

Girsanov Reweighting Enhanced Sampling Technique (GREST): On-the-Fly Data-Driven Discovery of and Enhanced Sampling in Slow Collective Variables

Kirill Shmilovich and Andrew L. Ferguson*

*Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637,
United States*

E-mail: andrewferguson@uchicago.edu

Abstract

Molecular dynamics simulations of microscopic phenomena are limited by the short integration time steps required for numerical stability but which limit the practically achievable simulation time scales. Collective variable (CV) enhanced sampling techniques apply biases to pre-defined collective coordinates to promote barrier crossing, phase space exploration, and sampling of rare events. The efficacy of these techniques is contingent on the selection of good CVs correlated with the molecular motions governing the long-time dynamical evolution of the system. In this work, we introduce Girsanov Reweighting Enhanced Sampling Technique (GREST) as an adaptive sampling scheme that interleaves rounds of data-driven slow CV discovery and enhanced sampling along these coordinates. Since slow CVs are inherently dynamical quantities, a key ingredient in our approach is the use of both thermodynamic and dynamical Girsanov reweighting corrections for rigorous estimation of slow CVs from biased simulation data. We demonstrate our approach on a toy 1D 4-well potential, a simple

biomolecular system alanine dipeptide, and the Trp-Leu-Ala-Leu-Leu (WLALL) pentapeptide. In each case GREST learns appropriate slow CVs and drives sampling of all thermally accessible metastable states starting from zero prior knowledge of the system. We make GREST accessible to the community via a publicly available open source Python package.

1 Introduction

Molecular dynamics (MD) simulations serve as a virtual microscope to provide microscopic understanding of the atomic-scale structure, thermodynamics, and dynamics of molecular and condensed-matter systems. By numerically integrating the classical equations of motion, MD simulations can accurately describe the evolution of molecular systems.¹ Advances in computational hardware^{2,3} and simulation software⁴⁻⁶ have enabled simulations of systems containing up to trillions of atoms.⁷ While the accessible size of systems continues to grow, a fundamental limitation of MD is the need for short integration time steps commensurate with the fastest atomic motions – typically on the order of femtoseconds – that limits practically achievable time scales to approximately microseconds to commodity hardware^{8,9} and milliseconds on specialized hardware.³ Many interesting processes such protein folding,¹⁰ protein-protein association,¹¹ ligand binding¹² and molecular self-assembly¹³ occur near or far beyond these attainable time scales, which has motivated the development of computational strategies to facilitate the study of slow molecular motions and rare events.

Techniques like the string method,¹⁴ transition path sampling,¹⁵ and forward flux sampling¹⁶ enable the study of rare events using astute initialization of simulation trajectories to explore the configuration space along paths connecting pre-defined “reactant” and “product” states (e.g., folded and unfolded states of a protein). Given these two end states, these techniques present efficient means to sample the pathways and associated dynamical properties of reactive transitions. Another class of popular methods are enhanced sampling techniques that seek to accelerate sampling not of a single reactive transition, but to

promote global exploration of configurational space by applying biasing potentials along predefined collective variables (CVs).¹⁷ Well chosen CVs correlated with transitions of interest and approximately orthogonal to separating free energy barriers can result in an increased frequency of barrier crossing to drive rare transitions and improve global sampling of configurational phase space. By exploiting this feature and modifying the system Hamiltonian to sample configurational space more efficiently, thermodynamic reweighting can be employed to estimate thermodynamic averages under the unbiased Hamiltonian. A plethora of methods have been developed for constructing biasing potentials along predefined CVs including, for example, metadynamics,¹⁸ hyperdynamics,¹⁹ umbrella sampling,²⁰ thermodynamic integration,²¹ and variational enhanced sampling.²² While these methods are demonstratively powerful tools for practitioners, their success is contingent upon the availability of good CVs in which to accelerate sampling and surmount the free energy barriers separating the thermally relevant metastable states. Accelerating poorly chosen CVs typically leads to little to no improvement in sampling over unbiased MD simulations.

The manual definition of good CVs is challenging for large, complex molecular systems where intuition is frequently limited. To help ameliorate this issue of proper CV selection, data-driven techniques can deploy machine learning (ML) to automatically discover optimal CVs for biasing. The central premise behind these approaches is an assumption of emergent low-dimensionality, and it is a generic feature of molecular systems that cooperative couplings between the atomic degrees of freedom lead to emergent low-dimensionality on sufficiently long time scales.^{23–25,25–28} Geometrically, one can view the dynamical evolution within the $3N$ -dimensional configurational phase space comprising the Cartesian coordinates of the N constituent atoms to be approximately restrained to a low-dimensional manifold supporting a dynamical attractor parameterized by the leading slow modes of the system (e.g., large-scale folding/unfolding of a protein) and to which the fast degrees of freedom (e.g., side chain rotations) are annealed as effective noise. In this sense, CV discovery can be viewed as a form of dimensionality reduction for identifying variables from data that best

characterize the slowest evolving (i.e., maximally autocorrelated) degrees of freedom. A central challenge in data-driven approaches for CV discovery is a “chicken-and-egg problem”: performing good configuration space sampling requires high-quality CVs to bias, but discovery of high-quality CVs requires sufficient data obtained from good configuration space sampling.²⁸ Consequently, data-driven methods typically strive to overcome this conundrum by iterating between rounds of enhanced sampling and CV learning to incrementally improve configuration space coverage. To this end, a variety of methods performing on-the-fly discovery of CVs and enhanced sampling have been proposed, including: REinforcement learning based Adaptive samPLing (REAP),²⁹ Reweighted Autoencoded Variational Bayes for Enhanced sampling (RAVE),³⁰ Diffusion-Map-directed MD (DM-d-MD),³¹ intrinsic Map Dynamics (iMapD),³² and Molecular Enhanced Sampling with Autoencoders (MESA).³³ A commonality between many of these adaptive sampling strategies is that they typically seek to identify high-variance CVs correlated with the elongated directions in configurational phase space containing the highest variance in the simulation trajectory. One may typically expect the high-variance directions to also be slowly evolving since they generally correspond to large-scale delocalized motions, but for the purposes of accelerating sampling along the slow collective motions, one would typically prefer an approach capable of identifying them directly.

Data-driven learning of slow CVs can be accomplished by appealing to the Variational Approach to Conformational dynamics (VAC).³⁴ The VAC formalism presents a variational approach to numerically estimate the maximally autocorrelated CVs by empirical analysis of time series data. The VAC underpins a number of popular techniques for slow CV discovery, including Time-Lagged Independent Component analysis (TICA),^{35–37} kernel TICA (kTICA),³⁸ and State-free Reversible VAMPNets (SRVs),³⁹ also known as Deep-TICA.⁴⁰ In particular, SRVs have been demonstrated as a powerful deep learning-enabled means to estimate slow variables from simulation data that have been applied to produce high-resolution models of Trp-cage miniprotein dynamics⁴¹ and describe DNA (de)hybridization.⁴² Crucially,

since slow CVs are inherently dynamic observables, their recovery from biased simulation data requires the use of dynamical, in addition to thermodynamic, corrections which are sufficient for the reweighting of purely configurational quantities.^{43–45} Methods to rigorously apply these dynamical corrections are necessary to solve the chicken and egg problem by iterative rounds of slow CV discovery and enhanced sampling exploration of configurational phase space.

A theoretical basis for dynamical reweighting is provided by path reweighting methods rooted in the Girsanov theorem.^{43,45–48} These approaches are similar to dynamical reweighting schemes that reweight discrete transition probabilities,^{49–53} but more powerful and general in that it is applied directly to the transition density elements and so does not require the assumption of local equilibrium or a predefined discretization of configurational phase space.^{43,47,48} Conceptually, Girsanov reweighting is a technique to correct the probability weights of dynamical pathways collected under a perturbed system Hamiltonian to those under an unperturbed Hamiltonian by applying dynamical reweighting of the transition density elements along the dynamical path through phase space.⁴⁸ Expectations of dynamical observables can be formulated as expectations in a path ensemble, and Girsanov reweighting provides a means to estimate unbiased averages from trajectories collected under an applied bias. Appealing to precisely the same rationale for enhanced sampling techniques in the estimation of configurational observables, applied biases that improve sampling of the phase space by accelerating rare transitions can lead to higher accuracy and lower variance estimates of dynamical observables.

Path reweighting approaches present a powerful tool for the estimation of dynamical observables from biased trajectories, but do have two important limitations. First, analytical expressions for path reweighting factors are only available for particular choices of stochastic integrators, specifically the Euler-Maruyama (E-M) scheme to integrate the overdamped Langevin (i.e., Brownian dynamics) equation^{43,47,48} and the Izaguirre, Sweet, and Pande (ISP) scheme – a variant of the Langevin leapfrog scheme – to integrate the (underdamped)

Langevin equation.^{43,54} Elegant recent work by Kieninger and Keller developed a prescription to derive path weights for other Langevin integration schemes and presented evidence that approximate E-M reweighting expressions may offer high accuracy approximations for other Langevin integrators, thereby generalizing the applicability of the methodology to popular numerical schemes employed in the molecular simulation community.⁴³ Second, the path reweighting factors can become extremely large or extremely small for large applied biases risking numerical instability and overflow/underflow. This means that there typically exists a “Goldilocks region” for the intensity of the applied bias that is sufficiently high to promote good sampling but sufficiently low for stable convergence and low-variance estimation of the path ensemble expectations.^{45,47,48}

There have been relatively few applications of Girsanov reweighting to molecular systems. Viewing temperature as a simple modification of the system Hamiltonian, Chodera, Swope, Noé, Prinz, Shirts, Smith, and Pande performed path reweighting of simulation trajectories at various temperatures within a parallel tempering / replica exchange framework to estimate dynamical expectations from simulations conducted different temperatures under an Anderson thermostat, overdamped Langevin dynamics, and Langevin dynamics.^{44,55} Weber and Pande employed the Girsanov approach to construct Markov state models (MSMs) for the 20-residue Trp-cage fast-folding mini-protein using biased simulation trajectories in which the dynamical evolution of a key dihedral angle was accelerated using artificial biasing forces.⁴⁵ Donati and Keller presented an elegant theoretical basis and devised efficient numerical schemes that were applied to construct Markov state models of alanine dipeptide, valine dipptide, VGVAPG hexapeptide, and a β -hairpin peptide by on-the-fly Girsanov reweighting of metadynamics enhanced sampling trajectories in which accelerating biases were applied to pre-selected backbone or side-chain dihedral angles, key hydrogen bonds, or the molecular end-to-end distance.^{47,48} In all studies to date, the CVs to accelerate, or more generally the variables/parameters in which the Hamiltonian was modified, were pre-defined. The success of the CV enhanced sampling is predicated on knowledge of CVs coincident with

the slow motions of the system, and has limited applications to systems where this knowledge is *a priori* available.

In the present work, we present Girsanov Reweighting Enhanced Sampling Technique (GREST) as an adaptive sampling scheme that performs interleaved rounds of data-driven slow CV discovery and enhanced sampling along those coordinates. Our approach is analogous to our previously introduced Molecular Enhanced Sampling with Autoencoders (MESA) method for interleaved data-driven discovery and enhanced sampling in *high variance* CVs,³³ but by appealing to Girsanov reweighting GREST extends this approach to *slow* (i.e., maximally autocorrelated) CVs. We build upon the elegant theoretical and numerical work by Keller and co-workers^{47,48} to perform nonlinear CV discovery using SRVs as a deep learning-enabled version of TICA that employs both thermodynamic corrections and Girsanov path weights to recover slow CVs from the biased simulation trajectories.³⁹ This modification of the SRV objective function enables the formally correct numerical estimation of slow collective variables from short, biased, and discontinuous trajectories. This estimation process is similar to that proposed by Bonati et al.,⁴⁰ but we explicitly include the Girsanov path weights in addition to the thermodynamic corrections to rigorously account for the influence of the bias upon the dynamics, and, rather than a one-shot discovery paradigm, we perform discovery within a virtuous iterative loop to simultaneously converge sampling of configurational phase space and the learned slow CVs. We demonstrate GREST in applications to two biomolecular systems – alanine dipeptide and the WLALL pentapeptide – and make the approach freely available to the community via publicly available open source Python packages accompanied with example notebooks. We anticipate that GREST may be a useful tool for rapid exploration of the configurational phase space of molecular systems without the requirement for prior knowledge of the important underlying slow dynamical modes, and a means to advance fundamental understanding of the slow molecular modes, guide rational molecular engineering, and furnish slow subspaces for the parameterization of efficient reduced-dimensional dynamical simulators.^{56,57}

2 Methods

2.1 Transfer operator theory

The transfer operator $\mathcal{T}_\tau$ of a dynamical system possessing configurational microstates $\mathbf{x} \in \mathbb{R}^n$ propagates probability distributions over configurational microstates $q_t(\mathbf{x})$ normalized with respect to the equilibrium probability distribution $u_t(\mathbf{x}) = q_t(\mathbf{x})/\pi(\mathbf{x})$ under transition densities $p(\mathbf{x}_{t+\tau}|\mathbf{x}_t)$ at a lag time τ ,^{58,59}

$$u_{t+\tau}(\mathbf{x}) = \mathcal{T}_\tau \circ u_t(\mathbf{x}) = \frac{1}{\pi(\mathbf{x})} \int p_\tau(\mathbf{x}|\mathbf{x}') u_t(\mathbf{x}') \pi(\mathbf{x}') d\mathbf{x}'. \quad (1)$$

Here $p_\tau(\mathbf{x}|\mathbf{x}') = \mathbb{P}(\mathbf{x} = \mathbf{x}_{t+\tau}|\mathbf{x}' = \mathbf{x}_t)$ is the (normalized) conditional probability density that a system in microstate $\mathbf{x}' = \mathbf{x}_t$ at time t will advance to state $\mathbf{x} = \mathbf{x}_{t+\tau}$ at time $(t + \tau)$. Crucially, Eqn. 1 assumes Markovianity, which is an increasingly good approximation at large lag times τ , and time homogeneity of the transition densities $p_\tau(\mathbf{x}|\mathbf{x}')$. For equilibrium systems satisfying detailed balance, we have the additional condition that,

$$\pi(\mathbf{x}) p_\tau(\mathbf{x}'|\mathbf{x}) = \pi(\mathbf{x}') p_\tau(\mathbf{x}|\mathbf{x}'), \quad (2)$$

which implies the transfer operator $\mathcal{T}_\tau$, also known as the Perron-Frobenius operator describing the evolution of densities with respect to the equilibrium density, becomes identical to the Koopman operator describing the evolution of observables, and self-adjoint with respect

to the equilibrium distribution,⁵⁸

$$\begin{aligned}
\langle \mathcal{T}_\tau \circ g, h \rangle_\pi &= \int \mathcal{T}_\tau \circ g(\mathbf{x}) h(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x} \\
&= \int \left[\frac{1}{\pi(\mathbf{x})} \int p_\tau(\mathbf{x}|\mathbf{x}') g(\mathbf{x}') \pi(\mathbf{x}') d\mathbf{x}' \right] h(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x} \\
&= \int \int \frac{\pi(\mathbf{x})}{\pi(\mathbf{x}')} p_\tau(\mathbf{x}'|\mathbf{x}) h(\mathbf{x}) g(\mathbf{x}') \pi(\mathbf{x}') d\mathbf{x}' d\mathbf{x} \\
&= \int \left[\frac{1}{\pi(\mathbf{x}')} \int p_\tau(\mathbf{x}'|\mathbf{x}) h(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x} \right] g(\mathbf{x}') \pi(\mathbf{x}') d\mathbf{x}' \\
&= \int \mathcal{T}_\tau \circ h(\mathbf{x}') g(\mathbf{x}') \pi(\mathbf{x}') d\mathbf{x}' \\
&= \langle g, \mathcal{T}_\tau \circ h \rangle_\pi.
\end{aligned} \tag{3}$$

The self-adjoint nature of $\mathcal{T}_\tau$ under the detailed balance condition implies that it admits diagonalization into a set of eigenfunctions $\{\psi_i(\mathbf{x})\}$ forming a complete orthonormal basis, and that the associated eigenvalues $\{\lambda_i\}$ are real and bounded from above by $1 = \lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots$,^{58,60–62}

$$\mathcal{T}_\tau \circ \psi_i(\mathbf{x}) = \lambda_i \psi_i(\mathbf{x}), \quad \langle \psi_i(\mathbf{x}) | \psi_j(\mathbf{x}) \rangle_\pi = \delta_{ij}. \tag{4}$$

As a result, any state function $\xi_t(\mathbf{x})$ can be represented within the eigenfunction basis $\xi_t(\mathbf{x}) = \sum_i \langle \psi_i(\mathbf{x}) | \xi_t(\mathbf{x}) \rangle_\pi \psi_i(\mathbf{x})$, and we can describe the dynamical evolution of ξ_t in increments of τ via consecutive applications of $\mathcal{T}_\tau$,^{39,56}

$$\begin{aligned}
\xi_{t+k\tau}(\mathbf{x}) &= \underbrace{\mathcal{T}_\tau \circ \mathcal{T}_\tau \circ \dots \circ \mathcal{T}_\tau}_k \circ \xi_t(\mathbf{x}) \\
&= \sum_i \langle \psi_i(\mathbf{x}) | \xi_t(\mathbf{x}) \rangle_\pi \mathcal{T}_\tau^k \circ \psi_i(\mathbf{x}) \\
&= \sum_i \langle \psi_i(\mathbf{x}) | \xi_t(\mathbf{x}) \rangle_\pi \lambda_i^k \psi_i(\mathbf{x}) \\
&= \sum_i \langle \psi_i(\mathbf{x}) | \xi_t(\mathbf{x}) \rangle_\pi e^{\left(-\frac{k\tau}{\tau_i}\right)} \psi_i(\mathbf{x}),
\end{aligned} \tag{5}$$

where we have transformed the eigenvalues into the implied time scales $t_i = -\tau / \ln(\lambda_i)$ associated with the relaxation of each of the orthonormal eigenfunctions in the expansion. These time scales t_i characterize the exponentially decaying contributions for the eigenfunctions $\{\psi_{i>0}(\mathbf{x})\}$, such that at sufficiently large lag times τ the contributions from basis associated with faster relaxation times can be neglected and ξ accurately approximated by keeping only the leading slowest terms. The leading pair ($\psi_0 = \mathbf{1}, \lambda_0 = 1$) therefore corresponds to the equilibrium distribution $\pi(\mathbf{x}) = \lim_{t \rightarrow \infty} q_t(\mathbf{x})$ and the remaining pairs to a hierarchy of increasingly quickly relaxing modes.

2.2 Approximating slow modes from data

The Variational Approach to Conformation dynamics (VAC) formalism provides a means to approximate the eigenfunctions $\{\psi_i(\mathbf{x})\}$ of the transfer operator $\mathcal{T}_\tau$. A variational principle underpins the VAC approach.^{38,39,61} This can be straightforwardly understood by starting from an identity that for the i^{th} eigenfunction $\psi_i(\mathbf{x})$ that follows directly from Eqn. 4,

$$\frac{\langle \psi_i(\mathbf{x}) | \mathcal{T}_\tau \circ \psi_i(\mathbf{x}) \rangle_\pi}{\langle \psi_i(\mathbf{x}) | \psi_i(\mathbf{x}) \rangle_\pi} = \lambda_i. \quad (6)$$

Consider now a trial function $\nu(\mathbf{x})$ that is orthogonal to the preceding $(i-1)$ eigenfunctions,

$$\langle \psi_0(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi = 0, \langle \psi_1(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi = 0, \dots, \langle \psi_{i-1}(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi = 0, \quad (7)$$

and may therefore be expanded within the eigenfunction basis as,

$$\nu(\mathbf{x}) = \sum_{j \geq i} \langle \psi_j(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi \psi_j(\mathbf{x}) = \sum_{j \geq i} \kappa_j \psi_j(\mathbf{x}), \quad (8)$$

where $\kappa_j = \langle \psi_j(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi$ are the linear expansion coefficients. Combining Eqns. 4, 6, and 8, we obtain the following inequality,

$$\begin{aligned}
\tilde{\lambda}_i &= \frac{\langle \nu(\mathbf{x}) | \mathcal{T}_\tau \circ \nu(\mathbf{x}) \rangle_\pi}{\langle \nu(\mathbf{x}) | \nu(\mathbf{x}) \rangle_\pi} \\
&= \frac{\langle \sum_{j \geq i} \kappa_j \psi_j(\mathbf{x}) | \mathcal{T}_\tau \circ \sum_{j \geq i} \kappa_j \psi_j(\mathbf{x}) \rangle_\pi}{\langle \sum_{j \geq i} \kappa_j \psi_j(\mathbf{x}) | \sum_{j \geq i} \kappa_j \psi_j(\mathbf{x}) \rangle_\pi} \\
&= \frac{\sum_{j \geq i} \kappa_j^2 \lambda_j}{\sum_{j \geq i} \kappa_j^2} \\
&\leq \frac{\sum_{j \geq i} \kappa_j^2 \lambda_i}{\sum_{j \geq i} \kappa_j^2} \\
&= \lambda_i,
\end{aligned} \tag{9}$$

such that $\tilde{\lambda}_i$ is bounded from above by the true eigenvalue λ_i and the equality $\tilde{\lambda}_i = \lambda_i$ holds if and only if the trial function $\nu(\mathbf{x})$ is equal to the true eigenfunction $\nu(\mathbf{x}) = \psi_i(\mathbf{x})$. As such, we are guaranteed to approach λ_i from below by variational optimization of a trial function $\nu(\mathbf{x})$ to maximize $\tilde{\lambda}_i$ in Eqn. 9 subject to the orthogonality conditions in Eqn. 7.

It is convenient to construct trial functions as a linear expansion of arbitrary (possibly non-orthogonal) basis functions $\{\chi_i(\mathbf{x})\}$ and optimize the expansion coefficients within this basis to develop optimal approximations $\{\phi_i(\mathbf{x})\}$ to the true eigenfunctions $\{\psi_i(\mathbf{x})\}$,

$$\phi_i(\mathbf{x}) = \sum_j \alpha_{ij} \chi_j(\mathbf{x}), \tag{10}$$

where α_{ij} are the linear expansion coefficients for basis function j within our approximation for eigenfunction i . Under this linear formulation, the VAC procedure can be shown to yield the following generalized eigenvalue problem,⁶¹

$$C^\tau \boldsymbol{\alpha}_i = \tilde{\lambda}_i S \boldsymbol{\alpha}_i, \tag{11}$$

where $\boldsymbol{\alpha}_i = \{\alpha_{ij}\} = [\alpha_{i0}, \alpha_{i1}, \alpha_{i2}, \dots]$ are the linear expansion coefficients for assembling the

approximate eigenfunction $\phi_i(\mathbf{x})$ with associated approximate eigenvalue $\tilde{\lambda}_i$. This generalized eigenvalue problem is isomorphic to the Roothaan-Hall equations in quantum mechanics that are used to represent the Hartree-Fock equation in a non-orthonormal basis set.⁶³ The matrix elements of C^τ and S are defined as,

$$C_{jk}^\tau = \langle \chi_j(\mathbf{x}) | \mathcal{T}_\tau \circ \chi_k(\mathbf{x}) \rangle_\pi, \quad (12)$$

$$S_{jk} = \langle \chi_j(\mathbf{x}) | \chi_k(\mathbf{x}) \rangle_\pi, \quad (13)$$

where C^τ is the time-lagged correlation matrix – analogous to the Fock matrix – and S the overlap matrix. Standard techniques⁶⁴ can be employed to solve the generalized eigenvalue problem in Eqn. 11, which will yield the linear expansion coefficients for the optimal approximate eigenfunctions $\{\phi_i(\mathbf{x})\}$ within the basis $\{\chi_i(\mathbf{x})\}$.

Using this VAC formalism we can numerically estimate $\{\phi_i(\mathbf{x})\}$ from trajectory data by replacing the exact inner products in Eqns. 12 and 13 with empirical estimates,

$$\begin{aligned} C_{jk}^\tau &= \langle \chi_j(\mathbf{x}) | \mathcal{T}_\tau \circ \chi_k(\mathbf{x}) \rangle_\pi \approx \mathbb{E}[\chi_j(\mathbf{x}_t) \chi_k(\mathbf{x}_{t+\tau})] \\ &\approx \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} \chi_j(\mathbf{x}_t) \chi_k(\mathbf{x}_{t+\tau}) \\ &\approx \frac{1}{L} \sum_{i=1}^L \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_{t+\tau}^{(i)}), \end{aligned} \quad (14)$$

$$\begin{aligned} S_{jk} &= \langle \chi_j(\mathbf{x}) | \chi_k(\mathbf{x}) \rangle_\pi \approx \mathbb{E}[\chi_j(\mathbf{x}_t) \chi_k(\mathbf{x}_t)] \\ &\approx \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} \chi_j(\mathbf{x}_t) \chi_k(\mathbf{x}_t) \\ &\approx \frac{1}{L} \sum_{i=1}^L \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_t^{(i)}), \end{aligned} \quad (15)$$

where the terminal summations defining the estimators in Eqns. 14 and 15 are carried out over $i = 1 \dots L$ training samples harvested from the trajectory data assembled as pairs of

states $(\mathbf{x}_t^{(i)}, \mathbf{x}_{t+\tau}^{(i)})$ separated by a lag time τ .

2.3 State-free Reversible VAMPnets (SRVs)

The quality of the eigenfunction approximations under the VAC is contingent on the choice of a sufficiently expressive and powerful basis. The Markov state model (MSM) formalism^{65,66} follows as a special case of the VAC under a choice of basis corresponding to indicator functions partitioning of the configurational phase space into n_c disjoint sets $\{S_j\}_{j=1}^{n_c}$ such that $\chi_j(\mathbf{x}) = 1$ for $\mathbf{x} \in S_j$ and $\chi_j(\mathbf{x}) = 0$ for $\mathbf{x} \notin S_j$.³⁴ In this case estimation of the matrix elements C_{jk}^τ simply amounts to counting transitions from trajectory data that carry states $\mathbf{x}_t \in S_j$ to $\mathbf{x}_{t+\tau} \in S_k$ in a lag time τ . The inherent orthonormality and binary nature of the indicator function basis $\{\chi_i\}$ implies that the matrix $\mathbf{S}$ is diagonal with matrix elements $S_{jj} = \pi_j$ corresponding to the stationary distribution of normalized state counts $\pi_j = \frac{\sum_{t=1}^{T-\tau} \chi_j(\mathbf{x}_t)}{T-\tau}$.⁶¹ The elements of the approximate MSM eigenfunctions $\{\phi_i(\mathbf{x})\}$ can then be interpreted as fluxes into or out of sets $\{S_j\}$ with each eigenfunction $\phi_i(\mathbf{x})$ characterizing slow modes as traversals along the discretized state-space $\{S_j\}$.⁶⁷ The quality of state-space discretization is a critical determinant of MSM precision, and is typically most amenable to metastable dynamics that can be effectively discretized into long-lived metastable states.^{61,68} State-space clustering to determine disjoint sets that define the basis $\{\chi_i\}$ is commonly performed within low-dimensional embeddings of the full-dimensional state-space using techniques such as TICA³⁵⁻³⁷ combined with k-means clustering. The different interacting components within MSM workflows typically suffer from requiring manual trial-and-error and human expertise to handle selecting the appropriate dimensionality reduction technique, the number of cluster states, choice of clustering method, and optimizing location of cluster centers.⁶¹ Mardt, Pasquali, Wu, and Noé developed VAMPnets as an end-to-end replacement for MSM construction pipeline using deep neural networks to integrate the featurization, dimensionality reduction, clustering, and kinetic modeling steps.⁶⁹ Indicator functions represent one rather restrictive choice of basis and the quality

of the approximate eigenfunction expansions depends sensitively upon the choice of phase space discretization. More general classes of basis functions offer the possibility of superior eigenfunction approximations, but it is challenging to intuit good choices of a basis set.

We previously introduced State free Reversible VAMPNets (SRVs) as a means to perform data-driven learning of good basis functions for the VAC by performing simultaneous learning of the basis functions and the linear expansion coefficients for approximation of the transfer operator eigenfunctions.³⁹ In brief, we learn $\{\chi_i\}$ from training trajectories using simple feedforward neural networks, pass these basis functions through the VAC to learn approximations for the leading transfer operator eigenfunctions $\{\phi_i(\mathbf{x})\}$, and quantify the quality of the learned eigenfunctions via a loss function based on the slowness of the associated implied time scales. We then iteratively update the neural networks to learn superior basis functions by backpropagating the gradient of the loss function through the entire computational graph.

Mathematically, the SRV neural networks $f^\theta(\mathbf{x}) \in \mathbb{R}^m$ are parameterized with model weights θ and tasked with learning a map from microstate configurations $\mathbf{x}$ into a m -dimensional output space, where m is the number of desired approximate eigenfunctions of the transfer operator (i.e., slow CVs) to learn. The neural network outputs $\chi_j(\mathbf{x}) = f_j^\theta(\mathbf{x})$ represent the learned basis functions, which we pass to the VAC and solve for the optimal linear expansion approximations of the transfer operator eigenfunctions $\phi_i(\mathbf{x}) = \sum_j \alpha_{ij} \chi_j(\mathbf{x})$ using Eqns. 11-15. The model weights θ are optimized via minimization of the VAMP-r loss function,

$$\mathcal{L}_{SRV} = - \sum_{j=0}^{m-1} \tilde{\lambda}_j^r, \quad (16)$$

with a typical choice of $r = 2$ corresponding to the VAMP-2 loss that can be interpreted as the cumulative kinetic variance.^{67,70} The loss function attains a minimum when the approximate eigenfunctions equal the true eigenfunctions and the approximation becomes exact.⁶⁹

Learning iterates between solving for the expansion coefficients $\{\alpha_{ij}\}$ and eigenvalues $\{\tilde{\lambda}_j\}$ by solving the generalized eigenvalue problem in Eqn. 11 and updating the network

weights θ by gradient descent steps of $\mathcal{L}_{SRV}$ in Eqn. 16. The generalized eigenvalue problem is formulated by estimating the matrices C^τ and S from the neural network outputs in batches according to Eqns. 14 and 15, and solved using Cholesky decomposition to provide stable gradients for the learning problem.³⁹ The neural network weights are updated by mini-batch gradient descent with the Adam⁷¹ optimizer. Neural network convergence is assessed by monitoring and ensuring plateauing of the loss $\mathcal{L}_{SRV}$ on a held-out validation set, which together with providing randomly shuffled training samples, also serves as a regularizer to also help prevent overfitting.

The trained SRV may then be used for out of sample inference on unseen configurations $\mathbf{x}^*$ to furnish the mappings into the approximate eigenfunctions $\mathbf{x}^* \rightarrow \{\phi_i(\mathbf{x}^*)\}_{i=0}^{m-1}$. This learned projection represents a data-driven embedding into a slow CV subspace spanned by our approximations to the leading eigenfunctions of the transfer operator. An appropriate number m of CVs to retain is typically informed by a gap in the eigenvalue spectrum separating the slowly relaxing modes from a rapidly relaxing continuum of fast modes that are effectively annealed to the slow subspace on sufficiently long time scales. An appropriate lag time τ is determined by convergence of the implied time scales of the learned slow modes and satisfaction of the Chapman-Kolmogorov test.³⁹ Full details of the architecture, training, and hyperparameter tuning of the SRV models employed in this work are provided in the Supporting Information.

2.4 Dynamical reweighting

Dynamical reweighting provides a formalism to estimate equilibrium dynamical properties from biased simulation data. Dynamical reweighting corrections for biased data have been analytically derived by Chodera and co-workers for the special case of temperature reweighting,^{44,72,73} and Keller and co-workers for arbitrary biases.^{43,47,48} The premise of dynamical reweighting lies in path reweighting methods, based mathematically on the Girsanov theorem,⁴⁶ in which phase space trajectories $\omega_{0 \rightarrow \tau} = ((\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_{\Delta t}, \mathbf{v}_{\Delta t}), \dots, (\mathbf{x}_{\tau-\Delta t}, \mathbf{v}_{\tau-\Delta t}), (\mathbf{x}_\tau, \mathbf{v}_\tau))$

generated under some simulation potential $V_{sim}(\mathbf{x})$ are associated with a path probability density $\mu_{\Omega_{\tau}, \mathbf{x}_0}^{sim}(\boldsymbol{\omega}_{0 \rightarrow \tau})$, where $\Omega_{\tau}, \mathbf{x}_0$ is the path space comprising the set of all paths of length τ that start at $\mathbf{x}_0$. Positing that the simulation potential $V_{sim}(\mathbf{x}) = V_{target}(\mathbf{x}) - U_{bias}(\mathbf{x})$ is a superposition of the underlying target potential $V_{target}(\mathbf{x})$ and a biasing potential $U_{bias}(\mathbf{x})$, path weights $M(\boldsymbol{\omega}_{0 \rightarrow \tau})$ can be used to recover the probabilities for paths collected under the simulation potential V_{sim} by reweighting to the target potential V_{target} ,

$$\mu_{\Omega_{\tau}, \mathbf{x}_0}^{target}(\boldsymbol{\omega}_{0 \rightarrow \tau}) = M(\boldsymbol{\omega}_{0 \rightarrow \tau}) \mu_{\Omega_{\tau}, \mathbf{x}_0}^{sim}(\boldsymbol{\omega}_{0 \rightarrow \tau}). \quad (17)$$

Until recently, analytical expressions of the path weights $M(\boldsymbol{\omega}_{0 \rightarrow \tau})$ for arbitrary biasing potentials were only available for the Euler-Maruyama scheme for overdamped Langevin dynamics (i.e., Brownian dynamics).^{43,74–77} More recently, analytical path weights have been derived for (underdamped) Langevin dynamics schemes⁴³ which are known to better reproduce the true dynamics for large molecular systems and are widely implemented within existing molecular dynamics frameworks.^{54,78–83} The molecular systems in this work are simulated in implicit solvent under (underdamped) Langevin dynamics that are numerically integrated using a variant of the Langevin leapfrog scheme developed by Izaguirre, Sweet, and Pande,⁵⁴ hereafter referred to as the ISP scheme, for which analytical path weights have been derived by Kieninger and Keller.⁴³ These authors have also implemented this integrator within OpenMM⁸⁴ along with numerically efficient implementations for on-the-fly calculation of the path weights $M(\boldsymbol{\omega}_{0 \rightarrow \tau})$ during runtime.

Co-opting this dynamical reweighting formalism within the SRV framework enables estimation of equilibrium slow modes from biased simulation data.^{47,48} Following the development of Donati and Keller,⁴⁸ we observe that the matrix elements C_{jk}^{τ} of the time-lagged correlation matrix of basis functions $\{\chi_i(\mathbf{x})\}$ under the target potential $V_{target}(\mathbf{x})$ can be viewed as a nested integral of a path ensemble average over paths $\boldsymbol{\omega}_{t \rightarrow t+\tau} \in \Omega_{\tau, \mathbf{x}_t}$ inside a

phase space ensemble average over path starting configurations $\mathbf{x}_t \in \Gamma$,

$$C_{jk}^\tau = \int_{\Gamma} \mu_{\Gamma}^{target}(\mathbf{x}_t) \chi_j(\mathbf{x}_t) \int_{\Omega_{\tau, \mathbf{x}_t}} \mu_{\Omega_{\tau, \mathbf{x}_t}}^{target}(\boldsymbol{\omega}_{t \rightarrow t+\tau}) \chi_k(\mathbf{x}_{t+\tau}) d\boldsymbol{\omega}_{t \rightarrow t+\tau} d\mathbf{x}_t, \quad (18)$$

where $\mu_{\Gamma}^{target}(\mathbf{x}_t)$ and $\mu_{\Omega_{\tau, \mathbf{x}_t}}^{target}(\boldsymbol{\omega}_{t \rightarrow t+\tau})$ are, respectively, the phase-space and path-space probability densities under $V_{target}(\mathbf{x})$. For simulation data collected under a potential $V_{sim}(\mathbf{x}) = V_{target}(\mathbf{x}) - U_{bias}(\mathbf{x})$, we are compelled to apply thermodynamic and dynamical corrections to estimate C_{jk}^τ under $V_{target}(\mathbf{x})$,

$$C_{jk}^\tau = \int_{\Gamma} h(\mathbf{x}_t) \mu_{\Gamma}^{sim}(\mathbf{x}_t) \chi_j(\mathbf{x}_t) \int_{\Omega_{\tau, \mathbf{x}_t}} M(\boldsymbol{\omega}_{t \rightarrow t+\tau}) \mu_{\Omega_{\tau, \mathbf{x}_t}}^{sim}(\boldsymbol{\omega}_{t \rightarrow t+\tau}) \chi_k(\mathbf{x}_{t+\tau}) d\boldsymbol{\omega}_{t \rightarrow t+\tau} d\mathbf{x}_t. \quad (19)$$

The thermodynamic reweighting correction accounts for the change in the configurational probability of the initial configuration of the path $\mathbf{x}_t$ under V_{target} relative to V_{sim} ,

$$\begin{aligned} h(\mathbf{x}_t) &= \frac{\mu_{\Gamma}^{target}(\mathbf{x}_t)}{\mu_{\Gamma}^{sim}(\mathbf{x}_t)} \\ &= \frac{\exp(-\beta V_{target}(\mathbf{x}_t))}{Z_{target}} \bigg/ \frac{\exp(-\beta V_{sim}(\mathbf{x}_t))}{Z_{sim}} \\ &= \frac{Z_{sim}}{Z_{target}} \exp(-\beta U_{bias}(\mathbf{x}_t)), \end{aligned} \quad (20)$$

where we operate in the canonical ensemble such that $Z = \int_{\Gamma} \exp(-\beta V(\mathbf{x})) d\mathbf{x}$ is the configurational part of the canonical partition function associated with a potential $V(\mathbf{x})$, and $\beta = (k_B T)^{-1}$, where k_B is the Boltzmann constant, and T is the temperature. The dynamical reweighting correction accounts for the change in the probability of the path under V_{target} relative to V_{sim} ,

$$M(\boldsymbol{\omega}_{t \rightarrow t+\tau}) = \frac{\mu_{\Omega_{\tau, \mathbf{x}_t}}^{target}(\boldsymbol{\omega}_{t \rightarrow t+\tau})}{\mu_{\Omega_{\tau, \mathbf{x}_t}}^{sim}(\boldsymbol{\omega}_{t \rightarrow t+\tau})}. \quad (21)$$

For underdamped Langevin dynamics numerically integrated using the ISP scheme,⁵⁴ an explicit expression for this dynamical reweighting factor is known (c.f., Eqns. 23, 57, 58, 59

and 60 in Ref. ⁴³), ^{43,47,48}

$$\begin{aligned}
M(\boldsymbol{\omega}_{t \rightarrow t+\tau}) = & \exp \left(- \sum_{\alpha=1}^N \sum_{\gamma=\{x,y,z\}} \sum_{k=0}^{\tau/\Delta t-1} \frac{(x_{t+(k+1)\Delta t}^{[\alpha,\gamma]} - x_{t+k\Delta t}^{[\alpha,\gamma]}) (\nabla_{[\alpha,\gamma]} V_{target}(\mathbf{x}_{t+k\Delta t}) - \nabla_{[\alpha,\gamma]} V_{sim}(\mathbf{x}_{t+k\Delta t}))}{k_B T \xi (1 + \exp(-\xi \Delta t)) \Delta t} \right) \\
& \cdot \exp \left(\sum_{\alpha=1}^N \sum_{\gamma=\{x,y,z\}} \sum_{k=0}^{\tau/\Delta t-1} \frac{v_{t+k\Delta t}^{[\alpha,\gamma]} (\nabla_{[\alpha,\gamma]} V_{target}(\mathbf{x}_{t+k\Delta t}) - \nabla_{[\alpha,\gamma]} V_{sim}(\mathbf{x}_{t+k\Delta t}))}{k_B T \xi (1 + \exp(\xi \Delta t))} \right) \\
& \cdot \exp \left(- \frac{\exp(\xi \Delta t) - 1}{\exp(\xi \Delta t) + 1} \cdot \sum_{\alpha=1}^N \sum_{\gamma=\{x,y,z\}} \sum_{k=0}^{\tau/\Delta t-1} \frac{((\nabla_{[\alpha,\gamma]} V_{target}(\mathbf{x}_{t+k\Delta t}))^2 - (\nabla_{[\alpha,\gamma]} V_{sim}(\mathbf{x}_{t+k\Delta t}))^2)}{2k_B T \xi^2 m_{[\alpha]}} \right), \tag{22}
\end{aligned}$$

where the summation $\alpha = 1 \dots N$ is carried over the N particles in the system, the summation $\gamma = \{x, y, z\}$ over the three spatial dimensions, $x_t^{[\alpha,\gamma]}$ and $v_t^{[\alpha,\gamma]}$ represent the γ components of the position and velocity of particle α , $\nabla_{[\alpha,\gamma]}$ is the partial derivative taken with respect to the γ spatial component of particle α , $m_{[\alpha]}$ is the mass of particle α , ξ is the Langevin friction coefficient, Δt the integration time step, and the lag time τ is restricted to be an integer multiple of Δt . This equation may be equivalently rewritten in terms of the bias potential U_{bias} and the random number sequence that generated the Langevin integration path $\{\eta_{t+k\Delta t}\}_{k=0}^{\tau/\Delta t-1}$ to present a simpler and potentially less computationally expensive expression to evaluate (c.f., Eqn. 27 in Ref. ⁴³), ^{43,47,48}

$$\begin{aligned}
M(\boldsymbol{\omega}_{t \rightarrow t+\tau}) = & \exp \left(- \frac{1 - \exp(-\xi \Delta t)}{\sqrt{1 - \exp(-2\xi \Delta t)}} \cdot \sum_{\alpha=1}^N \sum_{\gamma=\{x,y,z\}} \sum_{k=0}^{\tau/\Delta t-1} \frac{\nabla_{[\alpha,\gamma]} U_{bias}(\mathbf{x}_{t+k\Delta t}) \eta_{t+k\Delta t}}{\sqrt{k_B T \xi^2 m_{[\alpha]}}} \right) \\
& \cdot \exp \left(- \frac{(1 - \exp(-\xi \Delta t))^2}{1 - \exp(-2\xi \Delta t)} \cdot \sum_{\alpha=1}^N \sum_{\gamma=\{x,y,z\}} \sum_{k=0}^{\tau/\Delta t-1} \frac{(\nabla_{[\alpha,\gamma]} U_{bias}(\mathbf{x}_{t+k\Delta t}))^2}{2k_B T \xi^2 m_{[\alpha]}} \right). \tag{23}
\end{aligned}$$

Having defined the reweighting factors, we return to Eqn. 19 and estimate it empirically

from trajectories collected under the simulation potential $V_{sim} = V_{target} - U_{bias}$ from which we compile a batch of $i = 1 \dots L$ training samples comprising pairs of states $(\mathbf{x}_t^{(i)}, \mathbf{x}_{t+\tau}^{(i)})$ separated within a simulation trajectory by a lag time τ and connected by paths $\omega_{t \rightarrow t+\tau}^{(i)} = ((\mathbf{x}_t^{(i)}, \mathbf{v}_t^{(i)}), (\mathbf{x}_{t+\Delta t}^{(i)}, \mathbf{v}_{t+\Delta t}^{(i)}), \dots, (\mathbf{x}_{t+\tau-\Delta t}^{(i)}, \mathbf{v}_{t+\tau-\Delta t}^{(i)}), (\mathbf{x}_{t+\tau}^{(i)}, \mathbf{v}_{t+\tau}^{(i)}))$,

$$C_{jk}^\tau \approx \frac{1}{L} \sum_{i=1}^L h(\mathbf{x}_t^{(i)}) M(\omega_{t \rightarrow t+\tau}^{(i)}) \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_{t+\tau}^{(i)}). \quad (24)$$

As a corollary, the matrix elements S_{jk} of the instantaneous correlation matrix can be estimated over the same data as,

$$S_{jk} \approx \frac{1}{L} \sum_{i=1}^L h(\mathbf{x}_t^{(i)}) M(\omega_{t \rightarrow t+\tau}^{(i)}) \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_t^{(i)}). \quad (25)$$

The matrices C^τ and S are subsequently used to solve for the basis function expansion coefficients via the generalized eigenvalue problem defined in Eqn. 11. The thermodynamic reweighting factor $h(\mathbf{x}_t)$ (Eqn. 20) possesses a term $\left(\frac{Z_{sim}}{Z_{target}}\right)$ that serves as a multiplicative constant of each element of C^τ and S that is independent of state $\mathbf{x}_t$ and path $\omega_{t \rightarrow t+\tau}^{(i)}$ and may therefore be canceled out on each side of the generalized eigenvalue problem. As such, this partition function ratio never need be computed and can be dropped. Defining,

$$g(\mathbf{x}_t) = \exp(-\beta U_{bias}(\mathbf{x}_t)), \quad (26)$$

we can rewrite $h(\mathbf{x}_t) = \left(\frac{Z_{sim}}{Z_{target}}\right) g(\mathbf{x}_t)$ and formally drop the partition function ratio by replacing $h(\mathbf{x}_t)$ with $g(\mathbf{x}_t)$ in Eqns. 24 and 25 to yield,

$$C_{jk}^\tau \approx \frac{1}{L} \sum_{i=1}^L g(\mathbf{x}_t^{(i)}) M(\omega_{t \rightarrow t+\tau}^{(i)}) \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_{t+\tau}^{(i)}), \quad (27)$$

$$S_{jk} \approx \frac{1}{L} \sum_{i=1}^L g(\mathbf{x}_t^{(i)}) M(\omega_{t \rightarrow t+\tau}^{(i)}) \chi_j(\mathbf{x}_t^{(i)}) \chi_k(\mathbf{x}_t^{(i)}). \quad (28)$$

Finally, we define a combined thermodynamic and dynamic reweighting factor,

$$W(\boldsymbol{\omega}_{t \rightarrow t+\tau}^{(i)}) = g(\mathbf{x}_t^{(i)}) \times M(\boldsymbol{\omega}_{t \rightarrow t+\tau}^{(i)}), \quad (29)$$

and explicitly replace the basis functions $\{\chi_i(\mathbf{x})\}$ with the SRV neural network outputs $\chi_j(\mathbf{x}) = f_j^\theta(\mathbf{x})$ to obtain our final expressions for C_{jk}^τ and S_{jk} that we use in practice,

$$C_{jk}^\tau \approx \frac{1}{L} \sum_{i=1}^L W(\boldsymbol{\omega}_{t \rightarrow t+\tau}^{(i)}) f_j^\theta(\mathbf{x}_t^{(i)}) f_k^\theta(\mathbf{x}_{t+\tau}^{(i)}), \quad (30)$$

$$S_{jk} \approx \frac{1}{L} \sum_{i=1}^L W(\boldsymbol{\omega}_{t \rightarrow t+\tau}^{(i)}) f_j^\theta(\mathbf{x}_t^{(i)}) f_k^\theta(\mathbf{x}_t^{(i)}). \quad (31)$$

In the absence of bias (i.e., $U_{bias}(\mathbf{x}) = 0 \Rightarrow V_{sim}(\mathbf{x}) = V_{target}(\mathbf{x})$) the reweighting factors in Eqns. 26 and 23 both reduce to unity (i.e., $g(\mathbf{x}_t) = M(\boldsymbol{\omega}_{t \rightarrow t+\tau}) = 1 \Rightarrow W(\boldsymbol{\omega}_{t \rightarrow t+\tau}) = 1$) and the expressions for C_{jk}^τ and S_{jk} for biased simulation data (Eqns. 30 and 31) reduce to those for unbiased data (Eqns. 14 and 15) as they should.

Having adapted the SRV framework incorporating dynamical reweighting, training can now proceed to approximate the slow modes of a molecular system from short, discontinuous, and biased trajectories. Since the learned modes will correspond only to the slowest modes present within the training data, the true slow modes for systems containing high free energy barriers and rugged free energy landscapes will not be well approximated within unbiased simulation trajectories that have a propensity to be kinetically trapped. This is the primary motivation for the GREST approach presented herein that performs interleaved rounds of slow CV discovery and enhanced sampling until convergence is achieved in the explored phase space and the learned CVs. A Python package containing a PyTorch⁸⁵ implementation of SRVs with capabilities for training with biased simulation data incorporating these reweighting factors is publicly available via GitHub at <https://github.com/andrewlferguson/snrsv>.

2.5 GREST adaptive sampling workflow

The Girsanov Reweighting Enhanced Sampling Technique (GREST) is an adaptive sampling workflow combines slow CV discovery with SRVs and Girsanov path reweighting to perform simultaneous learning of the slow modes of a molecular system and efficient exploration of configurational space. This is accomplished by interleaving rounds of slow CV estimation with path reweighted SRVs and enhanced sampling along the discovered slow modes. Our approach is procedurally similar to MESA³³ but is designed to discover and accelerate slow rather than high variance CVs. It employs the path reweighting developed by Keller and co-workers to estimate dynamical observables from biased simulation trajectories,^{47,48} but uses this approach to iteratively discover slow CVs as opposed to construct MSMs from biased simulations conducted in pre-defined CVs that was the focus of work by both Keller and co-workers^{47,48} and Weber and Pande.⁴⁵ The first iterative round of GREST shares methodological similarities with the approach reported in Ref.⁴⁰ but we apply both dynamical and thermodynamic corrections to rigorously reweight our path ensemble estimators of the slow CVs and perform successive rounds of discovery and biasing to converge the CVs and exploration of configurational phase space. A schematic illustration of the four-step iterative workflow is presented in Fig. 1 and an algorithmic overview is shown in Algorithm 1.

Step 0: Generation of initial training data. Conduct an initial simulation to gather trajectory data for the first round of slow CV estimation. Slow CV estimation with SRVs is ultimately a data-driven technique, and therefore generally benefits from training on larger datasets that explore larger volumes of the configuration space. As this initial simulation serves as only a starting point for the adaptive sampling protocol, the length of the trajectory can judiciously be chosen by considering available computational resources. For larger, complex systems with rugged free energy surfaces susceptible to kinetic trapping, an unbiased simulation may be insufficient to provide adequate exploration of local configurational space and a biased simulation in intuited CVs based on expert knowledge may be conducted to provide more extensive configurational space coverage within this initial training dataset.




Figure 1: Schematic illustration of four-step iterative Girsanov Reweighting Enhanced Sampling Technique (GREST) workflow. **(a)** Initial training data is generated by unbiased or biased simulation data accelerated along intuitive CVs. **(b)** Slow CVs are estimated from the training data using Girsanov path reweighted State-free Reversible VAMPnets (SRVs). **(c)** Enhanced sampling calculations are performed within the learned slow CVs using a CV biasing technique such as metadynamics. **(d)** Convergence of the iterative cycle is assessed by evaluating the stability of the learned CVs in successive iterations.

Step 1: Slow CV estimation using SRVs. Train path reweighted SRV models using unbiased or biased simulation data to estimate slow CVs. The SRV learns a mapping from configuration space coordinates $\mathbf{x}$ to approximate slow eigenfunctions of the transfer operator $\mathbf{x} \rightarrow \{\phi_i(\mathbf{x})\}$. Since the dimensionality of the slow subspace may change over the course of the GREST iterative cycles as a result of the region and volume of configurational phase space that has been explored by the simulation trajectories, we recommend first training SRVs with a relatively high number of CVs (e.g., $n \sim 10$) to attempt to discern a spectral gap in the slow mode eigenvalue spectrum and make a principled inference of the dimensionality

Algorithm 1: Girsanov Reweighting Enhanced Sampling Technique (GREST) work-flow

```

1  $\mathbf{x} \leftarrow$  Initial short simulation // Step 0
  /* From an unbiased simulation or basing along intuited CVs */
2  $k \leftarrow 0$  // GREST round # counter
3 notConverged  $\leftarrow$  True
4 while notConverged do
5    $\{\psi_i^{(k)}(\mathbf{x})\} \leftarrow$  SRV( $\mathbf{x}$ ) // Step 1
6   /* Estimate slow CVs from simulation data using SRVs with Girsanov
     Reweighting if applicable */
7    $\mathbf{x} \leftarrow$  EnhancedSampling( $\{\psi_i^{(k)}(\mathbf{x})\}$ ) // Step 2
8   /* Perform enhanced sampling in the learned slow CVs to gather
     biased simulation data */
9   if ConvergenceCriterion( $\{\psi_i^{(k)}(\mathbf{x})\}, \{\psi_i^{(k-1)}(\mathbf{x})\}, \dots$ ) then
10    notConverged  $\leftarrow$  False // Step 3
    /* Analyze successive rounds of learned slow modes to determine
       convergence */
11     $k \leftarrow k + 1$  // Increment round # counter

```

of the slow subspace. The only exception to this recommendation is for the first round of sampling where exploration may be poor and we have found that it can be useful to simply choose a 2D latent space for the purposes of driving the first round of biased sampling.

Step 2: Enhanced sampling along slow CVs. Perform enhanced sampling within the SRV-determined CVs. The slow CVs learned by the SRV can be passed to any CV enhanced sampling technique and used to drive accelerated sampling in an off-the-shelf manner. For the molecular systems in this work, we choose to employ well-tempered metadynamics (WTMetaD).^{18,86} Converging WTMetaD can become computationally burdensome when sampling in more than two CVs, and in this case we recommend well-tempered parallel bias metadynamics (PBMetaD) as an approach that scales well to large numbers of CVs.^{87,88} The accumulation of bias throughout the metadynamics run can lead to large variances in the estimator of the dynamical path weights⁴⁸ and produce numerical instabilities during SRV training. As suggested in Refs.,^{45,48} we instead converge the metadynamics bias U_{meta} and construct a potential $V_{sim} = V_{target} - U_{bias} = V_{target} - \gamma U_{meta}$ under which we conduct

an additional round of metadynamics enhanced sampling in which we do not update the terminal bias potential and modulate its strength by a factor $\gamma = [0, 1]$. This constitutes enhanced sampling under an attenuated bias potential that eliminates the need to track the time varying nature of the bias and also enables us to judiciously tune the aggressiveness of the applied bias via the attenuation factor γ . Consistent with previous observations,⁴⁸ we have found that selecting an attenuation factor of $\gamma \approx 0.05 - 0.2$ tends to provide a good balance between enhancing sampling of configurational space while moderating the magnitude of the path weights to stabilize low-variance numerical estimation of path ensemble averages. In this work, we adopt a relatively conservative attenuation factor of $\gamma = 0.05$. This attenuation of the applied bias does limit the barrier crossing accelerations achievable by the technique. The exponential nature of barrier crossing dynamics, however, means that significant accelerations can be achieved even by relatively modest reductions in barrier heights. For example, assuming a simple Arrhenius dependence $k = A \exp\left(-\frac{E_a}{k_B T}\right)$, where k is the rate constant, A is the pre-exponential factor, E_a is the activation energy, k_B is the Boltzmann constant, and T is the temperature, a modest 5% decrease of a $E_a=10\ k_B T$ barrier leads to a 65% elevation of the rate constant.

Step 3: Evaluating convergence. Assess the convergence of our adaptive sampling workflow by determining if the learned SRV CVs have stabilized. Using all accumulated enhanced sampling data, we evaluate the Pearson correlation coefficient of the learned SRV CVs between consecutive enhanced sampling rounds. Plateauing of this metric suggests the slow modes identified with SRVs trained on our enhanced sampling data have approximately converged. Based on this metric we can evaluate if CV learning has halted and terminate our workflow. Otherwise, we perform another iterative round of GREST.

2.6 Simulation settings and parameters

2.6.1 One-dimensional four-well potential

Simulations under (underdamped) Langevin dynamics are performed for a single particle moving on the 1D 4-well potential $V_{4\text{-well}}(x) = 2x^8 + 1.6 \exp(-80x^2) + 0.4 \exp(-80(x - 0.5)^2) + \exp(-40(x + 0.5)^2)$ introduced by Schwantes and Pande.³⁸ The ISP scheme⁴³ – a full-step adaptation of the Langevin leapfrog algorithm developed by Izaguirre, Sweet, and Pande⁵⁴ – is used to integrate the dynamics forward in time. Simulations are performed using a time step of 0.01 s for a 1 kg particle with a friction coefficient of $\xi = 1 \text{ s}^{-1}$ at a temperature of 60 K. Configurations are saved for analysis at a period of 0.05 s.

2.6.2 Alanine dipeptide

Simulations of alanine dipeptide (ADP) are carried out at a temperature of 300 K in implicit solvent using OpenMM⁸⁴ with the AMBER99SB-ILDN⁸⁹ force field and the GBSA-OBC⁹⁰ solvation model. Nonbonded interactions were cutoff at distances beyond 1 nm. The system is numerically integrated under (underdamped) Langevin dynamics using the ISP scheme⁵⁴ for which Keller and co-workers developed analytical expressions and efficient numerical calculation schemes for the Girsanov path weights.^{43,47,48} A simulation time step of 2 fs is used with a friction coefficient of $\xi = 10 \text{ ps}^{-1}$. Configurations are saved for analysis at a period of 0.1 ps.

2.6.3 WLALL

Simulations of the Trp-Leu-Ala-Leu-Leu (WLALL) pentapeptide are conducted at a temperature of 300 K using OpenMM⁸⁴ and the CHARMM22^{91,92} force field with the GBSA-OBC⁹⁰ implicit solvent model. All non-bonded interactions are cutoff at 1 nm. The system is numerically integrated under (underdamped) Langevin dynamics using the ISP scheme⁵⁴ and path weights calculated using the approach developed by Keller and co-workers.^{43,47,48} A

simulation time step of 2 fs is used with a friction coefficient of $\xi = 10 \text{ ps}^{-1}$. Configurations are saved for analysis at a period of 0.1 ps.

3 Results

3.1 One-dimensional four-well potential

As a first illustrative toy example, we consider a single particle moving on the 1D 4-well potential (Fig. 2a).³⁸ We model single particle diffusion over this landscape under (under-damped) Langevin dynamics. This system is sufficiently simple that the four wells and three barrier hopping transitions can be comprehensively explored under long unbiased simulations to provide a gold standard baseline against which to benchmark GREST.^{38,39} We represent the system state of this simple 1D system to our path reweighted SRVs by simply furnishing the 1D coordinate $\mathbf{x} = x \in \mathbb{R}^1$. In contrast to linear algorithms such as TICA,^{35–37} the SRV deep learning architecture can learn nonlinear transformations of the $\mathbf{x}$ coordinate to learn a hierarchy of slowly relaxing CVs. Full details of the simulations are provided in the Supporting Information.

Pre-defined bias. We first demonstrate the capacity of the dynamically reweighted SRVs to learn the correct slow modes and associated time scales under a pre-defined bias. This can be conceived of as a first round of a GREST campaign in which the user applies an intuitive biasing potential to accelerate sampling of the system prior to any data-driven learning of the slow modes. The 4-well potential serves as the target potential under which we wish to estimate the slow modes $V_{target} = V_{4-well} = 2x^8 + 1.6 \exp(-80x^2) + 0.4 \exp(-80(x - 0.5)^2) + \exp(-40(x + 0.5)^2)$ but we perform simulations under a 2-well potential $V_{sim} = V_{2-well} = 2x^8 + 1.75 \exp(-80x^2)$ (Fig. 2a). Formally, since the simulation and target potentials are related via the applied bias as $V_{sim} = V_{target} - U_{bias}$, this amounts to adopting $U_{bias} = V_{4-well} - V_{2-well}$ as our biasing potential. The effect of this bias is to flatten the two smaller barriers to leave only the single high barrier at the origin intact. We




Figure 2: Application of GREST to the 1D 4-well potential as an illustrative toy example. **(Left column.)** Application of GREST under pre-defined bias. **(a)** The target 4-well potential V_{target} (blue) is shown alongside the 2-well potential under which the biased simulation is performed V_{sim} . The challenge is for path reweighted SRVs to correctly recover the slowly relaxing modes and associated time scales of the 4-well system in simulations conducted under the 2-well potential. **(b-d)** Comparison of the leading three leading slow eigenfunctions $\{\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \phi_3(\mathbf{x})\}$ recovered by path reweighted SRVs from long unbiased reference data V_{target} (green), alongside eigenfunctions learned using biased simulation data collected under V_{sim} to which we apply path weights comprising dynamical and thermodynamic corrections (orange), only thermodynamic corrections (red), and no reweighting corrections (blue). Eigenvectors are defined only up to scaling and sign so are presented under unit normalization and, if necessary, inversion of sign to align them with the unbiased reference data. **(e)** Associated implied time scales $t_i = -\tau/\ln(\lambda_i)$ for the SRV eigenfunctions in panels **(b-e)** with error bars showing the standard deviation over 10 independently trained models with different random seeds. **(Right column.)** Application of GREST under a learned bias. **(f)** A biasing potential U_{bias} is constructed by projecting the unbiased reference data into the leading slow mode $\phi_1(\mathbf{x})$ learned by a path reweighted SRV fit to the unbiased 4-well simulation trajectories. Application of the bias to the target potential results in a simulation potential $V_{sim} = V_{target} - U_{bias}$ that attenuates the highest central potential energy barrier by ~ 0.5 kJ/mol and the other potential energy barriers by ~ 0.25 kJ/mol. **(g-i)** Comparison between the reference eigenfunctions and those estimated from biased simulation data under V_{sim} in panel **(f)** to which we apply various path weights. **(j)** Associated implied time scales $t_i = -\tau/\ln(\lambda_i)$ for the SRV eigenfunctions in panels **(g-i)** with error bars showing the standard deviation over 10 independently trained models with different random seeds.

therefore challenge our path reweighted SRVs to recover the correct slow modes and time scales of the 4-well system under simulations conducted under the 2-well potential.

We present in Fig. 2b-d (green) approximations to the three leading slow eigenfunctions $\{\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \phi_3(\mathbf{x})\}$ learned using SRVs from long unbiased reference trajectories collected under $V_{target} = V_{4-well}$. The leading eigenfunction $\phi_1(\mathbf{x})$ changes sign at $\mathbf{x} = 0$ m and corresponds to the slowest occurring transitions over the largest potential barrier. The next eigenfunction $\phi_2(\mathbf{x})$ is approximately zero in the right half space and changes sign at $\mathbf{x} = (-0.5)$ m to characterize transitions over the next highest potential barrier. Finally, $\phi_3(\mathbf{x})$ is approximately zero in the left half space and changes sign at $\mathbf{x} = +0.5$ m to characterize transitions over the lowest potential barrier. The associated implied time scales for these relaxations are presented in Fig. 2e.

We also present in Fig. 2b-e the eigenfunctions and implied time scales learned by path reweighted SRVs from biased simulation data collected under $V_{sim} = V_{2-well}$ under different choices of path weights: dynamical and thermodynamic corrections (orange), only thermodynamic corrections (red), and no reweighting corrections (blue). The 2-well landscape differs from the 4-well landscape largely by flattening of the two lower potential barriers and leaving the largest central potential barrier largely intact. Recognizing that the slowest relaxing mode is associated with the same dynamical process of transitions over this central barrier, we anticipate good agreement between the leading slow mode estimated under the 4-well and 2-well simulation trajectories. Inspection of Fig. 2b shows that this is indeed the case: the estimate of $\phi_1(\mathbf{x})$ from unbiased reference data collected under V_{4-well} (green) is in good agreement in the uncorrected estimate of $\phi_1(\mathbf{x})$ computed from simulation trajectories collected under V_{2-well} (blue). Application of thermodynamic only ($W(\omega_{t \rightarrow t+\tau}) = g(\mathbf{x}_t)$, red) or thermodynamic and dynamical ($W(\omega_{t \rightarrow t+\tau}) = g(\mathbf{x}_t) \times M(\omega_{t \rightarrow t+\tau})$, orange) reweighting corrections within the SRV estimation (cf. Eqn. 29) therefore does little to improve the agreement with the $\phi_1(\mathbf{x})$ unbiased reference. Reweighting is, however, critical in accurate estimation of the correct implied time scale associated with this mode (Fig. 2e, left). The largest error in predicting the reference time scale of the leading mode (green) occurs when no reweighting corrections are applied (blue). Incorporating thermodynamic corrections alone improves agreement (red), but only by incorporating both the thermodynamic and dynamical corrections do we accurately recover the true time scale (orange).

The 2-well potential lacks potential barriers at $\mathbf{x} = (-0.5)$ m and $\mathbf{x} = +0.5$ m meaning that biased simulations conducted under $V_{sim} = V_{2-well}$ do not contain any signatures of relaxations associated with these barrier crossings and these modes are completely obfuscated due to obliteration of these barriers by the applied bias (Fig. 2b-e, blue). Nevertheless, $\phi_2(\mathbf{x})$ and $\phi_3(\mathbf{x})$ and the associated implied time scales are accurately recovered when incorporating both thermodynamic and dynamical reweighting corrections into SRV learning (orange). Interestingly, the thermodynamic correction alone does very little to improve esti-

mation of the true slow eigenfunctions and implied time scales, cautioning against omission of the dynamical path weights in estimating slow CVs.⁴⁰ This example illustrates that path reweighting can accurately learn the slow dynamical modes and time scales of a dynamical system from biased simulation trajectories.

Learned bias. We now test the capacity of the learned slow CVs to support an accelerating biasing potential and promote enhanced sampling of configurational phase space. This can be approximately conceived as the second round of a GREST campaign in which sampling is enhanced by accelerating biases in learned slow modes. We construct the bias potential by assuming that the projection of the unbiased reference simulation into the leading slow mode $P(\phi_1(\mathbf{x}))$ is Boltzmann distributed and may therefore be inverted to construct an effective potential $F(\phi_1(\mathbf{x})) = -k_B T \log(P(\phi_1(\mathbf{x})))$. This represents a good candidate biasing potential as it flattens out the portion of the potential associated with dynamical transitions in $\phi_1(\mathbf{x})$ and can therefore help accelerate sampling of these events. This leads us to define the biasing potential as the functional composition $U_{bias} = \mathcal{F}(\phi_1(\mathbf{x}))$, where $\mathcal{F}(\phi_1)$ is an analytical fit to $F(\phi_1)$ that we find to be well described by a super-Gaussian $F(\phi_1) \approx \mathcal{F}(\phi_1) = A \exp(-(\frac{(\phi_1 - x_0)^2}{2\sigma^2})^P)$ with $A = 3.1$ kJ/mol, $x_0 = 0.049$ m, $\sigma^2 = 0.058$ m², and $P = 2$ (see Supporting Information). This functional fit also admits continuous first derivatives required to compute the biasing forces on the particle through the chain rule, $\mathbf{f}_{bias} = -\frac{dU_{bias}}{d\mathbf{x}} = -\frac{d\mathcal{F}}{d\phi_1} \frac{d\phi_1}{d\mathbf{x}}$, where $\frac{d\mathcal{F}}{d\phi_1}$ is available analytically and $\frac{d\phi_1}{d\mathbf{x}}$ from automatic differentiation through the trained SRV neural network. In Fig. 2g, we illustrate $V_{target} = V_{4-well}$, U_{bias} , and $V_{sim} = V_{target} - \gamma U_{bias} = V_{4-well} - 0.2U_{bias}$, where we have chosen to set the attenuation factor to $\gamma = 0.2$ (cf. Section 2.5). The biasing potential depresses the height of the central highest potential barrier by ~ 0.5 kJ/mol, resulting in an observed $\sim 150\%$ increase in the number of $(\mathbf{x} < 0) \leftrightarrow (\mathbf{x} > 0)$ transitions. It also suppresses the two auxiliary barriers at $\mathbf{x} = -0.5$ m and $\mathbf{x} = 0.5$ m by ~ 0.25 kJ/mol, resulting in modest increases in the number of $(\mathbf{x} < -0.5) \leftrightarrow (\mathbf{x} > -0.5)$ and $(\mathbf{x} < 0.5) \leftrightarrow (\mathbf{x} > 0.5)$ transitions by $\sim 20\%$ and $\sim 14\%$, respectively.

We present in Fig. 2g-j the reference and learned eigenfunction approximations under different reweighting schemes. Due to the relatively small perturbation to the target potential, we see qualitative agreement with the unbiased reference eigenfunctions (green) even in the absence of any corrections (blue) or thermodynamic-only corrections (red), but the full thermodynamic and dynamical corrections (orange) are required to achieve quantitative agreement. The second and third implied time scales are in good agreement for all reweighting schemes since the barriers associated with these dynamical transitions are only moderately affected by the application of the bias, but we see that both the thermodynamic and dynamical corrections are imperative for accurate estimation of the leading time scale associated with transitions over the highest central potential barrier. This example demonstrates that an SRV-learned slow mode can be used to effectively accelerate slow dynamical transitions in a biased simulation from which path reweighted SRVs can accurately recover the true slow modes.

3.2 Alanine dipeptide

We now demonstrate our approach in an application to the widely studied biomolecular system alanine dipeptide (ADP). This system represents the “hydrogen atom of molecular simulation”⁹³ and we challenge GREST to simultaneously estimate the slow CVs governing its biomolecular dynamics and recover the underlying molecular free energy landscape without any prior knowledge of this system. Full details of the application of GREST to this system including all path reweighted SRV training settings and metadynamics hyperparameters are provided in the Supporting Information.

Unbiased reference benchmark. We first perform a long unbiased simulation to serve as a baseline reference for the thermodynamics and kinetics of our ADP system against which to benchmark the performance of GREST. We present in Fig. 3a-i the ADP free energy surface as a ϕ - ψ projection into the backbone dihedral angles that are known to be good CVs with which to differentiate the metastable conformational states of the system.^{61,94} To ensure

comprehensive sampling of configurational space, the reference simulation data was gathered from $8 \times 0.5 \mu\text{s}$ simulations initialized from configurations nearest k-means centroids identified by k-means clustering the ϕ - ψ projection of an initial 250 ns unbiased simulation. The slowest relaxing implied time scale for this system is ~ 1.4 ns, so these reference trajectories provide excellent sampling of all relevant transitions in this ADP system. We train an SRV over these unbiased simulation data in which we represent the instantaneous state of the ADP system through all $\binom{22}{2} = 231$ pairwise atomic distances to define a translationally and rotationally invariant featurization $\mathbf{x} \in \mathbb{R}^{231}$ of the system.

The slowest process captured by the first SRV mode $\phi_1(\mathbf{x})$ characterizes transitions between the triplet of metastable states $\{C_5, P_{II}, \alpha_R\}$ where $\phi < 0$ and the doublet of metastable states $\{\alpha_L, C_{7ax}\}$ where $\phi > 0$ (Fig. 3a-ii). The second slow mode $\phi_2(\mathbf{x})$ largely subpartitions the $\{C_5, P_{II}, \alpha_R\}$ triplet by characterizing transitions in ψ between $\{C_5, P_{II}\}$ and α_R (Fig. 3a-iii). Throughout our adaptive sampling workflow we constrain the path reweighted SRV to return two slow CVs since it is known that ADP possesses a 2D intrinsic manifold lying on the surface of a flat torus that is well parameterized by the ϕ and ψ backbone dihedrals.^{33,95–98} In general, one would determine the appropriate number of CVs to retain by searching for a gap in the SRV eigenvalue spectrum at each round of GREST. Our goal is to learn to recapitulate these kinetics by biasing along dynamically estimated slow modes to visit relevant macrostates and efficiently sample state-to-state transitions.

Application of GREST. We commence Round 0 of our adaptive sampling workflow by running an initial unbiased 5 ns simulation with the resultant free energy profile shown in (Fig. 3b-i). This initial simulation only samples the triplet of metastable states C_5 , P_{II} , and α_R residing in the low- ϕ /high- ψ superbasin and fails to experience any transitions to the $\phi > 0$ region. In the next step, we use this initial data to estimate the leading two SRV slow modes (Fig. 3b-ii,iii). The paucity of training data in this initial round contributes to learning a pair of seemingly degenerate slow modes that only capture transitions in ϕ . Equipped with these learned slow modes and initial simulation data, we perform k-means clustering in the




Figure 3: Application of GREST to ADP. **(a-i)** Reference free energy profile constructed from 4 μ s of unbiased simulation data projected into the ϕ - ψ backbone angles. The local free energy minima are annotated with their associated macrostate labels: C_5 , P_{II} , α_R , α_L , and C_{7ax} . **(a-ii,iii)** Leading two slow CVs $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$ identified by application of SRVs to the unbiased reference trajectories. $\phi_1(\mathbf{x})$ characterizes transitions in the ϕ dihedral between $\{C_5, P_{II}, \alpha_R\}$ and $\{\alpha_L, C_{7ax}\}$. $\phi_2(\mathbf{x})$ primarily characterizes transitions in the ψ dihedral between $\{C_5, P_{II}\}$ and α_R . Rows **(b)-(g)** index successive rounds of GREST, commencing with the initial unbiased Round 0 in **(b)** and ending with the terminal adaptive sampling Round 5 in **(g)**. Columns **(i)-(iv)** index various facets of each round of GREST: **(i)** illustrates the unbiased free energy landscape estimated by thermodynamic reweighting of the enhanced sampling simulations performed under the fixed terminal metadynamics bias projected into the ϕ - ψ backbone angles, **(ii)** and **(iii)** illustrate heatmaps of the path reweighted SRV slow modes $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$ learned from the enhanced sampling trajectories in the terminal metadynamics bias potential projected into ϕ - ψ space, and **(iv)** illustrates the phase space sampling density projected in ϕ - ψ space from well-tempered metadynamics simulations performed in the two learned slow modes $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$ illustrated in the preceding columns. The GREST iterations are terminated when the CVs learned in successive rounds stabilize, indicating no additional enhancement in configurational and dynamical exploration of phase space. Convergence of GREST at Round 4 means that additional enhanced sampling calculations are not conducted in Round 5.

feature space of the 231 pairwise atomic distances to identify configurations nearest 25 k-means centroids that are used to initialize parallel walker metadynamics along our learned SRV CVs in the subsequent step of GREST. Each walker is simulated for 1 ns yielding an aggregate 25 ns of metadynamics simulation data. In Fig. 3b-iv we show the sampling density from this metadynamics simulation as a visualization of the phase space explored due to accelerating along the learned SRV slow modes $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$. The resultant phase space exploration along these learned CVs reveals that biasing has largely driven transitions in ψ across the free energy barrier bridging metastable states P_{II} and α_R . To complete Round 0, we perform enhanced sampling simulations using a time invariant and attenuated version of the terminal metadynamics bias potential, which we reduce in magnitude by an attenuation factor of $\gamma = 0.05$ to control the numerical instability associated with very large and very small magnitude reweighting factors during the subsequent SRV training.^{45,48} These enhanced sampling simulations are similarly executed in parallel with 25×1 ns simulations initialized from configurations identified with the same k-means clustering protocol as in the

previous step but applied to these 25 ns of metadynamics simulation trajectories.

In Fig. 3c, we illustrate Round 1 of our adaptive sampling protocol which proceeds analogously to the initial round except that the training data now comprises enhanced sampling trajectories biased in the learned slow CVs from the previous round under the $\gamma = 0.05$ attenuation factor. Although the applied biasing forces are relatively modest, the reweighted free energy surfaces in Fig. 3c-i demonstrate that the biased trajectories experience transitions in ϕ and now visit metastable states α_L and C_{7ax} that were absent in the initial unbiased data. Training path reweighted SRVs over these trajectories with thermodynamic and dynamical corrections exposes two leading slow modes illustrated in Fig. 3c-ii,iii that are now in much better agreement with those extracted from the long unbiased reference data. (We recall that the sign of the learned eigenfunctions is immaterial.) We also see the subsequent WTMetaD simulations performed along these SRV CVs by initializing 25×1 ns walkers by k-means clustering in the feature space of 231 pairwise atomic distances drives more comprehensive sampling of configurational phase space (Fig. 3c-iv). Visual analysis reveals that this time the enhanced sampling has driven transits over the free energy barriers connecting P_{II} and α_R , P_{II} and α_L , α_R and C_{7ax} , and α_L and C_{7ax} . To complete Round 1, we perform a final enhanced sampling simulation applying the terminal WTMetaD potential under a $\gamma = 0.05$ attenuation factor using 25×1 ns walkers initialized by k-means clustering in the feature space of 231 pairwise atomic distances, and pass this forwards as the biased trajectory training data for Round 2.

We iterate this process to complete Rounds 2-5 of GREST that are illustrated in Fig. 3d-g. Convergence is assessed by tracking the round-to-round change in the learned slow SRV modes. Stabilization of the learned CVs between successive GREST rounds indicates that no important new dynamical modes or regions of configurational phase space are being explored under additional rounds of biased sampling. We quantify convergence by examining the Pearson correlation coefficient of SRV slow modes learned in consecutive rounds. Importantly, we maximize the correlation coefficient under rigid rotations of the subspace spanned by the

slow modes in successive rounds to account for possible rotational symmetries in which the same subspace could be spanned by an equivalent rotation of the learned slow eigenfunctions. Mathematically, this corresponds to identifying the affine transformation that maximizes the correlation coefficient. Convergence is defined by stabilization in the number n of recovered slow modes and round-to-round Pearson correlation coefficients $\{\rho_i > 0.8\}_{i=1}^n$ for all modes. In the present case the former criterion is satisfied by default since we fix the dimensionality of the SRV bottleneck layer to 2D. Consistent with expectations from inspection of the free energy surfaces and slow CVs in Fig. 3, we illustrate in Fig. 4 that the latter criterion is quickly satisfied after only Round 2, but we continue for three additional rounds for the purposes of illustration.




Figure 4: Convergence assessment of GREST applied to ADP by measuring the Pearson correlation coefficient between SRV slow CVs discovered in successive rounds. Stabilization of the CVs to $\rho > 0.8$ indicates that no important new dynamical modes or regions of configurational phase space are being explored under additional rounds of biased sampling.

The excellent agreement between the free energy surface and slow modes of the terminal round (Fig. 3g) and the unbiased reference data (Fig. 3a) indicate that the GREST iter-

ations successfully recovered the true slow dynamical modes and used these to efficiently enhance sampling of configurational phase space. While the ADP system does represent a demonstrative and instructive molecular example for which the ground truth is well understood, its small size and efficiency to simulate make it a relatively simple test of the GREST methodology. In particular, expending the same computational effort expended within the GREST pipeline (250 ns of cumulative simulation time) in unbiased simulations would accurately recover the free energy surface by brute force and enable SRV discovery of the slow modes without path reweighting. Moreover, the relatively low free energy barriers mean that it is difficult to assess the gains realized by the relatively mild biases offered by the $\gamma = 0.05$ attenuated enhanced sampling simulations since these barriers can be surmounted by sufficiently long unbiased trajectories. In our next application, we apply GREST to a larger, more complex molecular system where we demonstrate that the protocol does realize substantial efficiency gains.

3.3 WLALL pentapeptide

As a more complex biomolecular system, we apply GREST to the Trp-Leu-Ala-Leu-Leu (WLALL) pentapeptide.⁹⁹ This system exhibits a number of metastable states separated by free energy barriers that are sufficiently high to demonstrate the capabilities of GREST to identify and surmount these barriers while still being sufficiently low to enable determination of a ground truth benchmark by long unbiased simulations.

Unbiased reference benchmark. We generate unbiased reference simulation data for the WLALL system in a similar manner to ADP. We first conduct an initial unbiased 250 ns simulation then apply k-means clustering in a 2D TICA embedding using pairwise backbone distances as features to seed initial configurations for $8 \times 0.75 \mu\text{s}$ independent simulations and produce $6 \mu\text{s}$ of aggregated reference data. As we will see, the slowest relaxing time scale in this system is on the order of 19.6 ns, so these data provide comprehensive sampling of configurational space. For the WLALL system, and many other systems of practical

interest, intuitive low dimensional coordinates suitable for visualization, such as the two backbone dihedrals for ADP, are not known *a priori*. For convenient visualization of the WLALL free energy surface, we use TICA^{36,61,70} as a simple and convenient dimensionality reduction technique to recover a two-dimensional representation of the WLALL phase space independently of our SRVs. We represent the WLALL system to both the TICA and SRV approaches as the $\binom{20}{2} = 190$ pairwise atomic distances between the 20 backbone atoms to define a translationally and rotationally invariant featurization $\mathbf{x} \in \mathbb{R}^{190}$. The free energy surface resulting from projection of the unbiased reference trajectories into the leading two TICA coordinates TIC_1 and TIC_2 is presented in Fig. 5a and reveals three predominant macrostates: the unfolded ensemble (labeled 1), the folded state (2), and the misfolded state (3). With the global free energy minima located in the unfolded ensemble under the prevailing thermodynamic conditions, the compact folded state is differentiated from the misfolded state with a comparatively lower free energy minima.

We next train an SRV over the unbiased reference data. An implied time scale analysis reveals a gap in the eigenvalue spectrum after 2-3 leading slow modes (Fig. S3 in the Supporting Information). For simplicity and visualization convenience, choose to retain the leading two slow modes throughout our analyses of the WLALL system. As we shall see, this results in empirically good performance in recapitulating and enhancing sampling of the pentapeptide dynamics. The first slow mode $\phi_1(\mathbf{x})$ possesses an implied time scale of $t_1 = 19.6$ ns and characterizes transitions in and out of the misfolded state (Fig. 5b). The second slow mode $\phi_2(\mathbf{x})$ possesses an implied time scale of $t_2 = 11.3$ ns and corresponds to transitions in and out of the folded state (Fig. 5c). Having established this baseline understanding of the structure, thermodynamics, and kinetics from a 6 μ s unbiased reference simulation, we now challenge GREST to automatically discover these slow modes and exploit them for efficient enhanced sampling.

Application of GREST. We follow the same adaptive sampling routine for the WLALL system as we demonstrated with ADP, beginning with a short unbiased 20 ns simulation to




Figure 5: Unbiased reference simulations of WLALL pentapeptide. (a) The free energy surface resulting from projection of 6 μ s of unbiased simulation data into the leading two TICA coordinates identifies three predominant metastable minima marked by golden stars and annotated with their corresponding relative free energy. Insets show a collection of representative structures from each minima identifying the (1) unfolded state, (2) folded state, and (3) misfolded state. (b) Application of SRVs to the unbiased reference data resolves a leading slow mode $\phi_1(\mathbf{x})$ with an implied time scale of $t_1 = 19.6$ ns characterizing transitions in and out of the misfolded state. (c) The next slowest mode $\phi_2(\mathbf{x})$ possesses an implied time scale of $t_2 = 11.3$ ns and corresponds to transitions in and out of the folded state.

initialize Round 0. We illustrate the resulting free energy surface in Fig. 6a-i by projecting these data into the two leading TICA coordinates TIC_1 and TIC_2 identified from our analysis of the 6 μ s unbiased reference simulations. (We emphasize that the TICA coordinates are used exclusively to provide a consistent and interpretable visual presentation of our data as we progress through the GREST iterations, and the unbiased simulation data and TICA analysis thereof play no role in the GREST analysis pipeline.) This initial Round 0 remains trapped in the unfolded ensemble and fails to explore any of the folded or misfolded states, which results in learned SRV slow modes that only characterize transitions within the unfolded ensemble (Fig. 6a-ii,iii). The subsequent WTMetaD run (Fig. 6a-iv) and subsequent simulation employing a time invariant $\gamma = 0.05$ attenuated terminal metadynamics bias are executed in the same manner as ADP, but in this instance we perform k-means clustering in the space of our pairwise backbone distance featurization $\mathbf{x} \in \mathbb{R}^{190}$ to initialize 25×1 ns

parallel walkers. In Round 1, we observe that the system escapes the unfolded ensemble and successfully explores the folded state (Fig. 6b-i), which is now captured in the leading SRV slow mode (Fig. 6b-ii) and results in accelerated metadynamics sampling along the transition path connecting the unfolded and folded states (Fig. 6b-iv). This process continues with the discovery and SRV characterization of the misfolded state in Round 2 (Fig. 6c), and convergence of the SRV slow modes after approximately five rounds (Fig. 6a-f, Fig. 7).

Visual comparison of the free energy surface in the terminal round of GREST (Fig. 6f-i) with that computed from the long unbiased reference data (Fig. 5a) demonstrates that GREST learned to identify the relevant transitions between states and to recover all of the important metastable states. Relative to the ~ 20 ns relaxation time of the slowest mode in the system, GREST also accurately learned the leading slow relaxations from quite parsimonious simulation data. A path reweighted SRV fitted to the accumulated 125 ns of biased GREST simulation data identifies two slow modes that exhibit Pearson correlations of $\rho_1 = 0.91$ and $\rho_2 = 0.99$, respectively, with those estimated from the 6 μ s of unbiased reference data. An implied time scale analysis of the biased data using a path reweighted SRV also reveals converged timescales of the leading two slow modes in good agreement with the unbiased reference data (Fig. S4 in the Supporting Information).

Finally, we verify that the enhanced sampling procedure effected within GREST does substantially improve sampling even under the $\gamma = 0.05$ attenuation factor in the metadynamics bias potential by surgically directing sampling along the relevant slow modes albeit with relatively modest accelerations to preserve numerical stability. Employing the same computational budget of 250 ns of simulation time, we demonstrate that GREST outperforms (i) a single long 250 ns unbiased simulation, (ii) 125 ns of unbiased simulation + 10×12.5 ns of unbiased simulations seeded by k-means clustering in pairwise backbone atom distances, and (iii) 250 ns of metadynamics performed in the squared molecular radius of gyration R_g^2 as an intuitive CV (Fig. S5 in the Supporting Information). These alternative approaches fail to identify all relevant metastable states and result in free energy surfaces




Figure 6: Application of GREST to WLALL pentapeptide. Rows (a)-(f) index successive rounds of GREST, commencing with the initial unbiased Round 0 in (a) and ending with the terminal adaptive sampling Round 5 in (f). Columns (i)-(iv) index various facets of each round of GREST: (i) illustrates the unbiased free energy landscape estimated by thermodynamic reweighting of the enhanced sampling simulations performed under the fixed terminal metadynamics bias data projected into the two leading TICA coordinates TIC_1 and TIC_2 , (ii) and (iii) illustrate heatmaps of the path reweighted SRV slow modes $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$ learned from the enhanced sampling trajectories in the terminal metadynamics bias potential projected into TIC_1 - TIC_2 space, and (iv) illustrates the phase space sampling density projected into TIC_1 - TIC_2 space from well-tempered metadynamics simulations performed in the two learned slow modes $\phi_1(\mathbf{x})$ and $\phi_2(\mathbf{x})$ illustrated in the preceding columns. The GREST iterations are terminated when the CVs learned in successive rounds stabilize, indicating no additional enhancement in configurational and dynamical exploration of phase space. Convergence of GREST at Round 4 means that additional enhanced sampling calculations are not conducted in Round 5.




Figure 7: Convergence assessment of GREST applied to WLALL pentapeptide by measuring the Pearson correlation coefficient between SRV slow CVs discovered in successive rounds. Stabilization of the CVs to $\rho > 0.8$ indicates that no important new dynamical modes or regions of configurational phase space are being explored under additional rounds of biased sampling.

with inferior coverage of the thermally accessible phase space compared to GREST.

4 Discussion & Conclusions

In this work we introduce Girsanov Reweighting Enhanced Sampling Technique (GREST) as an adaptive sampling scheme that interleaves rounds of data-driven slow collective variable discovery and enhanced sampling along these coordinates. We employ state-free reversible VAMPNets (SRVs) to perform data-driven discovery of slow collective variables. The estimation of dynamical observables under an unbiased Hamiltonian from biased trajectories collected under a perturbed Hamiltonian is enabled by the Girsanov formalism^{43,45–48} that prescribes how to apply reweighting of trajectories in the biased path ensemble according to both thermodynamic and integrator-specific dynamical path weights. This unlocks the ability to bias along slow coordinates to accelerate exploration of configurational phase space and then recover improved estimates for these slow coordinates under the enhanced sampling by applying Girsanov reweighting to the biased trajectories. This capability is the key feature underpinning GREST that interleaves successive rounds of path reweighted SRV slow collective variable discovery and metadynamics enhanced sampling in the learned coordinates until we achieve convergence.

In an application to a toy one-dimensional four-well system, we developed intuition for the behavior of GREST and showed how biasing along SRV learned CVs serve to reduce potential energy barriers mediating slow dynamical processes and accelerate sampling of phase space. In applications to two biomolecular systems – alanine dipeptide and the WLALL pentapeptide – we demonstrated how GREST can accelerate sampling of configurational phase space and furnish accurate estimates of the slow collective variables without any prior knowledge of the system. We achieve significant speedups in sampling the configurational phase space of the WLALL pentapeptide that enable GREST to discover all relevant metastable states whereas unbiased simulations and enhanced sampling in intuitive coordinates that are

not well correlated with the slow dynamical relaxations fail to explore all metastable states in an equivalent simulation time.

In future work, we anticipate that GREST may be a valuable approach in realizing significant efficiency gains in exploring the configurational phase space of large biomolecular and macromolecular systems. We also observe that enhanced sampling need not be limited to just those learned slow coordinates, and we propose that it may be profitable to augment GREST with sampling in both intuited and learned coordinates as an inexpensive means of incorporating prior knowledge into the approach. GREST is generically compatible with any off-the-shelf collective variable biasing technique, but may also benefit from integration with parallel tempering / replica exchange / Hamiltonian exchange formalisms to perform generic acceleration of “all” system degrees of freedom as well as targeted enhanced sampling along the learned slow modes.^{100–105} Path weights under Hamiltonian exchange can be treated under the same formalism presented in this work; analytical expressions for temperature reweighting and efficient numerical solution schemes have been previously reported by Chodera and co-workers.^{44,72,73} Another future extension of GREST would involve application to explicit solvent systems. Exposing both solvent and solute degrees of freedom to the SRV network via permutationally invariant featurizations, such as the Permutation Invariant Vector (PIV) formalism proposed by Fabio Pietrucci and co-workers,^{106,107} would enable learning of slow CVs over the solute and (permutationally invariant) solvent degrees of freedom. However, Basconi and Shirts have shown that Langevin dynamics, along with other velocity randomizing thermostats, may not be ideal for simulating solvated systems because of their tendency to artificially dampen dynamical processes due to disturbances of the natural time correlations of particle velocities and that these artifacts can only be avoided by using weak coupling constants that can lead to imprecise temperature control.¹⁰⁸ As such, care must be taken that the Langevin integration schemes for which analytical expressions for Girsanov weights are available must be deployed with care. A further avenue of exploration would be to conduct a comparison of the Girsanov reweighting formalism

with the square root approximation (SqRA) to dynamical reweighting.^{109–111} The SqRA requires quite strong assumptions regarding the character of the potential energy landscape and constancy of the diffusivity under an applied bias and so can only offer an approximate reweighting scheme, but can convey advantages in terms of numerical robustness when these conditions are satisfied and it offers an accurate approximation to the exact Girsanov expressions. The SqRA has been explored as an alternative to Girsanov reweighting by Donati, Weber, and Keller in the context of Markov state models,¹¹¹ and it would be interesting to explore the quality of this approximation and its range of applicability in the context of GREST. Finally, we observe that GREST may readily be applied to generic dynamical systems beyond molecular dynamics to accelerate phase space exploration in numerical simulations in, for example, ecological modeling, climate science, process systems modeling, or financial markets. The absence of detailed balance guarantees in generic dynamical systems mean that the SRVs must be generalized to SNRVs (state-free non-reversible VAMPNets) using the non-reversible VAC formalism,^{17,112} but otherwise the approach may be deployed in a largely analogous manner to that for molecular systems herein.

Data availability

A Python package implementing SRVs with capabilities for training over biased simulation data is publicly available via GitHub at <https://github.com/andrewlferguson/snrv>. Example Jupyter notebooks illustrating the application of GREST to the 4-well potential, alanine dipeptide, and WLALL pentapeptide systems along with all associated simulation data are hosted on Zenodo at DOI:10.5281/zenodo.7552310.¹¹³

Conflicts of interest

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, and 63/479,378 and International Patent Applications PCT/US2020/035206 and PCT/US2020/050466.

Supporting Information

Description of SRV models and training settings, simulation details and parameters for the GREST applications; implied time scale analysis and comparison in phase space sampling of different simulation techniques for the WLALL system.

Acknowledgement

This material is based on work supported by the National Science Foundation under Grant No. CHE-2152521. K.S. was supported by a fellowship from the Molecular Sciences Software Institute under the National Science Foundation Grant No. CHE-2136142. 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. We thank Dr. Wei Chen and Dr. Hythem Sidky for fruitful conversations at the outset of this work. This work is dedicated to Prof. Pablo G. Debenedetti as part of a 70th birthday *Festschrift* celebrating his pioneering contributions in statistical thermodynamics. The marriage of mathematical beauty and rigor and its reduction to computational practice is a hallmark of Pablo’s scientific perspective, and its impression upon his doctoral advisees is his gift that we carry with us in celebration of his legacy.

References

- (1) Frenkel, D.; Smit, B.; Ratner, M. A. *Understanding molecular simulation: from algorithms to applications*; Academic press San Diego, 1996; Vol. 2.
- (2) Hirst, J. D.; Glowacki, D. R.; Baaden, M. Molecular simulations and visualization: introduction and overview. *Faraday Discuss.* **2014**, *169*, 9–22.
- (3) Shaw, D. E.; Grossman, J. P.; Bank, J. A.; Batson, B.; Butts, J. A.; Chao, J. C.; Deneroff, M. M.; Dror, R. O.; Even, A.; Fenton, C. H.; A, F. Anton 2: raising the bar for performance and programmability in a special-purpose molecular dynamics supercomputer. SC14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. 2014; pp 41–53.
- (4) Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kale, L.; Schulten, K. Scalable molecular dynamics with NAMD. *J. Comput. Chem.* **2005**, *26*, 1781–1802.
- (5) Chow, E.; Rendleman, C. A.; Bowers, K. J.; Dror, R. O.; Hughes, D. H.; Gullingsrud, J.; Sacerdoti, F. D.; Shaw, D. E. Desmond performance on a cluster of multicore processors. *DE Shaw Research Technical Report DESRES/TR-2008-01* **2008**,
- (6) Glaser, J.; Nguyen, T. D.; Anderson, J. A.; Lui, P.; Spiga, F.; Millan, J. A.; Morse, D. C.; Glotzer, S. C. Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.* **2015**, *192*, 97–107.
- (7) Tchipev, N.; Seckler, S.; Heinen, M.; Vrabec, J.; Gratl, F.; Horsch, M.; Bernreuther, M.; Glass, C. W.; Niethammer, C.; Hammer, N.; Krischok, B.; Resch, M.; Kranzlmüller, D.; Hasse, H.; Bungartz, H.-J.; Neumann, P. TweTriS: Twenty trillion-atom simulation. *Int J High Perform Comput Appl.* **2019**, *33*, 838–854.

- (8) Elber, R. Perspective: Computer simulations of long time dynamics. *J. Chem. Phys.* **2016**, *144*, 060901.
- (9) Harvey, M. J.; Giupponi, G.; Fabritiis, G. D. ACEMD: accelerating biomolecular dynamics in the microsecond time scale. *J. Chem. Theory Comput.* **2009**, *5*, 1632–1639.
- (10) Freddolino, P. L.; Harrison, C. B.; Liu, Y.; Schulten, K. Challenges in protein-folding simulations. *Nat. Phys.* **2010**, *6*, 751–758.
- (11) Plattner, N.; Doerr, S.; De Fabritiis, G.; Noé, F. Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling. *Nat. Chem* **2017**, *9*, 1005–1011.
- (12) Roy, J.; Laughton, C. A. Long-timescale molecular-dynamics simulations of the major urinary protein provide atomistic interpretations of the unusual thermodynamics of ligand binding. *Biophys. J.* **2010**, *99*, 218–226.
- (13) Frederix, P. W.; Patmanidis, I.; Marrink, S. J. Molecular simulations of self-assembling bio-inspired supramolecular systems and their connection to experiments. *Chem Soc Rev* **2018**, *47*, 3470–3489.
- (14) Weinan, E.; Ren, W.; Vanden-Eijnden, E. String method for the study of rare events. *Phys. Rev. B* **2002**, *66*, 052301.
- (15) Bolhuis, P. G.; Chandler, D.; Dellago, C.; Geissler, P. L. Transition Path Sampling: Throwing Ropes. *Annu. Rev. Phys. Chem* **2002**, *53*, 291–318.
- (16) Allen, R. J.; Valeriani, C.; Ten Wolde, P. R. Forward flux sampling for rare event simulations. *J. Phys. Condens. Matter* **2009**, *21*, 463102.
- (17) Sidky, H.; Chen, W.; Ferguson, A. L. Machine learning for collective variable discovery and enhanced sampling in biomolecular simulation. *Mol. Phys.* **2020**, *118*, e1737742.

- (18) Barducci, A.; Bussi, G.; Parrinello, M. Well-tempered metadynamics: a smoothly converging and tunable free-energy method. *Phys. Rev. Lett.* **2008**, *100*, 020603.
- (19) Voter, A. F. Hyperdynamics: Accelerated molecular dynamics of infrequent events. *Phys. Rev. Lett.* **1997**, *78*, 3908.
- (20) Torrie, G. M.; Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *J. Comput. Phys.* **1977**, *23*, 187–199.
- (21) Straatsma, T.; Berendsen, H. Free energy of ionic hydration: Analysis of a thermodynamic integration technique to evaluate free energy differences by molecular dynamics simulations. *J. Chem. Phys.* **1988**, *89*, 5876–5886.
- (22) Valsson, O.; Parrinello, M. Variational approach to enhanced sampling and free energy calculations. *Phys. Rev. Lett.* **2014**, *113*, 090601.
- (23) García, A. E. Large-amplitude nonlinear motions in proteins. *Phys. Rev. Lett.* **1992**, *68*, 2696.
- (24) Amadei, A.; Linssen, A. B.; Berendsen, H. J. Essential dynamics of proteins. *Proteins: Struct. Funct. Genet.* **1993**, *17*, 412–425.
- (25) Hegger, R.; Altis, A.; Nguyen, P. H.; Stock, G. How complex is the dynamics of peptide folding? *Phys. Rev. Lett.* **2007**, *98*, 028102.
- (26) Zhuravlev, P. I.; Materese, C. K.; Papoian, G. A. Deconstructing the native state: energy landscapes, function, and dynamics of globular proteins. *J. Phys. Chem. B* **2009**, *113*, 8800–8812.
- (27) Das, P.; Moll, M.; Stamati, H.; Kavraki, L. E.; Clementi, C. Low-dimensional, free-energy landscapes of protein-folding reactions by nonlinear dimensionality reduction. *Proc. Natl. Acad. Sci. U.S.A.* **2006**, *103*, 9885–9890.

- (28) Rohrdanz, M. A.; Zheng, W.; Clementi, C. Discovering mountain passes via torchlight: Methods for the definition of reaction coordinates and pathways in complex macromolecular reactions. *Annu. Rev. Phys. Chem.* **2013**, *64*, 295–316.
- (29) Shamsi, Z.; Cheng, K. J.; Shukla, D. Reinforcement learning based adaptive sampling: REAPing rewards by exploring protein conformational landscapes. *J. Phys. Chem. B* **2018**, *122*, 8386–8395.
- (30) Ribeiro, J. M. L.; Bravo, P.; Wang, Y.; Tiwary, P. Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). *J. Chem. Phys.* **2018**, *149*, 072301.
- (31) Zheng, W.; Rohrdanz, M. A.; Clementi, C. Rapid exploration of configuration space with diffusion-map-directed molecular dynamics. *J. Phys. Chem. B* **2013**, *117*, 12769–12776.
- (32) Chiavazzo, E.; Covino, R.; Coifman, R. R.; Gear, C. W.; Georgiou, A. S.; Hummer, G.; Kevrekidis, I. G. Intrinsic map dynamics exploration for uncharted effective free-energy landscapes. *Proc. Natl. Acad. Sci. U.S.A.* **2017**, *114*, E5494–E5503.
- (33) Chen, W.; Ferguson, A. L. Molecular enhanced sampling with autoencoders: On-the-fly collective variable discovery and accelerated free energy landscape exploration. *J. Comput. Chem.* **2018**, *39*, 2079–2102.
- (34) Noé, F. *Machine learning meets quantum physics*; Springer, 2020; pp 331–372.
- (35) Naritomi, Y.; Fuchigami, S. Slow dynamics in protein fluctuations revealed by time-structure based independent component analysis: the case of domain motions. *J. Chem. Phys.* **2011**, *134*, 02B617.
- (36) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. *J. Chem. Phys.* **2013**, *139*, 07B604.1.

- (37) Schwantes, C. R.; Pande, V. S. Improvements in Markov state model construction reveal many non-native interactions in the folding of NTL9. *J. Chem. Theory Comput.* **2013**, *9*, 2000–2009.
- (38) Schwantes, C. R.; Pande, V. S. Modeling molecular kinetics with tICA and the kernel trick. *J. Chem. Theory Comput.* **2015**, *11*, 600–608.
- (39) Chen, W.; Sidky, H.; Ferguson, A. L. Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets. *J. Chem. Phys.* **2019**, *150*, 214114.
- (40) Bonati, L.; Piccini, G.; Parrinello, M. Deep learning the slow modes for rare events sampling. *Proc. Natl. Acad. Sci. U.S.A.* **2021**, *118*.
- (41) Sidky, H.; Chen, W.; Ferguson, A. L. High-resolution Markov state models for the dynamics of Trp-cage miniprotein constructed over slow folding modes identified by state-free reversible VAMPnets. *J. Phys. Chem. B* **2019**, *123*, 7999–8009.
- (42) Jones, M. S.; Ashwood, B.; Tokmakoff, A.; Ferguson, A. L. Determining Sequence-Dependent DNA Oligonucleotide Hybridization and Dehybridization Mechanisms Using Coarse-Grained Molecular Simulation, Markov State Models, and Infrared Spectroscopy. *J. Am. Chem. Soc.* **2021**, *143*, 17395–17411.
- (43) Kieninger, S.; Keller, B. G. Path probability ratios for Langevin dynamics—Exact and approximate. *J. Chem. Phys.* **2021**, *154*, 094102.
- (44) Chodera, J. D.; Swope, W. C.; Noé, F.; Prinz, J.-H.; Shirts, M. R.; Pande, V. S. Dynamical reweighting: Improved estimates of dynamical properties from simulations at multiple temperatures. *J. Chem. Phys.* **2011**, *134*, 06B612.
- (45) Weber, J. K.; Pande, V. S. Potential-based dynamical reweighting for Markov state models of protein dynamics. *J. Chem. Theory Comput.* **2015**, *11*, 2412–2420.

- (46) Girsanov, I. V. On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory Probab. its Appl.* **1960**, *5*, 285–301.
- (47) Donati, L.; Hartmann, C.; Keller, B. G. Girsanov reweighting for path ensembles and Markov state models. *J. Chem. Phys.* **2017**, *146*, 244112.
- (48) Donati, L.; Keller, B. G. Girsanov reweighting for metadynamics simulations. *J. Chem. Phys.* **2018**, *149*, 072335.
- (49) Wu, H.; Mey, A. S.; Rosta, E.; Noé, F. Statistically optimal analysis of state-discretized trajectory data from multiple thermodynamic states. *J. Chem. Phys.* **2014**, *141*, 12B629_1.
- (50) Wu, H.; Paul, F.; Wehmeyer, C.; Noé, F. Multiensemble Markov models of molecular thermodynamics and kinetics. *Proc. Natl. Acad. Sci. U.S.A.* **2016**, *113*, E3221–E3230.
- (51) Rosta, E.; Hummer, G. Free energies from dynamic weighted histogram analysis using unbiased Markov state model. *J. Chem. Theory Comput.* **2015**, *11*, 276–285.
- (52) Tiwary, P.; Parrinello, M. From metadynamics to dynamics. *Phys. Rev. Lett.* **2013**, *111*, 230602.
- (53) Casanovas, R.; Limongelli, V.; Tiwary, P.; Carloni, P.; Parrinello, M. Unbinding kinetics of a p38 MAP kinase type II inhibitor from metadynamics simulations. *J. Am. Chem. Soc.* **2017**, *139*, 4780–4788.
- (54) Izaguirre, J. A.; Sweet, C. R.; Pande, V. S. *Biocomputing 2010*; World Scientific, 2010; pp 240–251.
- (55) Prinz, J.-H.; Chodera, J. D.; Pande, V. S.; Swope, W. C.; Smith, J. C.; Noé, F. Optimal use of data in parallel tempering simulations for the construction of discrete-state Markov models of biomolecular dynamics. *J. Chem. Phys.* **2011**, *134*, 06B613.

- (56) Sidky, H.; Chen, W.; Ferguson, A. L. Molecular latent space simulators. *Chem. Sci.* **2020**, *11*, 9459–9467.
- (57) Vlachas, P. R.; Zavadlav, J.; Praprotnik, M.; Koumoutsakos, P. Accelerated Simulations of Molecular Systems through Learning of Effective Dynamics. *J. Chem. Theory Comput.* **2021**, *18*, 538–549.
- (58) Klus, S.; Nüske, F.; Koltai, P.; Wu, H.; Kevrekidis, I.; Schütte, C.; Noé, F. Data-driven model reduction and transfer operator approximation. *J. Nonlinear Sci.* **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) Noé, F.; Nuske, F. A variational approach to modeling slow processes in stochastic dynamical systems. *Multiscale Model Simul.* **2013**, *11*, 635–655.
- (61) Nuske, F.; Keller, B. G.; Pérez-Hernández, G.; Mey, A. S.; Noé, F. Variational approach to molecular kinetics. *J. Chem. Theory Comput.* **2014**, *10*, 1739–1752.
- (62) Wu, H.; Noé, F. Variational approach for learning Markov processes from time series data. *J. Nonlinear Sci.* **2020**, *30*, 23–66.
- (63) Szabo, A.; Ostlund, N. S. *Modern quantum chemistry: introduction to advanced electronic structure theory*; Courier Corporation, 2012.
- (64) Watkins, D. S. *The matrix eigenvalue problem: GR and Krylov subspace methods*; SIAM, 2007.
- (65) 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.

- (66) Wehmeyer, C.; Scherer, M. K.; Hempel, T.; Husic, B. E.; Olsson, S.; Noé, F. Introduction to Markov state modeling with the PyEMMA software [Article v1.0]. *Living J. Comp. Mol. Sci.* **2019**, *1*, 5965.
- (67) Husic, B. E.; Pande, V. S. Markov state models: From an art to a science. *J. Am. Chem. Soc.* **2018**, *140*, 2386–2396.
- (68) Sarich, M.; Noé, F.; Schütte, C. On the approximation quality of Markov state models. *Multiscale Model Simul* **2010**, *8*, 1154–1177.
- (69) Mardt, A.; Pasquali, L.; Wu, H.; Noé, F. VAMPnets for deep learning of molecular kinetics. *Nat. Commun* **2018**, *9*, 1–11.
- (70) Noé, F.; Clementi, C. Kinetic distance and kinetic maps from molecular dynamics simulation. *J. Chem. Theory Comput.* **2015**, *11*, 5002–5011.
- (71) Kingma, D. P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* **2014**,
- (72) Shirts, M. R.; Chodera, J. D. Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **2008**, *129*, 124105.
- (73) Minh, D. D.; Chodera, J. D. Optimal estimators and asymptotic variances for nonequilibrium path-ensemble averages. *J. Chem. Phys.* **2009**, *131*, 134110.
- (74) Zuckerman, D. M.; Woolf, T. B. Dynamic reaction paths and rates through importance-sampled stochastic dynamics. *J. Chem. Phys.* **1999**, *111*, 9475–9484.
- (75) Woolf, T. B. Path corrected functionals of stochastic trajectories: towards relative free energy and reaction coordinate calculations. *Chem. Phys. Lett.* **1998**, *289*, 433–441.
- (76) Zuckerman, D. M.; Woolf, T. B. Efficient dynamic importance sampling of rare events in one dimension. *Phys. Rev. E* **2000**, *63*, 016702.

- (77) Schütte, C.; Nielsen, A.; Weber, M. Markov state models and molecular alchemy. *Mol. Phys.* **2015**, *113*, 69–78.
- (78) Van Gunsteren, W.; Berendsen, H. Algorithms for Brownian dynamics. *Mol. Phys.* **1982**, *45*, 637–647.
- (79) Bussi, G.; Parrinello, M. Accurate sampling using Langevin dynamics. *Phys. Rev. E* **2007**, *75*, 056707.
- (80) Ceriotti, M.; Bussi, G.; Parrinello, M. Langevin equation with colored noise for constant-temperature molecular dynamics simulations. *Phys. Rev. Lett.* **2009**, *102*, 020601.
- (81) Goga, N.; Rzepiela, A.; De Vries, A.; Marrink, S.; Berendsen, H. Efficient algorithms for Langevin and DPD dynamics. *J. Chem. Theory Comput.* **2012**, *8*, 3637–3649.
- (82) Leimkuhler, B.; Matthews, C. Robust and efficient configurational molecular sampling via Langevin dynamics. *J. Chem. Phys.* **2013**, *138*, 05B601-1.
- (83) Fass, J.; Sivak, D. A.; Crooks, G. E.; Beauchamp, K. A.; Leimkuhler, B.; Chodera, J. D. Quantifying configuration-sampling error in Langevin simulations of complex molecular systems. *Entropy* **2018**, *20*, 318.
- (84) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Comput. Biol.* **2017**, *13*, e1005659.
- (85) Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; Chintala, S. In

- Advances in Neural Information Processing Systems 32*; Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc., 2019; pp 8024–8035.
- (86) Barducci, A.; Bonomi, M.; Parrinello, M. Metadynamics. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2011**, *1*, 826–843.
- (87) Pfendtner, J.; Bonomi, M. Efficient sampling of high-dimensional free-energy landscapes with parallel bias metadynamics. *J. Chem. Theory Comput.* **2015**, *11*, 5062–5067.
- (88) Prakash, A.; Fu, C. D.; Bonomi, M.; Pfendtner, J. Biasing smarter, not harder, by partitioning collective variables into families in parallel bias metadynamics. *J. Chem. Theory Comput.* **2018**, *14*, 4985–4990.
- (89) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Struct. Funct. Genet.* **2010**, *78*, 1950–1958.
- (90) Onufriev, A.; Bashford, D.; Case, D. A. Exploring protein native states and large-scale conformational changes with a modified generalized born model. *Proteins: Struct. Funct. Genet.* **2004**, *55*, 383–394.
- (91) MacKerell, A. D. J.; Bashford, D.; Bellott, M.; Dunbrack, R. L. J.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F. T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; Roux, B.; Schlenkrich, M.; Smith, J. C.; Stote, R.; Straub, J.; Watanabe, M.; Wiórkiewicz-Kuczera, J.; Yin, D.; Karplus, M. All-atom empirical potential for molecular modeling and dynamics studies of proteins. *J. Phys. Chem. B* **1998**, *102*, 3586–3616.

- (92) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. How robust are protein folding simulations with respect to force field parameterization? *Biophys. J.* **2011**, *100*, L47–L49.
- (93) Ferguson, A. L. BayesWHAM: A Bayesian approach for free energy estimation, reweighting, and uncertainty quantification in the weighted histogram analysis method. *J. Comput. Chem.* **2017**, *38*, 1583–1605.
- (94) Vitalini, F.; Mey, A. S.; Noé, F.; Keller, B. G. Dynamic properties of force fields. *J. Chem. Phys.* **2015**, *142*, 02B611_1.
- (95) Preto, J.; Clementi, C. Fast recovery of free energy landscapes via diffusion-map-directed molecular dynamics. *Phys. Chem. Chem. Phys.* **2014**, *16*, 19181–19191.
- (96) Hashemian, B.; Arroyo, M. Topological obstructions in the way of data-driven collective variables. *J. Chem. Phys.* **2015**, *142*, 044102.
- (97) Bolhuis, P. G.; Dellago, C.; Chandler, D. Reaction coordinates of biomolecular isomerization. *Proc. Natl. Acad. Sci. U.S.A.* **2000**, *97*, 5877–5882.
- (98) Jákli, I.; Jensen, S. J. K.; Csizmadia, I. G.; Perczel, A. Variation of conformational properties at a glance. True graphical visualization of the Ramachandran surface topology as a periodic potential energy surface. *Chem. Phys. Lett.* **2012**, *547*, 82–88.
- (99) Scherer, M. K.; Trendelkamp-Schroer, B.; Paul, F.; Pérez-Hernández, G.; Hoffmann, M.; Plattner, N.; Wehmeyer, C.; Prinz, J.-H.; Noé, F. PyEMMA 2: A software package for estimation, validation, and analysis of Markov models. *J. Chem. Theory Comput.* **2015**, *11*, 5525–5542.
- (100) Earl, D. J.; Deem, M. W. Parallel tempering: Theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.* **2005**, *7*, 3910–3916.
- (101) Hansmann, U. H. Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.* **1997**, *281*, 140–150.

- (102) Sugita, Y.; Okamoto, Y. Replica-exchange molecular dynamics method for protein folding. *Chem. Phys. Lett.* **1999**, *314*, 141–151.
- (103) Sugita, Y.; Okamoto, Y. Replica-exchange multicanonical algorithm and multicanonical replica-exchange method for simulating systems with rough energy landscape. *Chem. Phys. Lett.* **2000**, *329*, 261–270.
- (104) Fukunishi, H.; Watanabe, O.; Takada, S. On the Hamiltonian replica exchange method for efficient sampling of biomolecular systems: Application to protein structure prediction. *J. Chem. Phys.* **2002**, *116*, 9058–9067.
- (105) Bussi, G. Hamiltonian replica exchange in GROMACS: a flexible implementation. *Mol. Phys.* **2014**, *112*, 379–384.
- (106) Gallet, G. A.; Pietrucci, F. Structural cluster analysis of chemical reactions in solution. *J. Chem. Phys.* **2013**, *139*, 074101.
- (107) Pipolo, S.; Salanne, M.; Ferlat, G.; Klotz, S.; Saitta, A. M.; Pietrucci, F. Navigating at will on the water phase diagram. *Phys. Rev. Lett.* **2017**, *119*, 245701.
- (108) Basconi, J. E.; Shirts, M. R. Effects of temperature control algorithms on transport properties and kinetics in molecular dynamics simulations. *J. Chem. Theory Comput.* **2013**, *9*, 2887–2899.
- (109) Bicout, D.; Szabo, A. Electron transfer reaction dynamics in non-Debye solvents. *J. Chem. Phys.* **1998**, *109*, 2325–2338.
- (110) Wang, Y.; Tiwary, P. Understanding the role of predictive time delay and biased propagator in RAVE. *J. Chem. Phys.* **2020**, *152*, 144102.
- (111) Donati, L.; Weber, M.; Keller, B. G. A review of Girsanov reweighting and of square root approximation for building molecular Markov state models. *J. Math. Phys.* **2022**, *63*, 123306.

- (112) Noé, F. Machine Learning for Molecular Dynamics on Long Timescales. *arXiv preprint arXiv:1812.07669* **2018**,
- (113) Shmilovich, K.; Ferguson, A. L. Supporting data for "Girsanov Reweighting Enhanced Sampling Technique (GREST): On-the-fly data-driven discovery of and enhanced sampling in slow collective variables". 2023; <https://doi.org/10.5281/zenodo.7552310>.

TOC Graphic

