13490–13497.

(113) ten Wolde, P. R.; Chandler, D. Drying-induced hydrophobic polymer collapse. *Proc. Natl. Acad. Sci. U.S.A.* **2002**, *99*, 6539–6543.

(114) Athawale, M. V.; Goel, G.; Ghosh, T.; Truskett, T. M.; Garde, S. Effects of length-scales and attractions on the collapse of hydrophobic polymers in water. *Proc. Natl. Acad. Sci. U.S.A.* **2007**, *104*, 733–738.

(115) Bonati, L.; Rizzi, V.; Parrinello, M. Data-driven collective variables for enhanced sampling. *J. Phys. Chem. Lett.* **2020**, *11*, 2998–3004.

(116) Bonati, Luigi, mlcvs: Machine learning collective variables. <https://github.com/luigibonati/mlcvs>, (Accessed April 2023).

(117) Rico, P. F. Z.; Schneider, L.; Perez-Lemus, G.; Alessandri, R.; Dasetty, S.; Menéndez, C. A.; Wu, Y.; Jin, Y.; Nguyen, T.; Parker, J.; Ferguson, A. L.; Whitmer, J. K.; de Pablo, J. J. PySAGES: flexible, advanced sampling methods accelerated with GPUs. *arXiv preprint arXiv:2301.04835* **2023**,

Graphical TOC Entry

