

# Tutorial on Molecular Latent Space Simulators (LSS): Spatially and Temporally Continuous Data-Driven Surrogate Dynamical Models of Molecular Systems

Michael S. Jones,<sup>†</sup> Kirill Shmilovich,<sup>†,‡</sup> and Andrew L. Ferguson<sup>\*,†</sup>

<sup>†</sup> *Pritzker School of Molecular Engineering, The University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois 60637, United States*

<sup>‡</sup>*Present address: Prescient Design, Genentech Research and Early Development (gRED) Computational Science, Genentech, South San Francisco, California 94080, United States*

E-mail: andrewferguson@uchicago.edu

## Abstract

The inherently serial nature and requirement for short integration time steps in the numerical integration of molecular dynamics (MD) calculations places strong limitations on the accessible simulation time scales and statistical uncertainties in sampling slowly relaxing dynamics and rare events. Molecular latent space simulators (LSS) are a data-driven approach to learn a surrogate dynamical model of the molecular system from modest MD training trajectories that can generate synthetic trajectories at a fraction of the computational cost. The training data may comprise single long trajectories or multiple short, discontinuous trajectories collected over, for example, distributed computing resources. Provided the training data provide has sufficient sampling of the relevant thermodynamic states and dynamical transitions to robustly learn the underlying microscopic propagator, an LSS furnishes a global model of the dynamics capable of producing temporally and spatially continuous molecular trajectories. Trained LSS models have produced simulation trajectories at up to six orders of magnitude lower cost than classical MD to enable dense sampling of molecular phase space and large reduction of the statistical errors in structural, thermodynamic, and kinetic observables. The LSS employs three deep learning architectures to solve three independent learning problems over the training data: (i) an encoding of the high-dimensional molecular dynamics into a low-dimensional slow latent space using state-free reversible VAMP-nets (SRVs), (ii) a propagator of the microscopic dynamics within the low-dimensional latent space using mixture density networks (MDNs), and (iii) a generative decoding of the low-dimensional latent coordinates back to the original high-dimensional molecular configuration space using conditional Wasserstein generative adversarial networks (cWGANs) or denoising diffusion probability models (DDPMs). In this software tutorial, we introduce the mathematical and numerical background and theory of LSS and present example applications of a user-friendly Python package software implementation to alanine dipeptide and a 28-residue Beta-Beta-Alpha (BBA) protein within simple Python notebooks.

# 1 Background and Theory

Classical molecular dynamics (MD) simulations model the microscopic dynamical evolution of atomistic and molecular systems by integrating Newton's equations of motion.<sup>1</sup> Modern high-performance computing hardware and efficient parallel software implementations have expanded the length scales accessible to all-atom simulations to tens of trillions of atoms.<sup>2</sup> However, the inherently serial nature of numerical integration together with the femtosecond time steps required to capture the fastest microscopic motions has limited accessible time scales to milliseconds.<sup>3</sup> This limitation, often referred to as the “time scale barrier” or “sampling challenge” in molecular simulation, places strong restrictions on the simulation of rare events and slowly relaxing dynamical processes via direct simulation by unbiased molecular dynamics calculations. The inability to (densely) sample rare but important configurational states and/or dynamical transitions, makes simulation trajectories subject to large statistical uncertainties in structural, thermodynamic, and kinetic observables.<sup>4</sup>

A number of strategies have been developed to engage the sampling challenge. *Coarse graining* sacrifices atomistic resolution in service of computational efficiency by developing simplified higher-level models that integrate out atomistic degrees of freedom, typically by lumping multiple atoms together into coarse-grained beads.<sup>5–7</sup> A variety of highly successful coarse-grained models have been developed for biological<sup>8</sup> and condensed matter<sup>9</sup> systems, and protocols developed to parameterize coarse-grained models from both bottom-up (i.e., fitting to all-atom data) and/or top-down (i.e., fitting to experimental observables) perspectives.<sup>5</sup> Although good strategies exist to ensure thermodynamic consistency between all-atom and coarse-grained models, preservation of dynamical consistency has proven a more difficult challenge such that the kinetics of coarse-grained models may be artificially accelerated in an uncontrolled fashion relative to the all-atom systems.<sup>10</sup> *Enhanced sampling techniques* present an alternative strategy to accelerate sampling via collective variable biasing,<sup>11,12</sup> tempering,<sup>13–15</sup> or path-based techniques.<sup>16–19</sup> Collective variable biasing approaches apply artificial potentials to accelerate barrier crossing and improve sampling. The application of

*a posteriori* analytical corrections can recover unbiased thermodynamic averages, but, except in special cases,<sup>20–24</sup> it is not generally possible to recover unbiased dynamical averages. Tempering approaches combine various thermodynamic ensembles within a single expanded ensemble to achieve improved sampling at the thermodynamic state of interest. In general, this involves exchanges between thermodynamic states that produces short, discontinuous trajectory segments as opposed to a single, long, temporally and spatially continuous trajectory. Path-based techniques employ a variety of approaches to generate transition pathways between predefined reactant and product states, but are generally limited in applicability to single-barrier transitions as opposed to global sampling of the thermally-accessible phase space. *Boltzmann generators* represent a relatively new sampling paradigm wherein a normalizing flow is trained to transform a simple easy-to-sample Gaussian prior distribution into a difficult-to-sample Boltzmann distribution over molecular states.<sup>25,26</sup> Once trained, the model can be used to efficiently sample from the target distribution and compute free energy profiles, but is not designed to generate unbiased dynamical trajectories for the estimation of kinetic or path-based observables. *Markov state models* (MSMs) exploit a separation of time scales to model the long-time system dynamics as probabilistic and memoryless jumps between discrete states within which the dynamical relaxations are fast.<sup>27,28</sup> The rate constants for the dynamical jumps between pairs of states are estimated from molecular simulation trajectories. The inherently localized nature of these jumps between kinetically-linked states imbue MSMs with the very attractive feature that they do not require a single long trajectory against which to fit the transition rates, but may be parameterized by a series of short, discontinuous trajectories that sufficiently densely sample the relevant states and transitions and which may be independently generated on distributed computing resources. MSMs can be viewed as a dynamical coarse-graining of the configurational phase space into a set of discrete dynamical states amenable to a divide-and-conquer parameterization of inter-state transition rates.

Molecular latent space simulators (LSS) were introduced in 2020 as a means to learn a

data-driven surrogate model of the underlying dynamics of a molecular system and generate spatially and temporally continuous unbiased trajectories at a fraction of the cost of standard MD.<sup>29,30</sup> Mathematically, MD may be viewed as an algorithm to propagate the state of a molecular system  $\mathbf{x}_t$  at time  $t$  to a configuration  $\mathbf{x}_{t+\tau}$  at time  $(t + \tau)$  via a set of transition density elements  $\mathbf{x}_{t+\tau} \sim p_\tau(\mathbf{x}_{t+\tau}|\mathbf{x}_t)$ .<sup>31,32</sup> For deterministic dynamics and  $\mathbf{x}$  containing the full-dimensional state of the system (i.e., all particle coordinates and momenta) the transition density element  $p_\tau(\mathbf{x}_{t+\tau}|\mathbf{x}_t)$  is a Dirac delta function lying on the single deterministic configuration  $\mathbf{x}_{t+\tau}$  to which the dynamics evolve at time  $(t + \tau)$ . If the state vector contains a reduced representation of the system (e.g., only the configurational variables) or the dynamical evolution is stochastic (e.g., temperature is maintained by a stochastic thermostat), then the transition density elements become distributions. In either case, the transition density elements are computed on-the-fly by accumulating the forces on all of the particles and numerically integrating Newton's equations of motion. The motivation for the LSS approach is to learn an efficient surrogate model for the microscopic transition density elements that can be evaluated at vastly lower computational cost than in MD.

Attempts to directly learn the transition density elements  $p_\tau(\mathbf{x}_{t+\tau}|\mathbf{x}_t)$  have been reported for small systems and have met moderate success,<sup>33</sup> but the curse of dimensionality makes learning of high-dimensional distributions for large systems computationally intractable. The fundamental premise of the LSS approach is to learn a dynamical encoding  $E : \mathbf{x}_t \rightarrow \psi_t$  of molecular system into a latent space spanned by its leading slow modes, learn a surrogate model  $P : \psi_t \rightarrow \psi_{t+\tau}$  to propagate the dynamics autonomously within this low-dimensional slow subspace, and learn a decoding  $D : \psi_{t+\tau} \rightarrow \mathbf{x}_{t+\tau}$  to generate molecular configurations from the dynamical trajectory generated within the latent space. Crucially, molecular systems generically exhibit a separation of time scales (i.e., a spectral gap) arising from cooperative couplings between the constituent atomic degrees of freedom.<sup>11,34-38</sup> This engenders an emergent low-dimensionality of the long-time dynamical evolution within a slow subspace spanned by the leading maximally autocorrelated dynamical modes  $\psi$  that are ki-

netically decoupled from the fast degrees of freedom residing beyond the spectral gap.<sup>11,39–41</sup> If this slow subspace can be discovered and is sufficiently low-dimensional ( $\lesssim 10$ ), then learning of the transition density elements  $p_\tau(\psi_{t+\tau}|\psi_t)$  becomes a computationally tractable low-dimensional learning problem and decoding back to the molecular space  $\psi_{t+\tau} \rightarrow \mathbf{x}_{t+\tau}$  becomes a well-posed conditional generation problem that requires learning to sample from the annealed distribution of fast degrees of freedom conditioned on the slow variable state  $\mathbf{x}_{t+\tau} \sim p(\mathbf{x}_{t+\tau}|\psi_{t+\tau})$ .

Mathematically, the LSS can be expressed as an alternative pathway from  $\mathbf{x}_t$  to  $\mathbf{x}_{t+\tau}$ ,<sup>29,31,32</sup>

$$\begin{array}{ccc}
 \mathbf{x}_t & \xrightarrow{E} & \psi_t \\
 \text{MD} \downarrow & & \downarrow P \\
 \mathbf{x}_{t+\tau} & \xleftarrow{D} & \psi_{t+\tau}
 \end{array} \tag{1}$$

where the encoder  $E$ , propagator  $P$ , and decoder  $D$  can be learned from training data and, once trained, enable  $\mathbf{x}_{t+\tau}$  to be generated at a fraction of the cost of MD. Framed in this manner, the  $E$ ,  $P$ , and  $D$  learning problems are mathematically independent, permitting the three models can be trained separately but using the same MD training data. Furthermore, deep neural network architectures ideally suited to each of these learning problems have been previously developed and can be modularly deployed to serve as the three constituent components of the LSS.<sup>29,30</sup> A pedagogical introduction to the LSS approach and demonstration of a user-friendly Python package implementing the LSS pipeline (<https://github.com/Ferg-Lab/LSS>) is the subject of the present tutorial.

## 1.1 Strengths and Limitations

LSSs share many similarities with MSMs, and it can be instructive to draw a comparison between these approaches to illuminate similarities and differences and strengths and limitations within the context of the more established and familiar MSM formalism. The

underlying principle of LSSs is similar to MSMs, but whereas MSMs define a discrete partitioning of the configurational phase space and learn the transition rates of a jump process between these states, LSSs learn a slow subspace within which the long-time dynamics evolve and learn a continuous-time effective dynamical model within this space. In both cases, a separation of time scales motivates and enables parameterization of a dynamical model in a set of slow collective variables governing the long time dynamical evolution of the system, and discarding the quickly relaxing fast degrees of freedom that are effectively equilibrated to the slow variables. In MSMs, the fast variables rapidly relax within the discrete coarse-grained states that contain dynamically distinct configurations of the slow collective variables. In LSSs, the dynamical coarse-graining procedure is more analogous to the Born-Oppenheimer approximation wherein the fast electronic degrees of freedom are annealed to the slow nuclear dynamics<sup>42</sup> and embodies the Mori-Zwanzig projection operator formalism in which the effective dynamics of a dynamical system can be written in a subspace of slowly evolving collective variables to which the remaining degrees of freedom couple as noise.<sup>43-46</sup>

Like MSMs, LSSs can also be parameterized by short, discontinuous training trajectories generated by distributed computing, provided that these trajectories sufficiently densely sample the relevant states and transitions in the molecular phase space to enable learning of the underlying microscopic dynamics. Unlike MSMs, LSSs do not induce a discretization of the configurational phase space and can therefore furnish spatially and temporally continuous trajectories in the slow collective variables. Moreover, the LSS is generative in the sense that trajectories in the slow subspace can be decoded or backmapped to synthetic molecular dynamics trajectories using a trained generative model to in-paint the fast degrees of freedom conditioned on the state of the slow collective variables.

Finally, LSSs and MSMs are similar in that they learn the underlying probabilistic and memoryless transition density elements of the microscopic dynamics from the training data. The learned dynamical model can then be used to inexpensively generate temporally continuous – and in the case of LSSs a spatially continuous in the slow collective variables –

synthetic trajectories over the entire learned phase space. Since these trajectories are generated stochastically from the learned transition density elements, the generated trajectories are not simply carbon copies of the training trajectories, but are rather novel trajectories through phase space based on the microscopic transition density elements learned from the data. Novel synthetic trajectories can be generated at a fraction of the cost of MD simulations and may therefore be used to densely sample states and events that may have only appeared a few times within the training set. Subject to the quality of the LSS model learned from the data, inexpensive synthetic trajectories may be used to greatly reduce statistical uncertainties in structural, thermodynamic, and kinetic observables. By the same token, LSSs are, like MSMs, fundamentally data-driven models. As such, they are only as good as their training data – they cannot be expected to accurately parameterize states or transitions that are not well represented within the training data, and, while some modest extrapolation into new regions of phase space can be anticipated, they are not generally expected to prospectively discover novel states or transitions. Further, the surrogate dynamical model learned from the data pertains only to that particular molecular system under those particular thermodynamic conditions. As such, in the absence of physical inductive biases, the trained model is not necessarily expected to be transferable to other molecules or other thermodynamic conditions. Finally, the trained models are, by construction, subject to the same systematic errors that may be present in the molecular force fields used to furnish the training data.

We note that a few other methods sharing similarities with the LSS formalism have recently been proposed.<sup>47–51</sup> Of these, the LSS approach is perhaps most closely related to learning of effective dynamics (LED)<sup>50</sup> which uses an autoencoding neural network to learn a latent space representation and employs a recurrent structure to add long-term memory to the propagator.

## 1.2 Components of the LSS

LSSs employ three deep learning architectures to solve three learning problems over the training data: (i) an encoding of the high-dimensional molecular dynamics into a low-dimensional slow latent space using state-free reversible VAMPnets (SRVs), (ii) a propagator of the microscopic dynamics within the low-dimensional latent space using mixture density networks (MDNs), and (iii) a generative decoding of the low-dimensional latent coordinates back to the original high-dimensional molecular configuration space using conditional Wasserstein generative adversarial networks (cWGANS) or denoising diffusion probability models (DDPMs) (Fig. 1). The three learning problems are independent and can be trained over the same MD training set. Full details of these three components have previously been reported elsewhere and the specific architectures and hyperparameter choices depend on the target molecular system,<sup>29,30</sup> but we provide an overview herein of the mathematical and algorithmic underpinnings.

### 1.2.1 Encoder, $E$ : State-free reversible VAMPnets (SRVs)

State-free reversible VAMPnets (SRVs) were introduced in 2019 as a deep neural network approach to perform data-driven nonlinear unsupervised learning of the slowest evolving (i.e., maximally autocorrelated) collective variables in dynamical systems.<sup>39</sup> The name of the approach stems from its algorithmic kinship with variational approach for Markov processes networks (VAMPnets) introduced by Noé and co-workers<sup>52</sup> and can also be considered as a deep variant of time-lagged independent components analysis (Deep-tICA)<sup>53–56</sup> or deep canonical correlation analysis (DCCA).<sup>57</sup> The mathematical basis of SRVs rests upon the transfer operator  $\mathcal{T}$  as the mathematical object that propagates the probability distribution over the microstates of a dynamical system as a function of time according to the transition density elements  $p_\tau(\mathbf{x}_{t+\tau}|\mathbf{x}_t)$ , where  $\tau$  is the time increment between successive states.<sup>58,59</sup> Molecular systems at equilibrium obey detailed balance. This induces the transfer operator  $\mathcal{T}$  to become self-adjoint with respect to the equilibrium distribution and admit a spectrum of eigen-

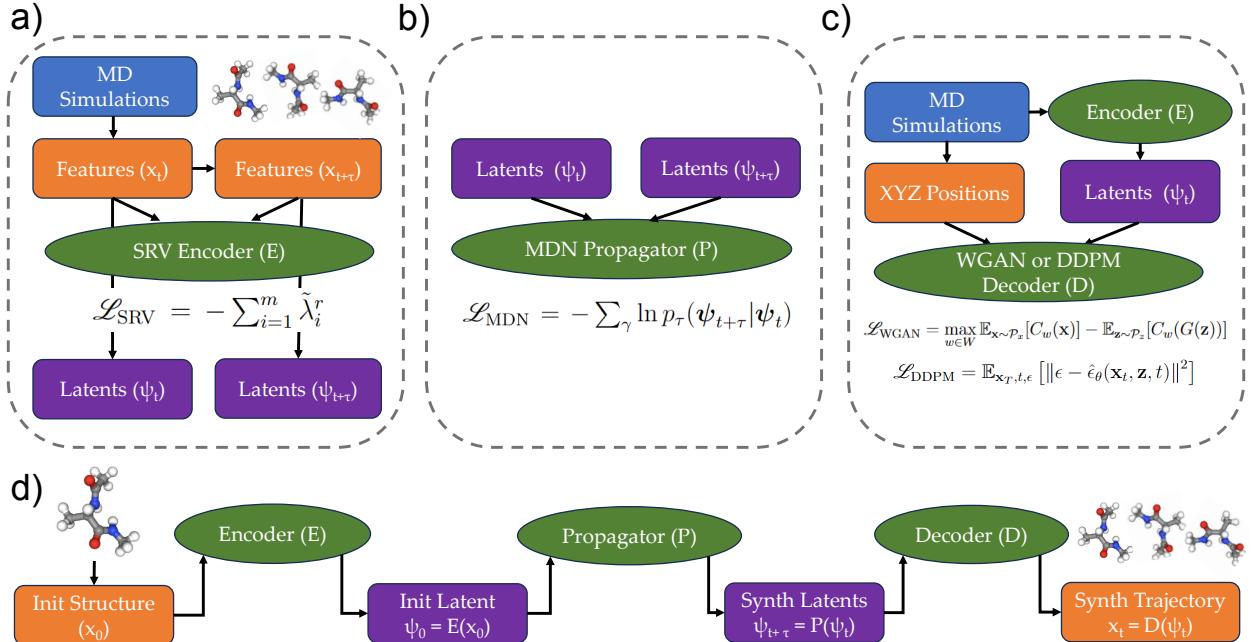


Figure 1: Overview of the molecular latent space simulator (LSS) approach. (a) The SRV encoder  $E$  is trained on molecular dynamics training data to learn a low-dimensional embedding into a latent space spanned by the maximally autocorrelated (i.e., slowest relaxing) dynamical modes  $\psi = E(\mathbf{x})$ . (b) The MDN propagator  $P$  is trained on time-lagged snapshots of the molecular dynamics training data projected into the latent space to learn transition density elements within the latent space that define the latent space propagator  $\psi_{t+\tau} = P(\psi_t) \sim p_\tau(\psi_{t+\tau}|\psi_t)$ . (c) The decoder  $D$  is trained to generate a realization of the molecular configuration conditioned on the latent space coordinates  $\hat{\mathbf{x}}_t = D(\psi_t) \sim p(\hat{\mathbf{x}}_t|\psi_t)$ . The decoding operation produces molecular configurations consistent with the slow degrees of freedom encoded by the latent space coordinate with a realization (i.e., in-painting) of the fast degrees of freedom consistent with the learned distribution of molecular configurations at that latent space embedding. We have implemented decoders based on cWGANS and DDPMs. (d) The final LSS model comprising the trained encoder  $E$ , propagator  $P$ , and decoder  $D$  can be deployed to generate novel synthetic molecular simulation trajectories consistent with the learned microscopic dynamics of the slow modes at approximately six orders of magnitude (i.e., 1 million fold) faster than standard molecular dynamics calculations.

functions  $\{\psi_i(\mathbf{x})\}$  with real eigenvalues  $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \dots$ <sup>29,31,39,40,58,60</sup> The leading eigenfunction corresponds to the equilibrium distribution over microstates and the higher-order eigenfunctions form a natural basis of increasingly more quickly relaxing deviations from the equilibrium distribution with associated relaxation times of  $t_i = -\tau / \ln \lambda_i$ .<sup>31,39,61</sup> For sufficiently long lag times  $\tau$ , the transfer operator becomes effectively low-dimensional as

manifested by a gap in the eigenvalue spectrum separating the leading eigenfunctions spanning a slow subspace from the rapidly relaxing higher order eigenfunctions.<sup>31</sup> This emergent low-dimensionality is what makes the entire LSS approach possible since it becomes numerically tractable to learn a low-dimensional surrogate dynamical model from limited MD trajectory data.

The variational approach to conformational dynamics (VAC) provides a route to optimally approximate the transfer operator eigenfunctions within a finite dimensional basis expansion,  $\tilde{\psi}_i(\mathbf{x}) = \sum_j s_{ij} \zeta_j(\mathbf{x})$ .<sup>39–41</sup> Given a particular choice of basis functions over the microstates  $\{\zeta_j(\mathbf{x})\}$ , the VAC prescribes the optimal expansion coefficients  $\{s_{ij}\}$  to follow from the solution of the generalized eigenvalue equation  $\mathbf{Cs}_i = \tilde{\lambda}_i \mathbf{Qs}_i$ , where  $\mathbf{C}$  is the time-lagged correlation matrix of the basis functions  $\{\zeta_j(\mathbf{x})\}$  under a time lag of  $\tau$ ,  $\mathbf{Q}$  is the non-time-lagged correlation matrix of the basis functions  $\{\zeta_j(\mathbf{x})\}$ ,  $\mathbf{s}_i$  is the (eigen)vector of linear expansion coefficients for the approximate eigenfunction  $\tilde{\psi}_i(\mathbf{x})$ , and  $\tilde{\lambda}_i$  is the corresponding approximate eigenvalue with associated implied relaxation time scale of  $\tilde{t}_i = -\tau / \ln \tilde{\lambda}_i$ .<sup>31,39</sup> In practice,  $\mathbf{C}$  and  $\mathbf{Q}$  are numerically estimated from training data comprising simulation trajectories of the molecular system and we solve the generalized eigenvalue problem by standard techniques. It is useful to observe that the mathematical underpinnings of the VAC are isomorphic to the Roothan-Hall equations in quantum mechanics where one can identify  $\mathbf{C}$  as the Fock matrix and  $\mathbf{Q}$  as the overlap matrix: instead of finding the lowest energy wavefunction of the Hamiltonian operator within a defined basis, we are finding the slowest eigenfunction of the transfer operator within a defined basis.<sup>31,41,42</sup> We also observe that an MSM is a special case of the VAC formalism when the basis functions are selected to be indicator functions over a discrete partitioning of configurational states.<sup>31,58</sup>

The VAC guarantees that the approximated eigenfunctions will be no slower than the true eigenfunctions, and, as a corollary, asserts that better choices of basis functions are those that result in slower leading eigenfunctions.<sup>39–41</sup> The fundamental idea underpinning SRVs is to use deep neural networks as flexible, nonlinear, and data-driven functional ap-

proximators to discover basis functions that result in the slowest possible approximations to the leading transfer operator eigenfunctions. As opposed to generic and pre-selected basis functions (e.g., pairwise distances, atom-centered symmetry functions) the networks learn basis functions tailored to each molecular system that can result in superior approximations to the true leading eigenfunctions of the transfer operator. In this sense, SRVs can simply be viewed as a VAC with a neural network bolted on at the front end whose task is to perform data-driven discovery of bespoke basis functions from MD training data. The networks are trained to learn optimal basis functions leading to the slowest  $m$  eigenfunctions by maximizing their implied time scales via the loss function  $\mathcal{L}_{\text{SRV}} = -\sum_{i=1}^m \tilde{\lambda}_i^r$ , where we typically adopt  $r = 2$  and an appropriate value of  $m$  is selected by resolving a gap in the eigenvalue spectrum (Fig. 1a). The learning task can be made end-to-end differentiable through the VAC. It is typically desirable to employ molecular featurizations or symmetrization operations that ensure the basis functions respect the underlying symmetries of the molecular system (e.g., translational invariance, rotational invariance, permutational invariance).<sup>62</sup> In our applications to date, we have found that simple fully connected feedforward neural networks comprising a handful of layers containing a few hundred neurons per layer have proven adequate for very satisfactory performance.<sup>29,30,39</sup> Conceptually, the success of very simple neural networks can perhaps be understood that they are tasked only with finding good basis functions and that the mathematical heavy lifting is taken care of by the VAC.

For non-equilibrium systems that do not obey detailed balance, the VAC must be replaced by a more general variational principle termed the variational approach to Markov processes (VAMP).<sup>31,52</sup> This results in a more complex mathematical development that rests upon singular vectors and singular values and underpins the more general state-free nonreversible VAMPnets (SNRV) approach.<sup>24</sup> For multi-molecular systems, we must contend with the combinatorial explosion in the state space resulting from the approximately independent nature of the dynamical evolution of each molecule when they are the non-interacting or only weakly interacting regions of configurational space.<sup>63,64</sup> This challenge can be engaged

by extending the LSS paradigm to contain multiple encoders for each independent and interacting subsystem, propagators to evolve the dynamics of each subsystem and the relative locations of each subsystem, and multiple decoders for each independent and interacting subsystem.<sup>30</sup> In this tutorial, we shall restrict our focus to the equilibrium scenario where the VAC applies and to applications to single molecular systems.

A full description of the mathematical underpinnings of VAC and VAMP<sup>31,52</sup> and the numerical implementation of S(N)RVs<sup>29,30,39</sup> are available in prior publications, and an open-source and user-friendly Python package implementing S(N)RVs is freely available from <https://github.com/Ferg-Lab/snrv>.

### 1.2.2 Propagator, $P$ : Mixture density networks (MDNs)

The SRV furnishes a slow subspace spanned by learned approximations to the  $m$  leading transfer operator eigenfunctions  $\{\psi_i(\mathbf{x})\}_{i=1}^m$  prior to the spectral gap. The separation of time scales delimited by the spectral gap means that at sufficiently long lag times  $\tau$ , a surrogate model for the dynamical evolution of the system can be constructed within the slow latent subspace. This assumes that rapid relaxation of the fast degrees of freedom contained in the eigenvectors beyond the spectral gap on sub- $\tau$  time scales, permitting the evolution of the slow variables to be accurately approximated as Markovian (i.e., memoryless) and, at equilibrium, stationary (i.e., time invariant).<sup>43–46</sup> The effective dynamics within the slow-subspace (i.e., approximations to the transition density elements  $p_\tau(\psi_{t+\tau}|\psi_t)$ ) can be learned by projection of the MD training data through the trained SRV model. The low dimensionality of the slow subspace  $k \ll 3N$  is induced by adopting a sufficiently long lag time, wherein only a small number of collective modes contribute to the long-time system evolution. This allows us to break the curse of dimensionality and define a tractable low-dimensional learning problem for a surrogate model of the effective molecular dynamics. Beyond the low dimensionality, the learning problem is also significantly simplified within the transfer operator eigenfunction basis since, by construction, these eigenfunctions diagonalize

the transfer operator and make the dynamical evolution linear within the eigenfunction basis.<sup>31</sup>

Conceptually, the transition density elements in the latent space,  $p_\tau(\psi_{t+\tau}|\psi_t)$ , define the jump probabilities to any other location in the latent space after time  $\tau$  given that the system currently exists at location  $\psi_t$ . We recall that the loss of configurational degrees of freedom under projection and neglect of the velocity information means that these transition density elements are typically distributions over microstates that reflect the stochastic nature of the effective dynamics in the slow subspace. Mixture density networks (MDNs) combine Gaussian mixture models with deep neural networks to efficiently approximate multi-modal distributions,<sup>65,66</sup> and the LSS approach uses MDNs to learn the transition density elements  $p_\tau(\psi_{t+\tau}|\psi_t)$  via linear combinations of  $C$   $m$ -dimensional Gaussian kernels  $\phi_c$ ,

$$p_\tau(\psi_{t+\tau}|\psi_t) = \sum_{c=1}^C \alpha_c(\psi_t) \phi_c(\psi_{t+\tau}; \boldsymbol{\mu}_c(\psi_t), \boldsymbol{\sigma}_c(\psi_t)). \quad (2)$$

During training, the MDN is trained to learn the  $\psi_t$ -dependent means  $\boldsymbol{\mu}_c$  and variances  $\boldsymbol{\sigma}_c$  of the constituent Guassians and  $\psi_t$ -dependent linear mixing coefficients  $\alpha_c$  from the projected MD training trajectories to minimize the loss function  $\mathcal{L}_{\text{MDN}} = -\sum_\gamma \ln p_\tau(\psi_{t+\tau}|\psi_t)$  over time-lagged pairs of training points (Fig. 1b). Once trained, the MDN is then pressed into service as the latent space propagator  $P$ , by iteratively sampling from the learned  $p_\tau(\psi_{t+\tau}|\psi_t)$  distributions to drive the dynamics of the system through the slow latent space in time increments of  $\tau$ .

Importantly, the MDN propagator learns the microscopic transition density elements from the training data as a surrogate model of the effective latent space dynamics and then stochastically samples from them to generate novel trajectories. As such, the trajectories generated by MDN are novel in the sense that they obey the learned microscopic dynamics and are not just carbon copies of the MD training trajectories. On the other hand, the MDN is unlikely to be able to extrapolate far beyond the MD training data so is unlikely

to spontaneously discover novel states or transitions not present in the training ensemble. Sampling from the MDN is extremely computationally efficient. Numerical benchmarks on the systems we have studied to date indicate that propagating the dynamical evolution of the system through the latent space is approximately six orders of magnitude (i.e., 1 million fold) faster than standard molecular dynamics computations, primarily due to the fact that the learned MDN model does not require the expensive force calculations at each time step that are inherently required in standard MD.<sup>29,30</sup> Importantly, propagation of the system dynamics occurs in a time-invariant, autonomous manner completely within the slow latent subspace. The learned MDN surrogate model for the effective dynamics is expressed exclusively in  $\psi$  so there is no requirement to decode the system back up to the molecular space and re-encode back into the slow subspace at each time step in order to propagate the dynamics. This is a valuable attribute of the approach since decoding and encoding during each step of the propagator is computationally slow and can also result in the accumulation of destabilizing errors into the dynamical evolution of the system.<sup>31,67</sup>

In practice, we have found that simple MDN networks comprising two hidden layers of 100 neurons and approximately 50 Gaussian kernels have been sufficient for the systems we have studied to date.<sup>29,30</sup> An open-source and user-friendly Python package implementing MDNs is freely available from [https://github.com/Ferg-Lab/MDN\\_Propagator](https://github.com/Ferg-Lab/MDN_Propagator).

### 1.2.3 Decoder, $D$ : Conditional Wasserstein GANs (cWGANs) and Denoising Diffusion Probabilistic Models (DDPMs)

The trained MDN is used to efficiently generate synthetic trajectories within the slow latent space by sampling from the learned microscopic transition density elements at a fraction of the cost of standard molecular dynamics. The final component of the LSS is a generative model to decode these latent space trajectories back into molecular configuration space. Importantly, this decoding is passive in the sense that there is no requirement that it be done contemporaneously with the MDN trajectory generation or that every frame of the

latent space trajectory must be decoded. By construction, the latent space trajectories describe the time evolution of the leading slow variables. The task of the decoder can be conceived of as to map the slow variables back to molecular configurational space by, in a sense, inverting the operation of the SRV encoder. To do so requires restoration of the fast degrees of freedom omitted in the slow latent space and which were assumed to be in quasi-equilibrium with the slow variables at each time step. Accordingly, we should expect the decoder to be able to generate a smooth, temporally continuous molecular trajectory in the slow collective variables since there is a bijective mapping between the molecular configurational state in these slow variables and each point in the latent space. Conversely, there are multiple realizations of the fast degrees of freedom associated with each point in the latent space, so the decoding operation is one-to-many. The trained decoder should therefore generate molecular configurations drawn from a distribution over the annealed, quasi-equilibrium distribution of fast degrees of freedom consistent with the state of the slow collective variables. In the limit of large training data volumes and perfect training of the decoder over a molecular system at equilibrium, the ensemble of fast degrees of freedom produced by the decoder is expected to approach the Boltzmann distribution.

Mathematically, the decoder is a conditional generative model tasked with learning to sample molecular configurations consistent with each possible state of the slow collective variables  $\mathbf{x}_{t+\tau} \sim p(\mathbf{x}_{t+\tau} | \boldsymbol{\psi}_{t+\tau})$ . Computationally, we have developed two alternative implementations to accomplish this generative decoding: conditional Wasserstein generative adversarial networks (cWGANS) and diffusion denoising probability models (DDPMs). While we have observed both models to achieve excellent reconstruction, we find DDPMs to be typically more stable to train and less sensitive to hyperparameters, whereas cWGANS tend to be more computationally efficient during inference. Given overall better performance in adherence to conditioning variables and reconstruction quality, we strongly recommend employing the DDPM unless decoding speed is of the highest priority.

The cWGAN comprises two components: a generator  $G(\mathbf{z})$  that is tasked to output

realizations of molecular configurations  $\mathbf{x}$  from inputs  $\mathbf{z}$ , and a critic  $C(\mathbf{x})$  that is tasked with evaluating the quality of a molecular configuration  $\mathbf{x}$ . The generator and critic are co-trained in an adversarial manner to minimizer the Wasserstein (i.e., earth mover's) distance,

$$\mathcal{L}_{\text{WGAN}} = \max_{w \in W} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}_x} [C_w(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim \mathcal{P}_z} [C_w(G(\mathbf{z}))], \quad (3)$$

where  $\mathcal{P}_x(\mathbf{x})$  is the distribution over molecular configurations sampled in the MD training data,  $\mathcal{P}_z(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{1}) \in \mathbb{R}^d$  is  $d$ -dimensional Gaussian noise, and  $\{C_w\}_{w \in W}$  is a family of  $K$ -Lipschitz functions enforced through a gradient penalty<sup>68,69</sup> (Fig. 1c). Conditioning is introduced by additionally passing the latent space location  $\psi$  we wish to decode into a molecular configuration to both the generator and critic.<sup>70</sup> As such, the generator is driven by both white noise  $\mathbf{z}$ , which induces diversity into the generated ensemble of fast degrees of freedom, and a conditioning variable  $\psi$ , which informs and restrains the state of the slow degrees of freedom encoded within the latent space. Once trained, the cWGAN critic is discarded and the generator serves as the LSS decoder. To date, we have typically been concerned with biomolecules in isotropic environments and so have trained our cWGAN implementations over rotationally and translationally aligned configurations  $\mathbf{x}$  from the MD training data. Alternatively, one could conceive of training the cWGAN to operate on internal molecular coordinates. For non-isotropic applications (e.g., interactions of molecules with surfaces) it may be desirable to preserve the center-of-mass translation and rotation within the decoding. For weakly coupled two-molecule systems, we trained cWGANS on each system independently oriented to their respective frames of reference.<sup>30</sup> So far, we have restricted our decoder to operate on biomolecular solutes and not tasked it to also decode the coordinates of solvent molecules. To do so would require engaging the permutational invariance of the solvent molecules using brute force combinatoric data augmentation or, as a more scalable and elegant solution, permutationally-invariant descriptors of the solvent environment.<sup>62,71–80</sup>

We have also developed an alternative decoder that leverages a DDPM to conditionally denoise samples from an isotropic Gaussian distribution. The model is based on the pioneering work of Ho et al.<sup>81</sup> and code is adapted from the work of Wang et al.<sup>82</sup> which used a DDPM to predict distributions of backbone dihedrals conditioned on temperature. Training consists of a forward noising process in which the Euclidean positions of MD configurations  $\mathbf{x}_T$  are gradually converted into an isotropic Gaussian  $\mathbf{x}_0$ . During inference, this process is reversed and realistic samples are generated over  $T$  steps by gradually denoising intermediate samples  $\mathbf{x}_t$  via predictions from a neural network which is exposed to the latent space coordinates  $\mathbf{z}$ . The forward diffusion process represents the conditional probability between subsequent steps  $p(\mathbf{x}_i|\mathbf{x}_{i-1})$  and has a Gaussian form represented by  $\mathcal{N}(\sqrt{1-\beta_t}\mathbf{x}_{i-1}, \beta_t\mathbf{I})$  where  $\beta_t$  is defined by a variance schedule and gradually ascends as a function of diffusion time  $t$ . An intermediate sample  $\mathbf{x}_t$  can be computed directly from  $\mathbf{x}_T$  and  $t$  by  $\mathbf{x}_t = \sqrt{\alpha_0}\mathbf{x}_T + \sqrt{1-\alpha_0}\epsilon$  where  $\alpha_t = \prod_{j=1}^{j=t} 1 - \beta_j$  and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ . This direct sampling trick is crucial, as it allows an arbitrary intermediate configuration to be retrieved during training. We use a neural network with a one-dimensional U-net architecture to make a prediction of the noise  $\hat{\epsilon}$  as a function of the intermediate configuration, diffusion time step, and latent space coordinates, and regress this prediction against the actual noise that is deposited (Fig. 1c),

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E}_{\mathbf{x}_T, t, \epsilon} [\|\epsilon - \hat{\epsilon}_\theta(\mathbf{x}_t, \mathbf{z}, t)\|^2]. \quad (4)$$

The reverse diffusion process estimates  $q_\theta(\mathbf{x}_{t+1}|\mathbf{x}_t)$  by drawing from another Gaussian distribution parameterized by  $\mathcal{N}(\mu_\theta(\mathbf{x}_t, \mathbf{z}, t), \sigma_t^2\mathbf{I})$ . The reverse time diffusion process is therefore conditioned on the time step with a learnable mean function  $\mu_\theta$  and an untrained time dependent variance  $\sigma_t^2\mathbf{I} = \beta_t\mathbf{I}$ .<sup>81</sup> Our learned model is essentially tasked with learning the mean function  $\mu_\theta$  of this reverse time denoising diffusion process.<sup>81</sup> Inference proceeds from a given latent space coordinate  $\mathbf{z}$  by sampling random noise  $\mathbf{x}_0$  and making a prediction of

the noise using the trained U-net to remove  $\hat{\epsilon}_\theta(\mathbf{x}_0, \mathbf{z}, t = 0)$ ,

$$\mu_\theta(\mathbf{x}_t, \mathbf{z}, t) = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \hat{\epsilon}_\theta(\mathbf{x}_t, \mathbf{z}, t) \right), \quad (5)$$

$$\sigma_t^2 \mathbf{I} = \beta_t \mathbf{I}, \quad (6)$$

$$\mathbf{x}_{t+1} \sim \mathcal{N}(\mu_\theta(\mathbf{x}_t, \mathbf{z}, t), \sigma_t^2 \mathbf{I}). \quad (7)$$

This process is repeated until  $t = T$ , at which point no additional noise is added to the mean prediction. By default the DDPM decoder uses 1000 diffusion steps during training and inference, however we have demonstrated successful results with a related model<sup>83</sup> using as few as 50 diffusion steps to realize a  $20\times$  acceleration in inference time. The number of diffusion steps acts as a hyperparameter during training and can be tuned based on the size and complexity of the encoded molecular system.

In practice, we have found that even for the largest systems studied to date, relatively simple networks were sufficient to achieve high generative performance: cWGAN networks comprising generators and discriminator networks comprising two hidden layers of 200 neurons each, and DDPM models employing a 1D U-Net architecture with 32 channels and three up/down sampling layers. An open-source and user-friendly Python package implementing the cWGAN and DDPM decoders is freely available from <https://github.com/Ferg-Lab/molgen>.

#### 1.2.4 Deployment

Once all three component parts of the LSS – the SRV encoder, MDN propagator, and cWGAN/DDPM decoder – have been trained, the LSS model can then be deployed for computationally efficient synthetic trajectory generation in three successive steps (Fig. 1d).

**(1)** An initial molecular configuration  $\mathbf{x}_{t=0}$  is deterministically encoded into the latent space via the trained SRV encoder to define the initial coordinates  $\psi_{t=0} = E(\mathbf{x}_{t=0})$  of the latent space trajectory.

**(2)** The MDN propagator then samples from the learned distribution of jump probabilities to define the next state in the latent space trajectory  $\psi_{t=\tau} = P(\psi_{t=0}) \sim p_\tau(\psi_{t+\tau}|\psi_t)$  and sampling is iteratively repeated to construct the succession of states  $[\psi_{t=0}, \psi_{t=\tau}, \psi_{t=2\tau}, \dots]$  defining the latent space trajectory in increments of  $\tau$ . Since the MDN stochastically samples from the learned jump probabilities distributions, the latent space trajectories are not simply copies of the training data, but generate novel trajectories that obey the statistics of the learned microscopic dynamics projected into the latent space and encoded within the trained MDN.

**(3)** The latent space trajectory can be passed to the trained cWGAN or DDPM decoder to generate corresponding molecular configurations consistent with each frame of the latent space trajectory  $\hat{\mathbf{x}}_t = D(\psi_t) \sim p(\hat{\mathbf{x}}_t|\psi_t)$ . Since the MDN trajectory generation proceeds autonomously within the latent space (i.e., successive frames depend only on the latent space coordinates of the prior frame) decoding is typically conducted after MDN trajectory generation has completed. There is no requirement to decode every frame but, since the decoding operations are independent, decoding is an embarrassingly parallel computational task. The decoding operation can be conceived of as in-painting a realization of the equilibrated fast degrees of freedom consistent with the state of the slow degrees of freedom encoded into the latent space coordinate  $\psi_t$ . As such, the decoded trajectory  $[\hat{\mathbf{x}}_{t=0}, \hat{\mathbf{x}}_{t=\tau}, \hat{\mathbf{x}}_{t=2\tau}, \dots]$  is expected to be temporally and structurally continuous in the slow degrees of freedom preserved by the latent space, but may be discontinuous in the fast degrees of freedom drawn from the learned distribution of molecular configurations consistent with each latent space coordinate  $p(\hat{\mathbf{x}}_t|\psi_t)$ . Accordingly, it may be desirable to generate multiple decoded realizations associated with each step of the latent space trajectory to produce an ensemble of  $K$  molecular configurations at each time step  $[\{\hat{\mathbf{x}}_{t=0}^{(i)}\}_{i=1\dots K}, \{\hat{\mathbf{x}}_{t=\tau}^{(i)}\}_{i=1\dots K}, \{\hat{\mathbf{x}}_{t=2\tau}^{(i)}\}_{i=1\dots K}, \dots]$ . It is also possible to select from the ensemble of configurations at each time step that which is most structurally consistent with that produced from the prior time step in order to generate a synthetic molecular trajectory that is temporally coherent in both the slow and fast degrees

of freedom.<sup>84</sup>

### 1.2.5 Uncertainty Quantification, Adaptive Retraining, and Model Transferability

The LSS is a fundamentally data-driven model that is only as good as the training data it is provided. As such, the trained LSS model will contain biases associated with systematic errors contained within the training trajectories due to approximations and biases inherent to the force field and the finite nature of the training data.<sup>4</sup> It can be instructive to quantify the epistemic uncertainties in the model by training ensemble of LSS models over different partitions of the training data. One useful way of doing so is to train the SRV encoder and cWGAN/DDPM decoder over the full training data, but an ensemble of MDN propagators over temporally contiguous stratifications of the training data. Analyzing the variability in the learned jump probability distributions across the ensemble of MDN propagators within a consistent latent space embedding can help expose regions of the latent space that are undersampled in the training data. This can also naturally inform an adaptive retraining paradigm, wherein additional molecular dynamics simulations are initialized in the vicinity of undersampled states and transitions within the latent space to improve the predictions of the model in the regions where it carries the most uncertainty.<sup>27</sup> Iterative cycles of model training and adaptive sampling can be conducted to optimally deploy computational collection of molecular dynamics training data and efficiently converge the LSS model.

The trained LSS model is also constrained to learning a latent space embedding and latent space transition probabilities consistent with the molecular dynamics training data for the molecular system and thermodynamic conditions under which it was collected. The LSS model, as presented, contains no physical model or inductive biases that assure its transferability to other thermodynamic state points (e.g., temperatures, pressures, salt concentrations, solvent conditions) or molecules (e.g., protein mutants). We have previously demonstrated transferability of the encoder and decoder across temperatures and salt con-

ditions for DNA oligomer hybridization/dehybridization, requiring only retraining of the propagator.<sup>30</sup> Very recent work by Dobers *et al.* employing a variant of an LSS using a SE(3)-invariant encoder-decoder has demonstrated limited transferability across an ensemble of dipeptides.<sup>85</sup> Realizing generic and scalable transferability to larger molecules will likely require more substantial innovations such as the inclusion of inductive biases and/or conditioning within the LSS paradigm.

## 2 Prerequisites and Installation

The LSS software is made freely available as open source Python packages implementing the SRV encoder, MDN propagator, and cWGAN or DDPM decoders via <https://github.com/Ferg-Lab/LSS>. The three packages are themselves require a small number of common publicly available Python packages to run, which can be easily installed via conda (<https://anaconda.org/>) or pip (<https://www.python.org/>). For the purposes of this tutorial, we have created demonstration exercises in Google Colab notebooks (<https://colab.research.google.com/>) where prerequisite packages are most straightforwardly installed via pip.

Within a Google Colab notebook, we first install the necessary prerequisite packages.

```
1 # install prerequisite packages
2 %pip install numpy scipy pandas scikit-learn jupyter ipywidgets==7.7.2
  widgetsnbextension jupyter_contrib_nbextensions matplotlib MDTraj tqdm
  pytest pyemma deeptime einops torch torchvision pytorch-lightning nglview
```

Next, we enable the widgets required to view molecular structures and trajectories using nglview.

```
1 # enable Jupyter widgets for nglview
2 !jupyter nbextension enable --py --sys-prefix widgetsnbextension
3 !jupyter nbextension enable nglview --py --sys-prefix
```

```
4 !nglview enable
```

Finally, we install the Python packages implementing the SRV encoder, MDN propagator, and cWGAN and DDPM decoders.

```
1 # install package for SRV encoder
2 %pip install git+https://github.com/andrewlferguson/snrv.git
```

```
1 # install package for MDN propagator
2 %pip install git+https://github.com/Ferg-Lab/mdn_propagator.git
```

```
1 # install package for cWGAN and DDPM decoders
2 %pip install git+https://github.com/Ferg-Lab/molgen.git
```

The installation of all prerequisites should require no more than a few minutes for download and installation under a typical high speed internet connection.

## 3 Exercises

### 3.1 Alanine Dipeptide

As a first exercise, we demonstrate the training and deployment of an LSS for the “hydrogen atom of protein folding”, alanine dipeptide. Full instructions and materials necessary to run this tutorial are available at [https://github.com/Ferg-Lab/IMSI\\_LSS](https://github.com/Ferg-Lab/IMSI_LSS). This tutorial was developed for and presented at the workshop “Learning Collective Variables and Coarse Grained Models” held April 22-26, 2024 at the Institute for Mathematical and Statistical Innovation (IMSI) at the University of Chicago.

**(1) Preparing environment.** We first load the various required components from the prerequisite packages installed above. To improve model training and inference speeds, GPU runtime can be optionally enabled in colab.

```

1 # loading required components
2
3 from mdn_propagator.propagator import Propagator
4
5 from molgen.models import DDPM
6
7 from snrv import Snrv
8
9 from snrv.utils import set_random_seed
10
11 import mdtraj as md
12
13 from pathlib import Path
14
15 import torch
16
17 import matplotlib.pyplot as plt
18
19 import numpy as np
20
21 import nglview as nv
22
23 from google.colab import output
24
25 output.enable_custom_widget_manager()

```

**(2) Loading and processing training data.** Next, we upload the alanine dipeptide training trajectory provided within the GitHub. The 250 ns trajectory contains 250,000 frames saved at 1 ps intervals.

```

1 # opening file upload dialog
2
3 # N.B. If file upload fails, try using alternate upload means by clicking on
4
5     file icon in left menu and directly uploading to colab session storage
6
7     or by uploading to and mounting Google Drive
8
9 from google.colab import files
10
11 files.upload()

```

We then proceed to process the trajectory data for LSS training. We use `mdtraj` to center the trajectory to the origin for visualization convenience and then proceed to extract the pairwise distances between all atoms in the molecule as a translationally and rotationally

invariant featurization of the molecular configuration defining the basis functions  $\{\zeta_j(\mathbf{x})\}$  that we pass to the SRV and from which it constructs approximations to the eigenfunctions of the transfer operator  $\tilde{\psi}_i(\mathbf{x}) = \sum_j s_{ij} \zeta_j(\mathbf{x})$ . A snapshot of the visualization is presented in Fig. 2.

```

1 # processing trajectory
2 trj_fnames = sorted([str(i) for i in Path('.').glob('alanine-dipeptide
3             -*-250ns-nowater.xtc')])
4 top_fname = 'alanine-dipeptide-nowater.pdb'
5 trjs = [md.load(t, top=top_fname).center_coordinates() for t in trj_fnames]
6 trjs
7 # visualizing
8 v = nv.show_mdtraj(trjs[0])
9 v
10 # extracting atomic pairwise distances
11 # N.B. Used commented lines to instead extract pairwise distances between
12     # only backbone atoms
13 coords_torch = list()
14 for trj in trjs:
15     #t_backbone = trj.atom_slice(trj.top.select('backbone')).center_coordinates()
16     #pdists = [torch.pdist(p)[None] for p in torch.tensor(t_backbone.xyz)]
17     pdists = [torch.pdist(p)[None] for p in torch.tensor(trj.xyz)]
18     coords_torch.append(torch.cat(pdists))
19 len(coords_torch), coords_torch[0].shape

```

**(2) Training SRV Encoder.** We now proceed to train the SRV encoder to learn the leading slow modes of the alanine dipeptide conformational dynamics. Training should require no more than a few minutes on a commodity GPU card. We elect to learn three

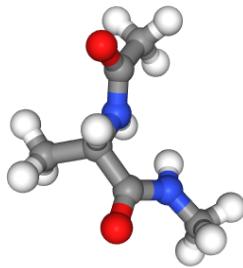


Figure 2: Snapshot of `nglview` visualization of alanine dipeptide training trajectory. Carbon atoms are colored in grey, hydrogen in white, oxygen in red, and nitrogen in blue.

slow modes. The first of these corresponds to the equilibrium distribution with a formally infinite implied time scale, which we discard. The next two correspond to the two slowest relaxing collective modes defining the SRV approximations to the dynamical processes with the largest autocorrelation times.

```
1 # setting random seed for reproducibility
2 set_random_seed(42)
3 # defining SRV training parameters
4 input_size = coords_torch[0].size()[1]
5 output_size = 3
6 hidden_depth = 2
7 hidden_size = 100
8 batch_norm = True
9 dropout_rate = 0.0
10 lr = 1E-2
11 weight_decay = 0.0
12 val_frac = 0.05
13 n_epochs = 30
14 batch_size = 25000
15 VAMPdegree = 2
```

```

16 is_reversible = True
17 num_workers = 0
18
19 model_snrv = Snrv(input_size, output_size, hidden_depth=hidden_depth,
20                     hidden_size=hidden_size,
21                     batch_norm=batch_norm, dropout_rate=dropout_rate, lr=lr,
22                     weight_decay=weight_decay,
23                     val_frac=val_frac, n_epochs=n_epochs, batch_size=batch_size,
24                     VAMPdegree=VAMPdegree, is_reversible=is_reversible, num_workers=
25                     num_workers,
26                     activation=torch.nn.GELU(), device=device)
27
28 model_snrv = model_snrv.to(device)
29
30 # defining lag time
31 lag_n = 10
32
33 # fitting model
34
35 model_snrv.fit(coords_torch, lag=lag_n, scheduler=0.9)
36
37 # flipping trained model from training to evaluation mode
38
39 model_snrv.eval()

```

Optionally, we can create visualizations of the training and validation loss curves and a bar plot of the learned leading implied time scales.

```

1 # plotting loss curves
2
3 fig, ax = plt.subplots()
4
5 ax.plot(np.arange(len(model_snrv.training_losses)), model_snrv.
6         training_losses)
7
8 ax.plot(np.arange(len(model_snrv.validation_losses)), model_snrv.
9         validation_losses)

```

```

5 ax.set_xlabel('epoch')
6 ax.set_ylabel('loss')
7 ax.legend(['training', 'validation'])
8 fig.tight_layout()
9 # plotting implied time scales
10 save_freq = 1 # ps
11 evals = model_snrv.evals.cpu().detach().numpy()
12 plt.bar(range(1,evals.size), -lag_n*save_freq/np.log(evals[1:]))
13 plt.ylabel('Implied timescale (ps)')
14 plt.xticks(range(1,evals.size))
15 plt.xlabel('Timescale index')
16 plt.axhline(lag_n*save_freq, color='k', linestyle=':')

```

It is well known that the conformational dynamics of alanine dipeptide are well represented within a Ramachandran plot of the backbone  $\phi$  and  $\psi$  dihedral angles. An intuition for the learned SRV slow modes can be gained by constructing the Ramachandran plot of the training data and generating a heat map of the slow modes.<sup>39</sup> We note that alanine dipeptide is a particularly simple and well understood system and physical interpretability can be more challenging to develop for systems for which good structural order parameters are not known *a priori*. The leading slow modes of the SRV spanning the LSS latent space are themselves, by construction, good order parameters for tracking the long time dynamical evolution of the system, but can be challenging to physically interpret as they are typically complicated nonlinear functions of the molecular featurization. We can optionally construct this visualization directly within the notebook, a copy of which is presented in Fig. 3.

```

1 # computing latent space coordinates of each training data frame
2 evecs = model_snrv.transform(torch.cat(coords_torch)).cpu().detach().numpy()
3 # plotting Ramachandran heatmaps

```

```

4 trj_cat = md.join(trjs)
5 phi = md.compute_phi(trj_cat)[1].flatten()
6 psi = md.compute_psi(trj_cat)[1].flatten()
7 fig, axes = plt.subplots(1, 2, figsize = (15, 7))
8 axes = axes.flatten()
9
10 for e in range(1, evecs.shape[1]):
11     evec = evecs[:, e]
12     ax = axes[e-1]
13
14     im = ax.scatter(phi, psi, c=evec, s=10, cmap='coolwarm')
15     ax.set_xlabel('$\phi$', fontsize=18)
16     ax.set_ylabel('$\psi$', fontsize=18)
17     ax.set_xlim(-np.pi, np.pi)
18     ax.set_ylim(-np.pi, np.pi)
19     cbar = plt.colorbar(im, ax=ax)
20     cbar.set_label(f'EV_{e}', size=18)
21
22 plt.tight_layout()

```

Finally, we extract the embeddings of the training data into the latent space coordinates

$\psi = (\psi_1, \psi_2) = E(\mathbf{x})$  in preparation for training of the MDN propagator and DDPM decoder.

```

1 # extracting projection of training data into latent space
2 CVs = [model_snrv.transform(x).cpu().detach()[:, 1:] for x in coords_torch]

```

**(3) Training and Deploying MDN Propagator.** We now train the MDN propagator to learn the latent space transition density elements  $p_\tau(\psi_{t+\tau}|\psi_t)$  from the projection of the training trajectories into the latent space and where we adopt a lag time of  $\tau = 10$  frames

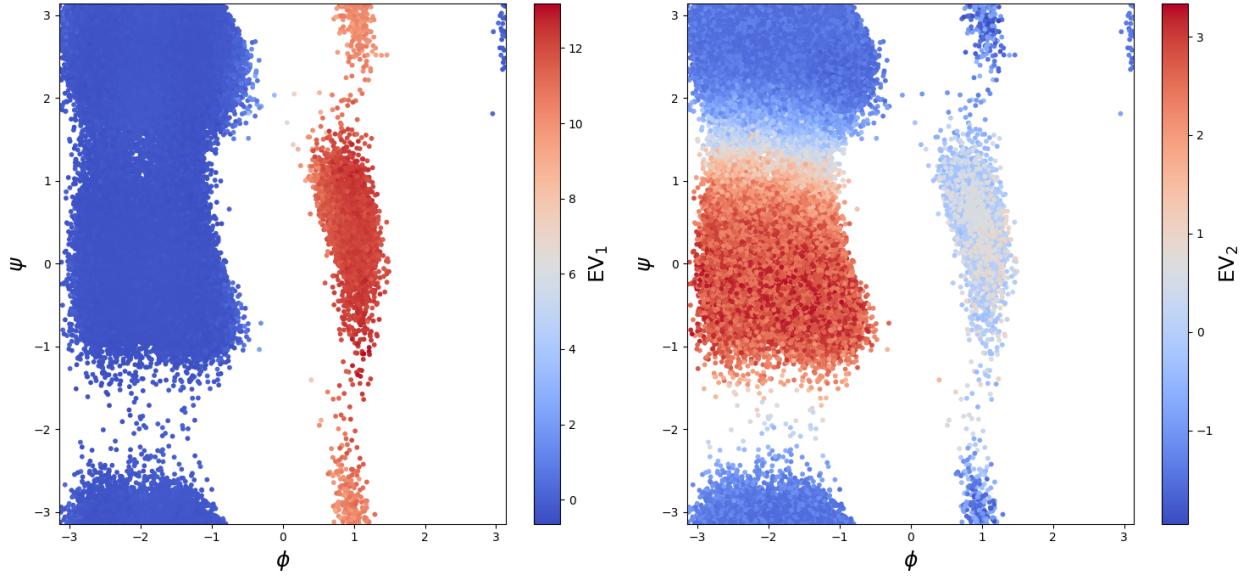


Figure 3: Ramachandran plots of the 250,000 frames constituting the alanine dipeptide training data embedded according to the backbone  $\phi$  and  $\psi$  dihedral angles and colored by the value of the SRV approximations to the first (left) and second (right) eigenfunctions of the transfer operator corresponding to the two leading most slowly relaxing dynamical modes of the molecular system. As expected, the leading mode is highly correlated with  $\phi$  and describes transitions between the triplet well at  $-\pi < \phi < 0$  comprising the  $\beta$ ,  $P_{||}$ , and  $\alpha$  states, from the doublet well at  $\phi \approx 1$  containing the  $\alpha_L$  and  $\gamma$  states, while the second leading mode is strongly correlated with  $\psi$  in the half space  $-\pi < \phi < 0$  and subdivides the triplet well to characterize transitions between the  $(\beta, P_{||})$  and  $\alpha$  states.<sup>39</sup>

$= 10$  ps. Training only takes a few minutes on a commodity GPU.

```

1 # training MDN propagator
2 model_mdn = Propagator(dim = CVs[0].size(1))
3 model_mdn.fit(CVs, lag = 10, max_epochs=10)

```

Once trained, the MDN can then be deployed to generate a synthetic trajectory in the latent space by defining an initial state within the latent space corresponding, arbitrarily, to the embedding of the first frame of the training data  $\psi_{t=0} = E(\mathbf{x}_{t=0})$ , then repeatedly sampling from the learned jump distributions  $\psi_{t=\tau} = P(\psi_{t=0}) \sim p_\tau(\psi_{t+\tau} | \psi_t)$  to generate a succession of 100 states within the latent space at  $\tau = 10$  ps intervals that defines the synthetic latent space simulation trajectory.

```

1 # deploying trained MDN propagator
2 n_steps = int(1E2)
3 x = CVs[0][0][None]
4 synthetic_traj_CVs = model_mdn.gen_synthetic_traj(x, n_steps)

```

**(4) Training and Deploying DDPM Decoder.** Finally, we train a DDPM decoder to learn to generatively decode molecular configurations from latent space coordinates by learning the relationship  $p(\hat{\mathbf{x}}_t|\psi_t)$ . To do so, we train over  $(\mathbf{x}_i, \psi_i = E(\mathbf{x}_i))$  pairs corresponding to translationally and rotationally aligned molecular configurations from the training trajectories and their projected images in the latent space. We note that the training trajectory is already translationally and rotationally aligned and so we do not need to conduct this operation here, but it is easy to do so using the functionality within the `mdtraj` package. Additionally, DDPM training can be quite computationally burdensome, so for the purposes of this tutorial, we reduce the molecular representation to only the backbone atoms to reduce training time. Having done so, the DDPM trains in just a few minutes on a commodity GPU card.

```

1 # preparing training data
2 xyz = list()
3 for trj in trjs:
4     t_backbone = trj.atom_slice(trj.top.select('backbone')).
5         center_coordinates()
6     n = trj.xyz.shape[0]
7     xyz.append(torch.tensor(t_backbone.xyz.reshape(n, -1)).float())
8 # training DDPM decoder
9 model_ddpm = DDPM(xyz[0].shape[1], CVs[0].shape[1])
10 model_ddpm.fit(xyz, CVs, max_epochs=3)

```

Once trained, the DDPM can then be deployed to decode the synthetic latent space tra-

jectory generated by the trained MDN by performing the generative decoding operation  $\hat{\mathbf{x}}_t = D(\psi_t) \sim p(\hat{\mathbf{x}}_t | \psi_t)$ . This only requires a couple of minutes on a commodity GPU card. We then visualize the resulting trajectory using `nglview` and save it to file. A snapshot from the synthetic trajectory is presented in Fig. 4. We note that the quality of the reconstruction can be somewhat improved by training the DDPM for additional epochs, but we choose to train for only three epochs in this tutorial to limit the required training time.

```

1 # deploying trained DDPM decoder
2 xyz_gen = model_ddpm.generate(synthetic_traj_CVs)
3 # visualizing decoded synthetic latent space trajectory
4 xyz_gen = xyz_gen.reshape(xyz_gen.size(0), -1, 3).numpy()
5 fake_trj = md.Trajectory(xyz = xyz_gen, topology=t_backbone.top)
6 v = nv.show_mdtraj(fake_trj)
7 v
8 # saving synthetic molecular trajectory to file
9 fake_trj.save_pdb('ADP_backbone_synthetic_traj.pdb')

```

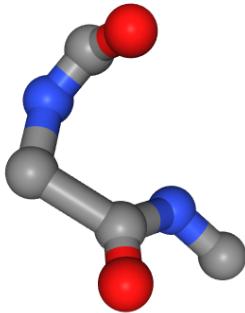


Figure 4: Snapshot of `nglview` visualization of backbone-only reconstruction of alanine dipeptide from the synthetic latent space trajectory produced by the trained MDN propagator. Carbon atoms are colored in grey, oxygen in red, and nitrogen in blue.

## 3.2 BBA

As a second, more sophisticated, exercise, we illustrate the training and deployment of an LSS for the *de novo* designed Beta-Beta-Alpha (BBA) protein (PDB: 1FME). BBA is a 28-residue fast-folding protein that was designed based on a Zinc finger template (PDB: 1AAY) to stabilize a native structure containing two short  $\beta$ -sheets and one  $\alpha$ -helix.<sup>86</sup> A long simulation of this protein was performed by Lindorff-Larsen et al.<sup>3</sup> at D.E. Shaw Research (DESRES) which uncovered numerous (un)folding events that traverse through an ensemble of metastable states. We use 200  $\mu$ s of this simulation data saved every 200 ps (resulting in 1,000,000 frames) to train our LSS. Given the larger size of BBA compared to alanine dipeptide, we reduce our representation of the system to the C $\alpha$  atoms only and use the set of pairwise distances between each of these atoms as our feature set. An even more aggressively parsimonious featurization may retain only every  $n^{\text{th}}$  C $\alpha$  atom. The training process is very similar to that of alanine dipeptide with the main difference being the scale and composition of features and the number of learned SRV coordinates. Based on the substantially larger size of training set and increased training time, we provide a walkthrough that loads pre-trained models for the SRV, MDN, and DDPM decoder, and show some quantitative results below. Full instructions and materials necessary to run this tutorial are available at [https://github.com/Ferg-Lab/LSS\\_BBA](https://github.com/Ferg-Lab/LSS_BBA) including additional code which can be used to train the models from scratch.

**(1) Loading and processing training data.** We have provided the reduced representation of DESRES trajectories in the `collab_files` directory, with frames strided at 10 ns (compared to 200 ps during training), reducing the number of frames from 1,000,000 to 20,000. These files can be downloaded locally and then loaded into the Jupyter or Colab notebook. Alternatively, users can pass in their custom trajectories and topologies by setting `trj_fnames` and `top_fname` to their corresponding paths. Access to complete BBA trajectories is available by request to D.E. Shaw Research. Trajectory data is contained in 10 `.dcd` files and the topology is specified by a single `.pdb` file. To simplify visualization, after data

is loaded, we select the first frame of the first trajectory as our reference configuration and superpose each subsequent frame with respect to this configuration.

```

1 # Path to local or uploaded trajectories
2 load_path = '/PATH/T0/DES'
3 trj_fnames = sorted(glob.glob(f'{load_path}/1FME-0-c-alpha-00*.dcf'))
4
5 # Path to local or uploaded topology
6 top_fname = f'{load_path}/1FME-0-c-alpha.pdb'
7 save_freq = 10 # ns
8
9 trjs = [md.load(t, top=top_fname).center_coordinates() for t in trj_fnames]
10 ref_frame = trjs[0][0]
11 trjs = [t.superpose(ref_frame, 0) for t in trjs]
```

We then use the `torch.pdist` function to calculate all pairwise distance between the 28 C $\alpha$  atoms, resulting in a  $(28 \times 27)/2 = 378$ -dimensional feature set.

```

1 coords_torch = list()
2 for trj in trjs:
3     pdists = [torch.pdist(p)[None] for p in torch.tensor(trj.xyz)]
4     coords_torch.append(torch.cat(pdists))
5 len(coords_torch), coords_torch[0].shape
```

**(2) Loading and Deploying SRV Encoder.** We load our pre-trained SRV, which is stored in the `collab_files` directory. We select a lag of 50 steps (10 ns) based on the approximate convergence of leading timescales (Fig. 5a). We trained our SRV for three epochs based on convergence of the validation loss (Fig. 5b) and project into the six leading non-trivial eigenvectors based on a spectral gap after the sixth mode (Fig. 5c). We recall that we discard the trivial leading eigenvector  $\psi_0$  corresponding to the equilibrium distribution

and possessing a formally infinite implied time scale. In addition to loading the model weights, we load the expansion coefficients and specify that the model has already been fitted. Code for retraining the SRV from scratch is also available in the notebook.

```
1 model_save_path = './model_snrv-3.pth'
2 ckp = torch.load(model_save_path)
3 model_snrv.load_state_dict(ckp['model_state_dict'])
4 model_snrv.evals = ckp['evals']
5 model_snrv.expansion_coefficients = ckp['expansion_coefficients']
6 model_snrv.is_fitted = True
```

We then encode our molecular features into SRV eigenvectors using the `snrv.transform` method. These eigenvectors define the latent space coordinates which will be used later to train both the propagator and decoder. We save another copy `CVs_cat` which concatenates latent space coordinates across all trajectories and converts the data to a `np.ndarray` in order to perform additional analysis.

```
1 CVs = [model_snrv.transform(x).cpu().detach()[:, 1:] for x in coords_torch]
2 CVs_cat = torch.cat(CVs).numpy()
```

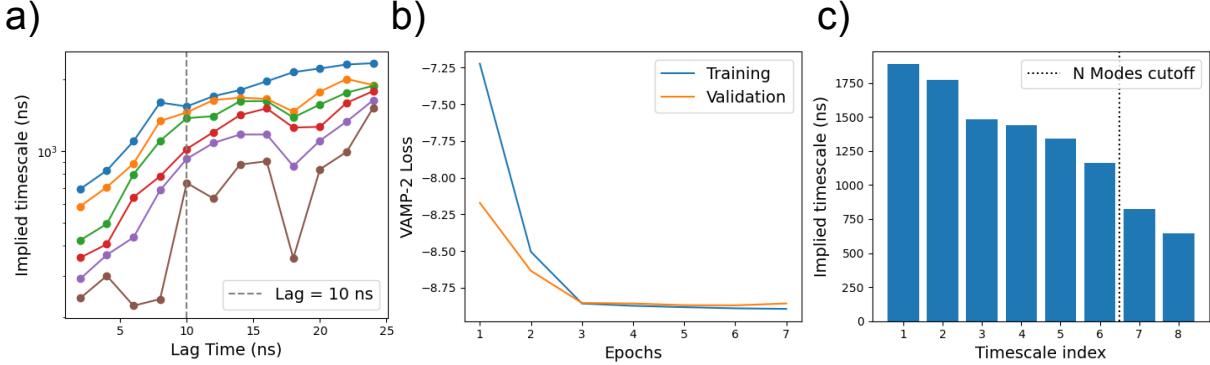


Figure 5: Hyperparameter selection for training SRV encoder applied to BBA trajectories. (a) Implied timescale of leading models based on SRV models trained at various lag times. Each curve corresponds to a different mode, where the blue curve shows the slowest process and brown shows the fastest process. A lag time of 10 ns (dashed line) was selected based on approximate convergence of these modes. (b) Training and validation loss calculated by the sum of leading eigenvalues squared (VAMP-2 loss). Training time of three epochs was selected based on convergence of validation loss. (c) Implied timescales corresponding to the leading eight non-trivial modes. A spectral gap is observed after the sixth mode, and we retrain our model and perform subsequent analysis using only these leading six modes.

Without access to good *a priori* collective variables, such as phi/psi angles for alanine dipeptide, it can be more challenging to interpret our learned SRV coordinates. As an initial check, we first plot comparisons against the leading modes of time-lagged independent component analysis (TICA).<sup>55</sup> TICA solves a linear time-lagged eigenvalue problem (without using a neural network to optimize the basis) and can serve as a reasonable linear approximation for the slowest modes. TICA modes also tend to possess higher variance than SRV modes, which can make them better suited to interpretable visualizations. We obtain the leading TICA coordinates using the `deeptime` Python package<sup>87</sup> as shown below. We specify a lag of one time step, corresponding to a lag time of 10 ns since we loaded the trajectories with a 10 ns stride.

```

1 import deeptime as dt
2
3 tica = dt.decomposition.TICA(lagtime=1, dim=2)
4 TICs = tica.fit_transform([a.numpy() for a in coords_torch]).reshape(-1, 2)

```

In Fig. 6a we show our leading SRV modes overlaid on the first two TIC coordinates. We see strong correlation between  $\psi_1$  and TIC<sub>1</sub>, corresponding to the global (un)folding process. Higher order eigenvectors differentiate smaller regions of TIC space from the greater ensemble, corresponding to transitions in and out of metastable states. Code for generating the SRV heatmaps projected into the leading TIC coordinates is provided below and in Fig. 6a we manually annotate these with representative molecular configurations.

```

1 fig, axes = plt.subplots(2, 3, figsize = (12, 6), sharey=True, sharex=True)
2 axes = axes.flatten()
3 for e in range(1, evecs.shape[1]):
4     evec = evecs[:, e]
5     ax = axes[e-1]
6     im = ax.scatter(TICs[::stride, 0], TICs[::stride, 1], c=evec, s=10, cmap
7         ='coolwarm')
8     ax.set_xlabel('$TIC_1$', fontsize=18)
9     ax.set_ylabel('$TIC_2$', fontsize=18)
10    cbar = plt.colorbar(im, ax=ax)
11    cbar.set_label(f'$\psi_{e}$', size=18)
12 plt.tight_layout()

```

Physical interpretation of the SRV coordinates can be aided by examining their linear correlations with candidate input features. In Fig. 6b we illustrate the Pearson correlation coefficients between the leading three modes and each of the 378 pairwise distance features used to train the encoder. We observe strong positive correlations between  $\psi_1$  and pairwise distances within the group of residues 15-23, which correspond to interactions within the  $\alpha$ -helix. These correlations appear to strongly emerge since the  $\alpha$ -helix tends to be the last secondary structural element that forms during the folding process. We also observe significant negative correlations between  $\psi_2$  and pairwise distances from residues (1-7) to residues (10-19) which correspond to changes in distance between the two  $\beta$ -sheets. Indeed we find

that  $\psi_2$  distinguishes partially unfolded states that contain  $\beta$ -sheet structure from those that do not. The  $\psi_3$  coordinate correlates with transitions into a more rare metastable state in which the protein is effectively folded in half and the termini are in close proximity. The strong negative correlations of  $\psi_3$  with pairwise distances from residues (1-12) ascending to residues (17-28) descending correlate with the relative proximity of these residue pairs extending from the termini (residues 1 and 28) down to the hinge at residues 14-16. Code to generate these correlation plots is provided below.

```

1 fig, axes = plt.subplots(1, 3, figsize=(14, 4))
2 all_feats = np.concatenate(coords_torch)
3 num_features = all_feats.shape[1]
4 nres = trjs[0].n_residues
5
6 for evec_idx, ax in enumerate(axes):
7     heatmap = np.zeros((nres, nres))
8     correlations = np.zeros(num_features)
9     cnt = 0
10    for i in range(num_features):
11        correlations[i], _ = pearsonr(all_feats[:, i], evecs[:, evec_idx +
12            1])
12    for i in range(len(heatmap)):
13        for j in range(i+1, nres):
14            heatmap[i, j] = correlations[cnt]
15            heatmap[j, i] = correlations[cnt]
16        cnt += 1
17    mask = np.triu(np.ones_like(heatmap, dtype=bool))
18    sns.heatmap(heatmap, mask=mask, cmap='coolwarm', ax=ax)
19    ax.set_title(f'$\psi_{evec_idx+1}$ Correlations', fontsize=18)

```

```

20     ax.set_xlabel('Residue index', fontsize=14)
21
22     ax.set_xticklabels(1 + 2*np.arange(14))
23
24     ax.set_yticklabels(1 + 2*np.arange(14))
25
26 axes[0].set_ylabel('Residue index', fontsize=14)

```

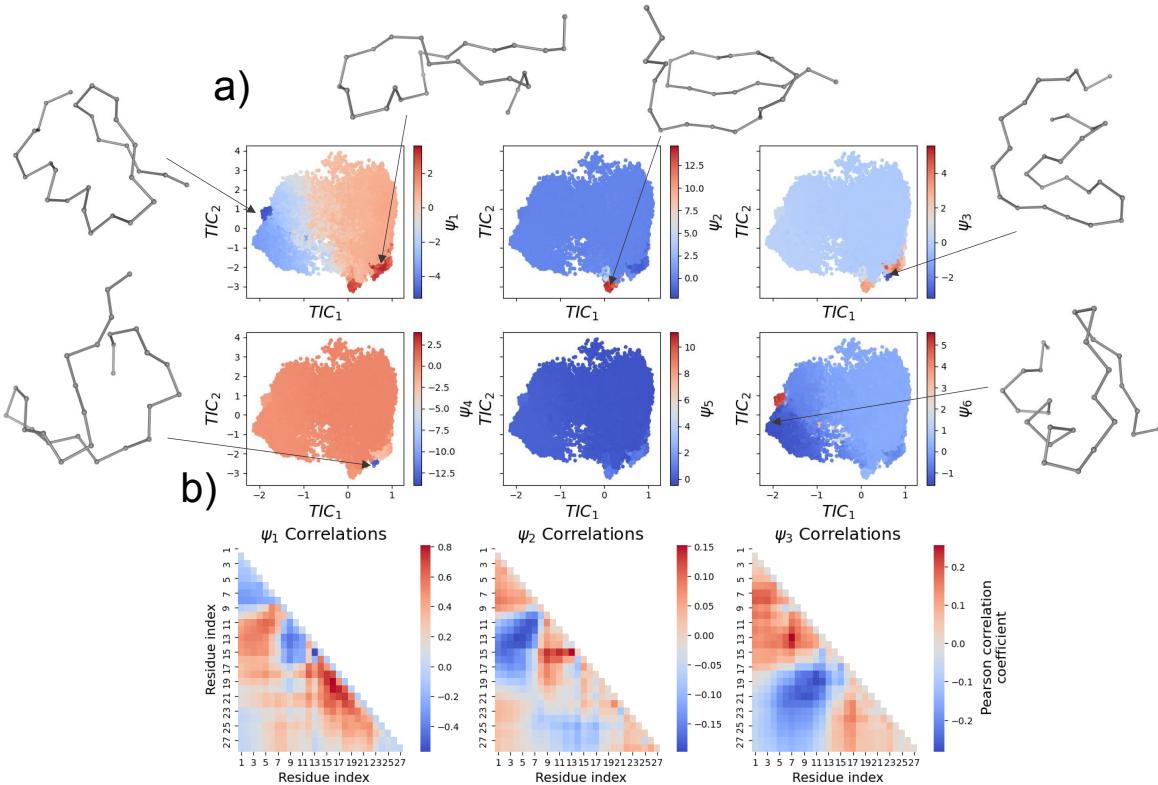


Figure 6: Projection of the 20,000 frames of the BBA training data into the leading two TICA modes. (a) Heat maps coloring the points by the values of the six leading SRV modes  $\{\psi_i\}_{i=1}^6$  onto which we have also projected representative molecular configurations. Changes in the value of the SRV modes expose interconversions between different configurational populations.  $\psi_1$ , for example, is strongly correlated with  $TIC_1$  and corresponds to global (un)folding. Noting the striking, but purely serendipitous, similarity of our TICA embedding to a map of Australia,  $\psi_2$  corresponds to transitions between the red partially unfolded  $\beta$ -sheet state located in Melbourne to the rest of the island, whereas  $\psi_3$  corresponds to transitions between the blue “folded-in-half” state in Canberra to the more populated red unfolded states in Melbourne and Sydney. (b) Pearson correlation coefficients between each of the 378 pairwise distance features and three leading SRV modes. Analysis of these linear correlations can help expose interpretability of the learned SRV modes.

**(3) Loading and Deploying MDN Propagator.** Next, we deploy our MDN prop-

agator to generate synthetic trajectories in our learned latent space. We load our pre-trained model which was run for 100 epochs and predicts conditional probabilities in the six-dimensional latent space. We trained the MDN using the same lag time as the SRV encoder (10 ns). It is not a requirement that the same lag time be used for training the SRV and the MDN, however, using a shorter MDN lag time risks breaking the Markovian (i.e., memoryless) assumption and longer lag times unnecessarily reduce the temporal resolution of the trajectory generated by the propagator.

```

1 model_mdn = Propagator(dim = CVs[0].size(1))
2 model_save_path = f'{load_path}/model_mdn-100.pth'
3 model_mdn.load_state_dict(torch.load(model_save_path))
4 model_mdn.is_fit = True

```

We use our pre-trained model to generate a synthetic trajectory that is  $10\times$  longer than the training data, corresponding to an effective simulation time of 2 ms. This timescale represents an extremely large computational expense at both all-atom and coarse-grained resolution, but the MDN generates this trajectory in less than two minutes on a single GPU. In order to ensure consistency with our MD trajectory, we calculate the number of integration steps by dividing by the MDN lag time. Our initial configuration  $x$  is set to the latent coordinates of the first frame of our SRV embedding to ensure a physically reliable starting point. Each subsequent frame is sampled from the predicted MDN distribution.

```

1 n_steps = 10 * len(CVs_cat) / lag_n
2 x = CVs[0][0][None]
3 synthetic_traj_CVs = model_mdn.gen_synthetic_traj(x, n_steps)

```

To inspect the consistency of our model with the training data, we plot in Fig. 7a the first 10% of our synthetic trajectory as a function of time (orange) and compare it to the SRV embedding of MD data (blue). We show each of the six leading modes on a separate axis, and observe similar dynamical trends between the corresponding modes of the two

trajectories. As described above,  $\psi_1$  corresponds to the global (un)folding process, and we observe several transitions occurring at distinct time points for each trajectory. Higher order modes mainly correspond to rare excursions through metastable states, and we see the abruptness of transitions and relatively short lifetime of these states recapitulated in the synthetic trajectory. Because synthetic trajectories are substantially cheaper to generate than standard molecular dynamics, we can sample many more transitions than are available in the training data and use these to greatly reduce the statistical uncertainty in structural, thermodynamic, and kinetic observables. Code to generate these trajectory comparisons is provided below.

```

1 fig, axes = plt.subplots(6, figsize=(8, 8), sharex=True)
2 x = 0.01*np.arange(len(evecs)//lag_n) # convert to ms
3 for evec_idx, ax in enumerate(axes):
4     ax.plot(x, CVs_cat[:, evec_idx], label='DES')
5     ax.plot(x, synthetic_traj_CVs[:len(x), evec_idx], label='MDN')
6     ax.set_ylabel(f'$\psi_{evec_idx+1}$', fontsize=16)
7 axes[-1].legend(fontsize=14)
8 axes[-1].set_xlabel('Simulation Time (ms)', fontsize=16)

```

Additionally, we may be interested a thermodynamic comparison between the MD trajectories and the synthetic LSS trajectories generated by the MDN. In Fig. 7b we plot a kernel density estimate of the probability distributions along the six leading SRV modes from each trajectory. In this plot, we include data from the complete synthetic trajectory, which effectively represents a  $10\times$  longer timescale than the MD data. As such we expect strong similarities between the synthetic and MD distributions but also expect differences that reflect both the stochasticity in the generated trajectory and the disparate data volumes. Indeed, there is a close correspondence overall, but, for example, we observe that the MDN predicts a slightly higher population in  $\psi_1 < -4$  region and slightly lower population near  $\psi_2 = 0$  than is seen in the MD data. Code to plot these distributions is provided below.

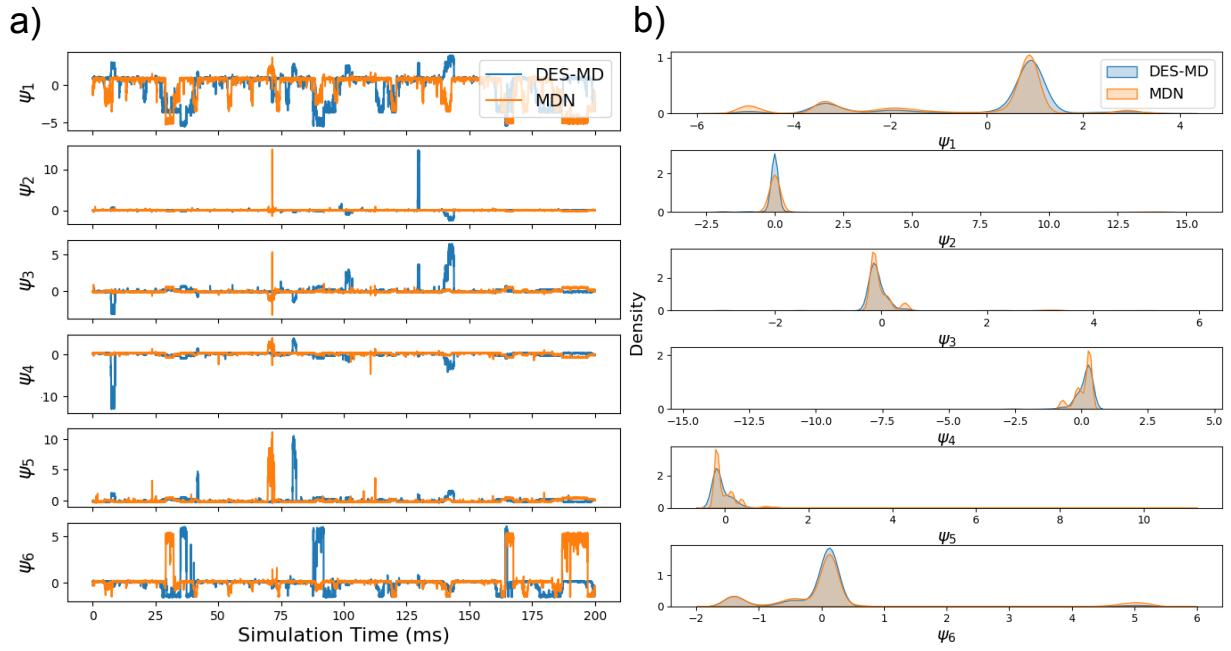


Figure 7: Comparison of synthetic LSS MDN trajectory dynamics and thermodynamics with ground-truth D.E. Shaw molecular dynamics (DES-MD) training data for BBA. (a) DES-MD trajectory data (blue) projected into the leading six SRV modes as a function of time. The LSS MDN trajectory (orange) initialized from the same starting coordinates at the MD data and propagated for an equivalent timescale, where each step is separated by a lag of 10 ns corresponding to a total simulation time of 200  $\mu$ s. (b) Probability distributions of each leading mode for the MD trajectory (blue) and the synthetic LSS MDN trajectory (orange) constructed by kernel density estimation.

```

1 fig, axes = plt.subplots(6, figsize=(8, 8), sharex=False)
2
3 evecs = np.array(evecs)
4 synthetic_traj_CVs = np.array(synthetic_traj_CVs)
5 evecs[np.isinf(evecs)] = np.nan
6 synthetic_traj_CVs[np.isinf(synthetic_traj_CVs)] = np.nan
7
8 for evec_idx, ax in enumerate(axes):
9     sns.kdeplot(evecs[::lag_n, evec_idx+1], ax=ax, label='DES', fill=True)
10    sns.kdeplot(synthetic_traj_CVs[:, evec_idx], ax=ax, label='MDN', fill=

```

```

    True)

10  ax.set_ylabel('')

11  ax.set_xlabel(f'$\psi_{\text{evec\_idx+1}}$', fontsize=16)

12

13 axes[0].legend(fontsize=14)

14 axes[3].set_ylabel(' *20 + 'Density', fontsize=16)

15 plt.tight_layout()

```

**(4) Loading and deploying DDPM Decoder.** Finally, we load our pre-trained DDPM decoder in order to restore the  $C\alpha$  resolution of the original MD training data. This model was trained for 40 epochs and uses 100 diffusion steps. The first frame was selected as a reference configuration to align all subsequent frames. We tested several reference configurations and did not find a difference in DDPM performance in this case, but, in general, we recommend selecting a frame close to the native state or an otherwise highly populated state. To ensure the model was not overfit, we performed a 90/10 split on segments of each of our 10 training trajectories. All subsequent analysis will be shown on the 10% hold-out test data.

```

1 test_perc = 0.1

2 n_test = int(trjs[0].xyz.shape[0]*test_perc)

3 n_train = (trjs[0].xyz.shape[0] - n_test)

4

5 xyz_train = [torch.tensor(trj.xyz[n_test:]).reshape(n_train, -1)).float() for
    trj in trjs]

6 CVs_train = [cv[n_test:] for cv in CVs]

7 xyz_test = [torch.tensor(trj.xyz[:n_test]).reshape(n_test, -1)).float() for
    trj in trjs]

8 CVs_test = [cv[:n_test] for cv in CVs]

```

```

9
10 # Training a cWGAN offers faster decoding but lower structural quality
11 # model = WGANGP(xyz[0].shape[1], CVs_train[0].shape[1])
12
13 n_timesteps = 100
14 model_ddpm = DDPM(xyz_train[0].shape[1], CVs_train[0].shape[1], timesteps=
15     n_timesteps)
16
17 model_save_path = f'{load_path}/model_ddpm-40-eps-{n_timesteps}-ts.pth'
18 model_ddpm.load_state_dict(torch.load(model_save_path))
model_ddpm.is_fit = True

```

To evaluate the quality of our decoder we can first apply the model to SRV embeddings of our MD data and compare against the corresponding reference structure. In the following code, we aggregate the trajectory frames corresponding to our test set and generate synthetic trajectories conditioned on each SRV embedding.

```

1 trj_ref = md.join([trj[:n_test] for trj in trjs])
2 CVs_ref = torch.Tensor(np.concatenate([cv.numpy() for cv in CVs_test]))
3
4 xyz_gen = model_ddpm.generate(torch.Tensor(CVs_ref))
5 xyz_gen = xyz_gen.reshape(xyz_gen.size(0), -1, 3).numpy()
6 fake_trj = md.Trajectory(xyz = xyz_gen, topology=trjs[0].top)
7 fake_trj = md.join([ft.superpose(rf) for ft, rf in zip(fake_trj, trj_ref)])

```

To compare generated structures against their corresponding reference configurations, in Fig. 8a we visualize three structures from distinct regions of SRV latent space to compare the generated structure (red) superposed on the reference structure (blue). We do not expect these structures to be identical given that information about fast degrees of freedom

is lost during the SRV encoding and the model is based on a denoising diffusion process that will stochastically produce distinct samples for identical conditionings, but, as expected, we observe structural similarities associated with the slow dynamics governing the (un)folding of the protein. In Fig. 8b, we illustrate synthetic configurations decoded from randomly selected frames of an MDN trajectory. In this case, no corresponding reference configurations exist, but configurations are physically plausible with respect to the MD ensemble.

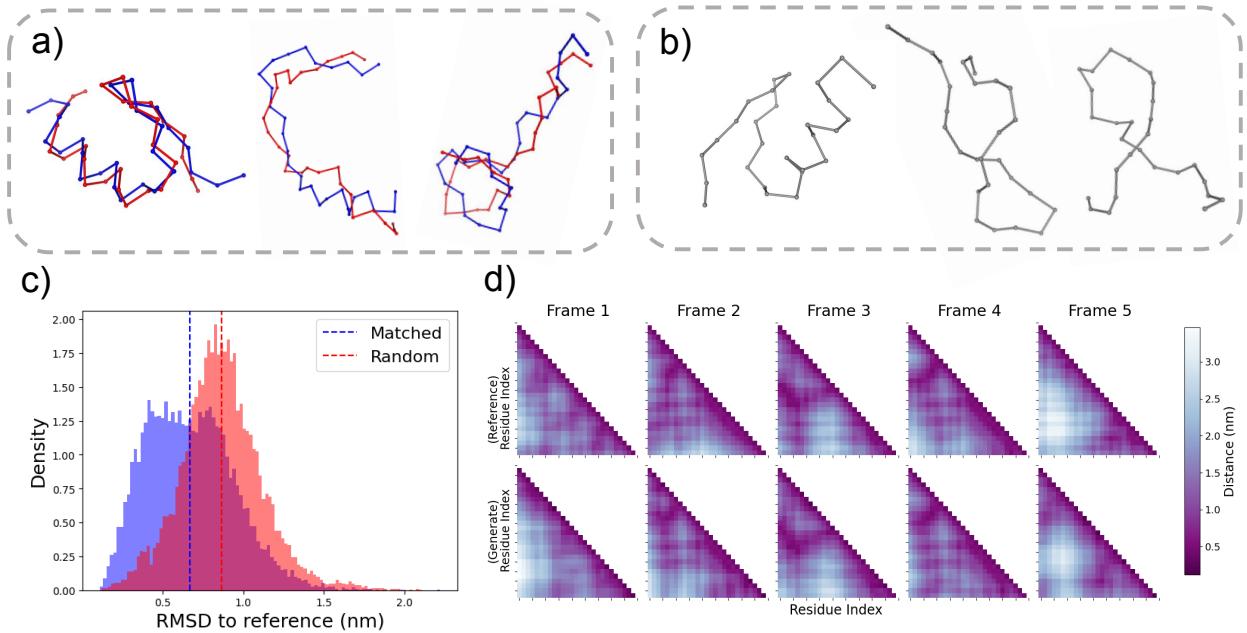


Figure 8: Performance of the DDPM decoder applied to BBA. (a) Synthetic configurations (red) conditionally generated from the SRV latent space coordinates superposed on the corresponding reference configurations (blue). (b) Synthetic configurations generated from randomly selected frames of an MDN trajectory. No corresponding reference configurations exist, but configurations are physically plausible with respect to the MD ensemble. (c) Distribution of RMSDs of all conditionally generated structures with respect to their corresponding (matched) reference structures (blue) and with respect to randomly selected structures (red). The matched structures have on average significantly lower RMSD, indicating that the decoder is paying attention to the conditioning provided by the latent space coordinates. (d) All pairwise distances between C $\alpha$  carbons for five evenly spaced frames within the test trajectory of reference (top row) and generated (bottom row) configurations. Given correct adherence to SRV conditioning, pairwise distance maps of reference and generated structures are expected to approximately match.

To verify that the model is paying attention to the latent space coordinate conditioning, we compared the root mean squared deviation (RSMD) of rototranslationally aligned generated samples with respect to their corresponding matched reference structures versus their alignment to randomly selected frames. We show a comparison of these two distributions in Fig. 8c which, as expected, reveals significantly lower mean RSMD for the matched samples than for the random ones. The code to perform this analysis is provided below.

```

1 rmsds_aligned = []
2 for f_gen, f_ref in zip(fake_trj, trj_ref):
3     rmsds_aligned.append(md.rmsd(f_gen, f_ref)[0])
4
5 rmsds_random = []
6 rand_list = list(range(len(trj_ref)))
7 random.shuffle(rand_list)
8 for f_gen, f_ref in zip(fake_trj, [trj_ref[i] for i in rand_list]):
9     rmsds_random.append(md.rmsd(f_gen, f_ref)[0])
10
11 x = np.arange(len(rmsds_aligned))
12 plt.hist(rmsds_aligned, bins=100, alpha=0.5, color='blue', density=True)
13 plt.hist(rmsds_random, bins=100, alpha=0.5, color='red', density=True)
14
15 plt.axvline(np.mean(rmsds_aligned), linestyle='dashed', c='blue', label='
16     Matched')
17
18 plt.axvline(np.mean(rmsds_random), linestyle='dashed', c='red', label='
19     Random')
20
21 plt.legend(fontsize=16)
22 plt.xlabel('RMSD to reference (nm)', fontsize=18)

```

```
20 plt.ylabel('Density', fontsize=18)
```

As an additional comparison, we visualize all C $\alpha$  carbon pairwise distances between reference and generated configurations. Although not identical, we would expect overall structural trends to remain consistent at samples corresponding to the same regions of SRV space. In Fig. 8d we observe strong similarity between reference (top row) and generated (bottom row) structures for five evenly spaced frames in the test trajectory. Code to generate these plots is provided below.

```
1 pdists_ref = [torch.pdist(p)[None] for p in torch.tensor(trj_ref.xyz
2   [::2000])]
3
4 pdists_gen = [torch.pdist(p)[None] for p in torch.tensor(fake_trj.xyz
5   [::2000])]
6
7
8 fig, axes = plt.subplots(2, 5, figsize=(15, 6))
9 cmap = 'BuPu_r'
10
11 global_min = np.inf
12 global_max = -np.inf
13
14 for p_ref, p_gen in zip(pdists_ref, pdists_gen):
15     global_min = min(global_min, p_ref.min(), p_gen.min())
16     global_max = max(global_max, p_ref.max(), p_gen.max())
17
18 # Create heatmaps
19
20 for f, (p_ref, p_gen, ax_row) in enumerate(zip(pdists_ref, pdists_gen, axes.
21   T)):
22
23     heatmap_ref = np.zeros((nres, nres))
24     heatmap_gen = np.zeros((nres, nres))
```

```

17  cnt = 0
18
19  for i in range(len(heatmap_ref)):
20
21      for j in range(i + 1, nres):
22
23          heatmap_ref[j, i] = p_ref[0, cnt]
24
25          heatmap_gen[j, i] = p_gen[0, cnt]
26
27      cnt += 1
28
29  mask = np.triu(np.ones_like(heatmap_ref, dtype=bool))
30
31  sns.heatmap(heatmap_ref, mask=mask, cmap=cmap, ax=ax_row[0], cbar=False,
32
33      vmin=global_min, vmax=global_max)
34
35  sns.heatmap(heatmap_gen, mask=mask, cmap=cmap, ax=ax_row[1], cbar=False,
36
37      vmin=global_min, vmax=global_max)
38
39  ax_row[0].set_title(f'Frame {f+1}', fontsize=22)

```

After verifying that our DDPM is correctly adhering to the conditioning and predicting physically plausible structure, we employ the model to decode our synthetic MDN trajectory to high-resolution molecular structures. As visualized in Fig. 8b these structures are physically robust and retain characteristics of the MD data. The transformation from latent space trajectories to  $\text{C}\alpha$ -resolution synthetic structures is performed by the code below.

```

1 xyz_syn = model_ddpm.generate(torch.Tensor(synthetic_traj_CVs[:1000]))
2 xyz_syn = xyz_syn.reshape(xyz_syn.size(0), -1, 3).numpy()
3 trj_syn = md.Trajectory(xyz = xyz_syn, topology=trjs[0].top)

```

## 4 Conclusions

The LSS approach offers a powerful tool to expand the effective timescales of molecular dynamics simulations. The model leverages deep learning architectures to construct a surrogate dynamical model from modest training data, enabling the generation of synthetic trajectories

at a drastically reduced computational cost on the order of six orders of magnitude lower than classical MD. This approach facilitates dense sampling of molecular phase space and enables large reduction in the statistical errors in structural, thermodynamic, and kinetic observables. The LSS framework employs three distinct deep learning architectures to encode high-dimensional dynamics into a low-dimensional latent space, propagate dynamics within this space, and decode these dynamics back to the original molecular configuration. This software tutorial introduces the theoretical underpinnings and practical applications of LSS, with examples demonstrated on alanine dipeptide and the 28-residue BBA protein within accessible Python notebooks.

The LSS framework has already been extended to multi-molecular and highly distributed systems,<sup>30</sup> and going forward we are interested in developing conditioning on thermodynamic parameters such as temperature as well as chemical features like protein sequence. We also acknowledge current limitations in terms of the flexibility of the encoding and decoding frameworks. For example, molecular featurizations are currently selected by the user but could instead be learned by a structure-agnostic graph encoding similar to Schnet.<sup>88</sup> Additionally, the decoder currently requires alignment to a reference configuration, which may be ambiguous for larger and more complex systems, and could be replaced by an equivariant decoding mechanism. Indeed, the field is developing rapidly, and an advantage of the LSS framework is that novel architectures can be modularly deployed to enhance or replace any of the three components. Going forward, we hope that the molecular modeling community will find utility in the LSS framework, collaboratively build upon this approach by experimenting with novel learning methods and network architectures, and apply these models to increasingly complex and challenging systems.

## Conflict of Interest Statement

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, 63/521,617, and 63/669,836, and International Patent Applications PCT/US2020/035206, PCT/US2020/050466, and PCT/US24/10805.

## Data Availability

The LSS Python package is available free and open source from <https://github.com/Ferg-Lab/LSS> and the alanine dipeptide and BBA exercises are hosted, respectively, at [https://github.com/Ferg-Lab/IMSI\\_LSS](https://github.com/Ferg-Lab/IMSI_LSS) and [https://github.com/Ferg-Lab/LSS\\_BBA](https://github.com/Ferg-Lab/LSS_BBA).

## Author Information

**Michael S. Jones** Pritzker School of Molecular Engineering, The University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois 60637, United States.

**Kirill Shmilovich** Pritzker School of Molecular Engineering, The University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois 60637, United States. Present address: Pre-scient Design, Genentech Research and Early Development (gRED) Computational Science, Genentech, South San Francisco, California 94080, United States.

**Andrew L. Ferguson** Pritzker School of Molecular Engineering, The University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois 60637, United States.

email: [andrewferguson@uchicago.edu](mailto:andrewferguson@uchicago.edu)

## Acknowledgements

This material is based on work supported by the National Science Foundation under Grant No. CHE-2152521. 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. Part of this research was performed while the author was visiting the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-1929348).

## References

- (1) Frenkel, D.; Smit, B. *Understanding Molecular Simulation: From algorithms to applications*; Elsevier, 2023.
- (2) Tchipev, N.; Seckler, S.; Heinen, M.; Vrabec, J.; Gratl, F.; Horsch, M.; Bernreuther, M.; Glass, C. W.; Niethammer, C.; Hammer, N.; et al., TweTriS: Twenty trillion-atom simulation. *The International Journal of High Performance Computing Applications* **2019**, *33*, 838–854.
- (3) Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. How fast-folding proteins fold. *Science* **2011**, *334*, 517–520.
- (4) Karplus, M.; Petsko, G. A. Molecular dynamics simulations in biology. *Nature* **1990**, *347*, 631–639.
- (5) Noid, W. G. Perspective: Coarse-grained models for biomolecular systems. *The Journal of Chemical Physics* **2013**, *139*, 090901.
- (6) Noid, W. Perspective: Advances, challenges, and insight for predictive coarse-grained models. *The Journal of Physical Chemistry B* **2023**, *127*, 4174–4207.
- (7) Jin, J.; Pak, A. J.; Durumeric, A. E.; Loose, T. D.; Voth, G. A. Bottom-up coarse-graining: Principles and perspectives. *Journal of Chemical Theory and Computation* **2022**, *18*, 5759–5791.
- (8) Saunders, M. G.; Voth, G. A. Coarse-graining methods for computational biology. *Annual Review of Biophysics* **2013**, *42*, 73–93.
- (9) Padding, J.; Briels, W. J. Systematic coarse-graining of the dynamics of entangled polymer melts: the road from chemistry to rheology. *Journal of Physics: Condensed Matter* **2011**, *23*, 233101.

(10) Rudzinski, J. F. Recent progress towards chemically-specific coarse-grained simulation models with consistent dynamical properties. *Computation* **2019**, *7*, 42.

(11) Sidky, H.; Chen, W.; Ferguson, A. L. Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation. *Molecular Physics* **2020**, *118*, e1737742.

(12) 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. *Annual Review of Physical Chemistry* **2013**, *64*, 295–316.

(13) Abrams, C.; Bussi, G. Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. *Entropy* **2013**, *16*, 163–199.

(14) Yang, Y. I.; Shao, Q.; Zhang, J.; Yang, L.; Gao, Y. Q. Enhanced sampling in molecular dynamics. *The Journal of Chemical Physics* **2019**, *151*, 070902.

(15) Lyubartsev, A.; Martsinovski, A.; Shevkunov, S.; Vorontsov-Velyaminov, P. New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles. *The Journal of Chemical Physics* **1992**, *96*, 1776–1783.

(16) Bolhuis, P. G.; Chandler, D.; Dellago, C.; Geissler, P. L. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annual Review of Physical Chemistry* **2002**, *53*, 291–318.

(17) Chong, L. T.; Saglam, A. S.; Zuckerman, D. M. Path-sampling strategies for simulating rare events in biomolecular systems. *Current Opinion in Structural Biology* **2017**, *43*, 88–94.

(18) Escobedo, F. A.; Borrero, E. E.; Araque, J. C. Transition path sampling and forward flux sampling. Applications to biological systems. *Journal of Physics: Condensed Matter* **2009**, *21*, 333101.

(19) Bolhuis, P.; Dellago, C. Practical and conceptual path sampling issues. *The European Physical Journal Special Topics* **2015**, *224*, 2409–2427.

(20) Tiwary, P.; Parrinello, M. From metadynamics to dynamics. *Physical Review Letters* **2013**, *111*, 230602.

(21) Piccini, G.; McCarty, J. J.; Valsson, O.; Parrinello, M. Variational flooding study of a SN2 reaction. *The Journal of Physical Chemistry Letters* **2017**, *8*, 580–583.

(22) Ray, D.; Ansari, N.; Rizzi, V.; Invernizzi, M.; Parrinello, M. Rare event kinetics from adaptive bias enhanced sampling. *Journal of Chemical Theory and Computation* **2022**, *18*, 6500–6509.

(23) Schäfer, J.-L.; Keller, B. G. Implementation of Girsanov Reweighting in OpenMM and Deeptime. *The Journal of Physical Chemistry B* **2024**, *128*, 6014–6027.

(24) 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. *The Journal of Physical Chemistry A* **2023**, *127*, 3497–3517.

(25) Noé, F.; Olsson, S.; Köhler, J.; Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science* **2019**, *365*, eaaw1147.

(26) Coretti, A.; Falkner, S.; Weinreich, J.; Dellago, C.; von Lilienfeld, O. A. Boltzmann Generators and the New Frontier of Computational Sampling in Many-Body Systems. *arXiv preprint arXiv:2404.16566* **2024**,

(27) Pande, V. S.; Beauchamp, K.; Bowman, G. R. Everything you wanted to know about Markov state models but were afraid to ask. *Methods* **2010**, *52*, 99–105.

(28) Wehmeyer, C.; Scherer, M. K.; Hempel, T.; Husic, B. E.; Olsson, S.; Noé, F. Introduction to Markov State Modeling with the PyEMMA Software. *Living Journal of Computational Molecular Science* **2019**, *1*, 1–12.

(29) Sidky, H.; Chen, W.; Ferguson, A. L. Molecular latent space simulators. *Chemical Science* **2020**, *11*, 9459–9467.

(30) Jones, M. S.; McDargh, Z. A.; Wiewiora, R. P.; Izaguirre, J. A.; Xu, H.; Ferguson, A. L. Molecular latent space simulators for distributed and multimolecular trajectories. *The Journal of Physical Chemistry A* **2023**, *127*, 5470–5490.

(31) Noé, F. Machine learning for molecular dynamics on long timescales. *Lecture Notes in Physics* **2020**, *968*, 331–372.

(32) Fernández, A. Deep learning unravels a dynamic hierarchy while empowering molecular dynamics simulations. *Annalen der Physik* **2020**, *532*, 1–5.

(33) Klein, L.; Foong, A.; Fjelde, T.; Mlodzeniec, B.; Brockschmidt, M.; Nowozin, S.; Noe, F.; Tomioka, R. Timewarp: Transferable Acceleration of Molecular Dynamics by Learning Time-Coarsened Dynamics. Advances in Neural Information Processing Systems. 2023; pp 52863–52883 <https://proceedings.neurips.cc/paper-files/paper/2023/file/a598c367280f9054434fdcc227ce4d38-Paper-Conference.pdf> (Accessed 11 Oct 2024).

(34) García, A. E. Large-amplitude nonlinear motions in proteins. *Physical Review Letters* **1992**, *68*, 2696.

(35) Amadei, A.; Linssen, A. B.; Berendsen, H. J. Essential dynamics of proteins. *Proteins: Structure, Function, and Bioinformatics* **1993**, *17*, 412–425.

(36) Hegger, R.; Altis, A.; Nguyen, P. H.; Stock, G. How complex is the dynamics of peptide folding? *Physical Review Letters* **2007**, *98*, 028102.

(37) Das, P.; Moll, M.; Stamati, H.; Kavraki, L. E.; Clementi, C. Low-dimensional, free-energy landscapes of protein-folding reactions by nonlinear dimensionality reduction.

*Proceedings of the National Academy of Sciences of the United States of America* **2006**, *103*, 9885–9890.

(38) Ferguson, A. L.; Panagiotopoulos, A. Z.; Debenedetti, P. G.; Kevrekidis, I. G. Systematic determination of order parameters for chain dynamics using diffusion maps. *Proceedings of the National Academy of Sciences of the United States of America* **2010**, *107*, 13597–13602.

(39) Chen, W.; Sidky, H.; Ferguson, A. L. Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets. *Journal of Chemical Physics* **2019**, *150*, 214114.

(40) Noé, F.; Nüske, F. A variational approach to modeling slow processes in stochastic dynamical systems. *Multiscale Modeling and Simulation* **2013**, *11*, 635–655.

(41) Nuske, F.; Keller, B. G.; Pérez-Hernández, G.; Mey, A. S.; Noé, F. Variational approach to molecular kinetics. *Journal of Chemical Theory and Computation* **2014**, *10*, 1739–1752.

(42) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to advanced electronic structure theory*; Courier Corporation, 2012.

(43) Mori, H. Transport, collective motion, and Brownian motion. *Progress of Theoretical Physics* **1965**, *33*, 423–455.

(44) Zwanzig, R. Nonlinear generalized Langevin equations. *Journal of Statistical Physics* **1973**, *9*, 215–220.

(45) Zwanzig, R. *Nonequilibrium Statistical Mechanics*; Oxford University Press: Oxford, 2001.

(46) Risken, H.; Frank, T. *The Fokker-Planck Equation: Methods of Solution and Applications*, 2nd ed.; Springer Verlag: Berlin Heidelberg New York, 2012.

(47) Fu, X.; Xie, T.; Rebello, N. J.; Olsen, B. D.; Jaakkola, T. Simulate time-integrated coarse-grained molecular dynamics with geometric machine learning. *arXiv preprint arXiv:2204.10348* **2022**,

(48) Lin, K.; Peng, J.; Xu, C.; Gu, F. L.; Lan, Z. Realization of the trajectory propagation in the MM-SQC dynamics by using machine learning. *arXiv preprint arXiv:2207.05556* **2022**,

(49) Tsai, S.-T.; Kuo, E.-J.; Tiwary, P. Learning molecular dynamics with simple language model built upon long short-term memory neural network. *Nature Communications* **2020**, *11*, 5115.

(50) Vlachas, P. R.; Zavadlav, J.; Praprotnik, M.; Koumoutsakos, P. Accelerated simulations of molecular systems through learning of effective dynamics. *Journal of Chemical Theory and Computation* **2022**, *18*, 538–549.

(51) Wu, H.; Noé, F. Reaction coordinate flows for model reduction of molecular kinetics. *The Journal of Chemical Physics* **2024**, *160*, 044109.

(52) Mardt, A.; Pasquali, L.; Wu, H.; Noé, F. VAMPnets for deep learning of molecular kinetics. *Nature Communications* **2018**, *9*, 1–11.

(53) Bonati, L.; Piccini, G.; Parrinello, M. Deep learning the slow modes for rare events sampling. *Proceedings of the National Academy of Sciences of the United States of America* **2021**, *118*, e2113533118.

(54) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. *Journal of Chemical Physics* **2013**, *139*, 015102.

(55) Schwantes, C. R.; Pande, V. S. Improvements in Markov state model construction reveal

many non-native interactions in the folding of NTL9. *Journal of Chemical Theory and Computation* **2013**, *9*, 2000–2009.

(56) Molgedey, L.; Schuster, H. G. Separation of a mixture of independent signals using time delayed correlations. *Physical Review Letters* **1994**, *72*, 3634–3637.

(57) Andrew, G.; Arora, R.; Bilmes, J.; Livescu, K.; Deep canonical correlation analysis, In *Proceedings of the 30th International Conference on Machine Learning*; Dasgupta, S., McAllester, D., Eds.; 2013; Vol. 28; pp 1247–1255 <https://proceedings.mlr.press/v28/andrew13.html> (Accessed 11 Oct 2024).

(58) Klus, S.; Nüske, F.; Koltai, P.; Wu, H.; Kevrekidis, I.; Schütte, C.; Noé, F. Data-driven model reduction and transfer operator approximation. *Journal of Nonlinear Science* **2018**, *28*, 985–1010.

(59) Koltai, P.; Wu, H.; Noé, F.; Schütte, C. Optimal data-driven estimation of generalized Markov state models for non-equilibrium dynamics. *Computation* **2018**, *6*, 22.

(60) Wu, H.; Noé, F. Variational ppproach for learning Markov processes from time series data. *Journal of Nonlinear Science* **2020**, *30*, 23–66.

(61) McCandlish, D. M. Visualizing fitness landscapes. *Evolution* **2011**, *65*, 1544–1558.

(62) Herringer, N. S.; Dasetty, S.; Gandhi, D.; Lee, J.; Ferguson, A. L. Permutationally Invariant Networks for Enhanced Sampling (PINES): Discovery of multimolecular and solvent-inclusive collective variables. *Journal of Chemical Theory and Computation* **2023**, *20*, 178–198.

(63) Hempel, T.; Olsson, S.; Noé, F. Markov field models: Scaling molecular kinetics approaches to large molecular machines. *Current Opinion in Structural Biology* **2022**, *77*, 102458.

(64) Mardt, A.; Hempel, T.; Clementi, C.; Noé, F. Deep learning to decompose macromolecules into independent Markovian domains. *Nature Communications* **2022**, *13*, 7101.

(65) Bishop, C. M. Mixture density networks. *Neural Computing Research Group Report* **1994**, 171–198.

(66) Bishop, C. M. *Pattern Recognition and Machine Learning*; Springer-Verlag: New York, Inc., Secaucus, NJ, USA, 2006.

(67) Pathak, J.; Hunt, B.; Girvan, M.; Lu, Z.; Ott, E. Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters* **2018**, *120*, 24102.

(68) Arjovsky, M.; Chintala, S.; Bottou, L. Wasserstein Generative Adversarial Networks. Proceedings of the 34th International Conference on Machine Learning. 2017; pp 214–223 <https://proceedings.mlr.press/v70/arjovsky17a.html> (Accessed 11 Oct 2024).

(69) Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A. C. Improved Training of Wasserstein GANs. Advances in Neural Information Processing Systems. 2017 <https://proceedings.neurips.cc/paper-files/paper/2017/file/892c3b1c6dccc52936e27cbd0ff683d6-Paper.pdf> (Accessed 11 Oct 2024).

(70) Mirza, M.; Osindero, S. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* **2014**,

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

(72) Braams, B. J.; Bowman, J. M. Permutationally invariant potential energy surfaces in high dimensionality. *International Reviews in Physical Chemistry* **2009**, *28*, 577–606.

(73) Xie, Z.; Bowman, J. M. Permutationally invariant polynomial basis for molecular energy surface fitting via monomial symmetrization. *Journal of Chemical Theory and Computation* **2010**, *6*, 26–34.

(74) Behler, J.; Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical Review Letters* **2007**, *98*, 146401.

(75) Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *Journal of Chemical Physics* **2011**, *134*, 074106.

(76) Huo, H.; Rupp, M. Unified representation of molecules and crystals for machine learning. *Machine Learning: Science and Technology* **2022**, *3*, 045017.

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

(78) Harrigan, M. P.; Shukla, D.; Pande, V. S. Conserve water: A method for the analysis of solvent in molecular dynamics. *Journal of Chemical Theory and Computation* **2015**, *11*, 1094–1101.

(79) 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**,

(80) Bartók, A. P.; Kondor, R.; Csányi, G. On representing chemical environments. *Physical Review B* **2013**, *87*, 184115.

(81) Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*. 2020; pp 6840–6851

<https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html> (Accessed 11 Oct 2024).

(82) Wang, Y.; Herron, L.; Tiwary, P. From data to noise to data for mixing physics across temperatures with generative artificial intelligence. *Proceedings of the National Academy of Sciences of the United States of America* **2022**, *119*, e2203656119.

(83) Jones, M. S.; Shmilovich, K.; Ferguson, A. L. DiAMoNDBack: Diffusion-Denoising Autoregressive Model for Non-Deterministic Backmapping of C $\alpha$  protein traces. *Journal of Chemical Theory and Computation* **2023**, *19*, 7908–7923.

(84) Shmilovich, K.; Stieffenhofer, M.; Charron, N. E.; Hoffmann, M. Temporally coherent backmapping of molecular trajectories from coarse-grained to atomistic resolution. *The Journal of Physical Chemistry A* **2022**, *126*, 9124–9139.

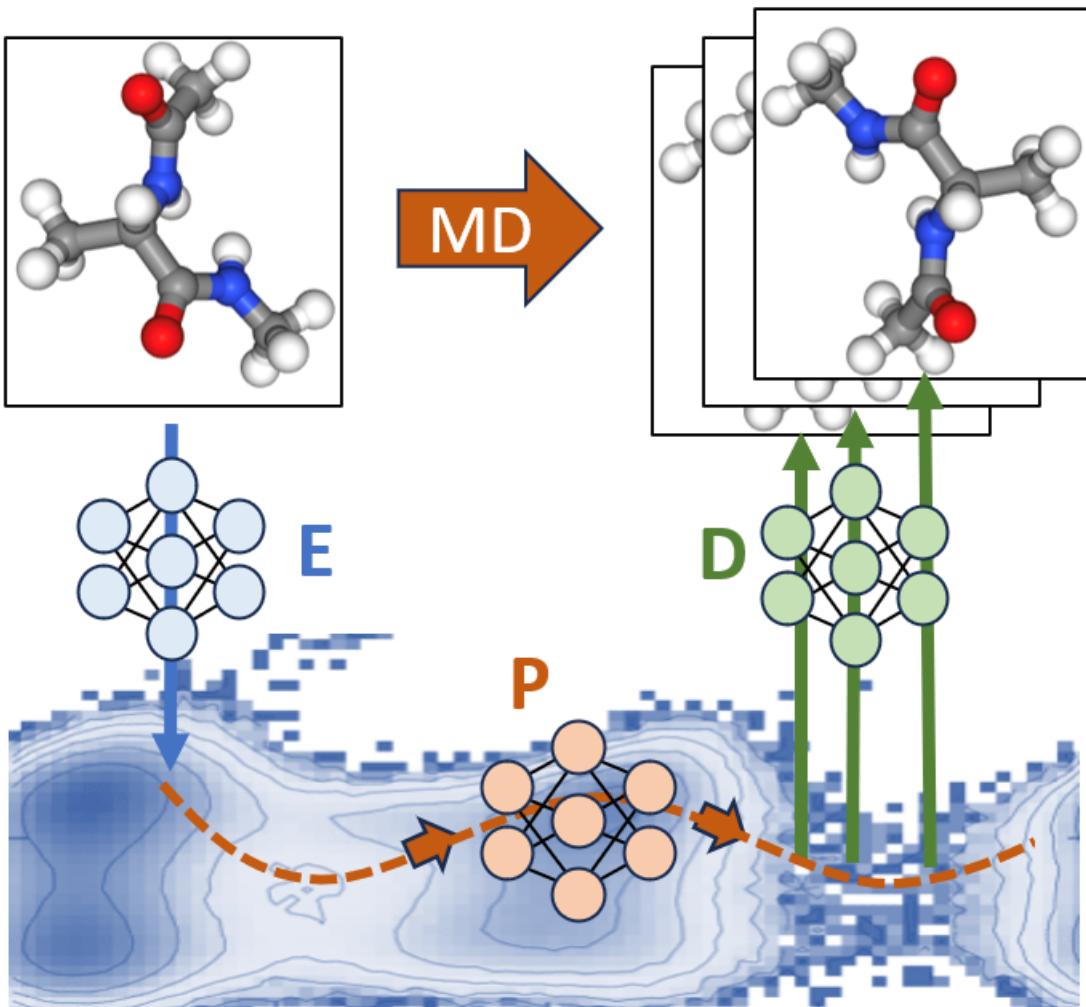
(85) Dobers, S.; Stark, H.; Fu, X.; Beaini, D.; Gündemann, S. Latent Space Simulator for Unveiling Molecular Free Energy Landscapes and Predicting Transition Dynamics. NeurIPS 2023 AI for Science Workshop. 2023 <https://openreview.net/forum?id=XIxclPy9c> (Accessed 11 Oct 2024).

(86) Sarisky, C. A.; Mayo, S. L. The  $\beta\beta\alpha$  fold: explorations in sequence space. *Journal of Molecular Biology* **2001**, *307*, 1411–1418.

(87) Hoffmann, M.; Scherer, M.; Hempel, T.; Mardt, A.; de Silva, B.; Husic, B. E.; Klus, S.; Wu, H.; Kutz, N.; Brunton, S. L.; Noé, F. Deeptime: a Python library for machine learning dynamical models from time series data. *Machine Learning: Science and Technology* **2022**, *3*, 015009.

(88) Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; Müller, K.-R. Schnet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics* **2018**, *148*, 241722.

# Graphical TOC Entry



TOC Graphic