

Implicit Transfer Operator Learning: Multiple Time-Resolution Surrogates for Molecular Dynamics

Mathias Schreiner*
DTU[†]

Chalmers University of Technology
matschreiner@gmail.com

Ole Winther
DTU

Simon Olsson[‡]
Chalmers University of Technology
simonols@chalmers.se

Abstract

Computing properties of molecular systems rely on estimating expectations of the (unnormalized) Boltzmann distribution. Molecular dynamics (MD) is a broadly adopted technique to approximate such quantities. However, stable simulations rely on very small integration time-steps (10^{-15} s), whereas convergence of some moments, e.g. binding free energy or rates, might rely on sampling processes on time-scales as long as 10^{-1} s, and these simulations must be repeated for every molecular system independently. Here, we present Implicit Transfer Operator (ITO) Learning, a framework to learn surrogates of the simulation process with multiple time-resolutions. We implement ITO with denoising diffusion probabilistic models with a new SE(3) equivariant architecture and show the resulting models can generate self-consistent stochastic dynamics across multiple time-scales, even when the system is only partially observed. Finally, we present a coarse-grained CG-SE3-ITO model which can quantitatively model all-atom molecular dynamics using only coarse molecular representations. As such, ITO provides an important step towards multiple time- and space-resolution acceleration of MD. Code is available at <https://github.com/olsson-group/ito>

1 Introduction

Numerical simulation of stochastic differential equations (SDE) is critical in the sciences, including statistics, physics, chemistry, and biology applications [1]. Molecular dynamics (MD) simulations are an important example of such simulations [2]. These simulations prescribe a set of mechanistic rules governing the time evolution of a molecular system through numerical integration of, for example, the Langevin equation [3]. MD grants mechanistic insights into experimental observables. These observables are expectations, including time-correlations, of observable functions (e.g., pairwise distances or angles) computed for the Boltzmann distribution $\hat{\mu}(\mathbf{x}) \propto \exp[-\beta U(\mathbf{x})]$ corresponding to the potential $U(\cdot) : \Omega \rightarrow \mathbb{R}$ of a M -particle molecular system, $\mathbf{x} \in \Omega \subset \mathbb{R}^{3M}$ kept at the inverse temperature $\beta = 1/kT$. However, stable numerical integration relies on time steps, τ , which are strictly smaller than the fastest characteristic time-

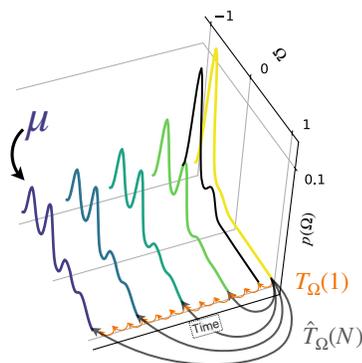


Figure 1: **Implicit Transfer Operator:** A multiple time-scale surrogate of stochastic molecular dynamics.

*Contributions to this work were done while visiting Chalmers University of Technology

[†]Technical University of Denmark

[‡]Corresponding author

scales of the molecular system (10^{-15} s, e.g., bond vibrations), yet many molecular systems are characterized by processes on much longer time-scales ($10^{-3} - 10^{-1}$ s, e.g. protein-folding, protein-ligand unbinding, regulation). Consequently, we need infeasibly long simulations to characterize many important processes quantitatively due to the slow mixing in Ω .

In this work, we present the implicit Transfer Operator (ITO, Fig. 1) as an effective way to learn multiple time-step surrogate models of the stochastic generating distribution of MD. To our knowledge, this is the first surrogate modeling approach that allows for the simultaneous generation of stochastic dynamics at multiple different time resolutions. By adopting an SE(3)-equivariant generative model, we further demonstrate stable long-time-scale dynamics in increasingly difficult settings where an increasing number of degrees of freedom are marginalized. Our approach can be several orders of magnitude more efficient than direct MD simulations and can be made asymptotically unbiased if the generative model permits exact likelihood evaluation. Our current results do not generalize across different thermodynamic ensembles or across chemical space, but show strong generalization across different time-scales.

Our main contributions are

1. the **Implicit Transfer Operator (ITO)** framework for learning generative models for multiple time resolution molecular dynamics simulations,
2. implementation of ITO using a denoising diffusion probabilistic model (DDPM) [4] with strong empirical results across resolutions: **SE(3)-equivariant ITO model (SE3-ITO)** gives stable long time-scale simulations and self-consistent dynamics across multiple time-scales for molecular benchmarks and **Coarse-grained SE3-ITO model (CG-SE3-ITO)** trained on large-scale protein folding data sets shows quantitative agreement with major dynamic and stationary observables of interest.

2 Background and Preliminaries

Notation Throughout this work, diffusion time, related to Diffusion Models (see Sec. 2), and physical time are represented using superscripts and subscripts, respectively.

Molecular dynamics and observables Molecular dynamics (MD) is a wide-spread simulation strategy in computational chemistry and physics. In this approach, the time-evolution of N particles configuration in Euclidean space $\mathbf{x} \in \Omega \subset \mathbb{R}^{3M}$, is modeled via a stochastic differential equation (SDE) with a drift term based on a potential energy model $U(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$. An important aim of MD is to compute:

1. **Stationary observables:** $O_f = \mathbb{E}_\mu[f(\mathbf{x})]$
2. **Dynamic observables:** $O_{f(t)h(t+\Delta t)} = \mathbb{E}_{\mathbf{x}_t \sim \mu}[\mathbb{E}_{\mathbf{x}_{t+\Delta t} \sim p_\tau(\mathbf{x}_{t+\Delta t} | \mathbf{x}_t)}[f(\mathbf{x}_t)h(\mathbf{x}_{t+\Delta t})]]$

where μ is the normalized Gibbs measure, and $p_\tau(\mathbf{x}_{t+\Delta t} | \mathbf{x}_t)$ is a conditional probability density function encoding the time-discrete evolution of the molecular system \mathbf{x} , with time-step $\Delta t = N\tau$ as prescribed by a dynamic model, e.g. *Langevin dynamics* [3], integrated with time-step, τ . N is typically a large integer. The functions $f, h : \Omega \rightarrow \mathbb{R}$ are observable functions or ‘*forward models*’ describing the microscopic observation process, e.g. computing a distance or an angle. The observables, $O_{f(t)}$, and $O_{f(t)h(t+\Delta t)}$, include binding affinities and binding rates of a drug to a protein, respectively. Conventionally, these observables are estimated from simulation trajectories using naive Monte Carlo estimators.

For illustrative purposes, we assume the temporal behavior of a state, \mathbf{x} , follows the Brownian dynamics SDE (Itô form)

$$d\mathbf{x}_t = -\nabla U(\mathbf{x}_t)\gamma^{-1} dt + \sqrt{2D}dW, \quad (1)$$

where $D = \gamma^{-1}\beta^{-1}$ is a diffusion constant, with friction γ and inverse-temperature β , and dW is a Wiener process. Using the Euler–Maruyama time-discretization, with time-step τ , simulating the SDE corresponds to simulating a Markov chain with the transition probability density

$$p(\mathbf{x}_{t+\tau} | \mathbf{x}_t, \tau) = \mathcal{N}(\mathbf{x}_{t+\tau} | \mathbf{x}_t - \tau\nabla U(\mathbf{x}_t)\gamma^{-1}, \tau\sqrt{2D}\mathbb{I}_{3M}) \quad (2)$$

where \mathcal{N} specifies the multi-variate Normal distribution, and \mathbb{I}_{3M} is the $3M$ -dimensional identity matrix. If τ is sufficiently small to allow stable simulation, the *invariant measure*, of the Markov chain

(eq. 2), is the Boltzmann distribution (normalized Gibbs measure) corresponding to the potential energy model $U(\mathbf{x})$ at β . Consequently, by simulating a large number of steps we can draw samples from μ to compute stationary observables and compute dynamic observables by simulating $\Delta t = N\tau$ steps enough times with initial states distributed according to μ . Explicit simulation make such computations extremely costly, and consequently, there’s much interest in speeding up the calculations of these quantities.

Transfer Operators Let ρ specify an initial condition, a probability density function on Ω . We can define a Markov operator $T_\Omega : L^1(\Omega) \rightarrow L^1(\Omega)$ using a transition density (e.g., 2):

$$[T_\Omega \circ \rho](\mathbf{x}_{t+\tau}) \triangleq \frac{1}{\mu(\mathbf{x}_{t+\tau})} \int_{x_t} \mu(\mathbf{x}_t) \rho(\mathbf{x}_t) p(\mathbf{x}_{t+\tau} | \mathbf{x}_t) d\mathbf{x}_t \quad (3)$$

which then describes the μ -weighed evolution of absolutely convergent probability density functions on Ω according to eq. 1, with time-step, τ . Such an operator is called the (Ruelle) Transfer Operator [5, 6]. We can express the operator using a spectral form

$$T_\Omega(\tau) = \sum_{i=0}^{\infty} \lambda_i(\tau) |\psi_i\rangle \langle \phi_i| \quad (4)$$

where only eigenvalues $\lambda_i(\tau) = \exp(-\tau\kappa_i)$ depend on the time-step, τ . κ_i are characteristic ‘relaxation’ rates associated the left and right eigenfunction pair, ϕ_i and ψ_i [7]. We can compute the operator with time-lag $N\tau$ via the Chapman-Kolmogorov equation (see Sec. A.1 for details)

$$T_\Omega(N\tau) = \sum_{i=0}^{\infty} \lambda_i(\tau)^N |\psi_i\rangle \langle \phi_i|. \quad (5)$$

Equivariant Message Passing Neural Networks In this work, we are concerned with MD, where the time-evolution of a molecule is governed by a force field $\mathcal{F}(\cdot) \triangleq -\nabla U(\cdot)$ derived from a central potential $U(\cdot)$. While $U(\cdot)$ is *invariant* to group-actions of the Euclidean group in three dimensions (E(3)), its corresponding force field is E(3)-*equivariant*. We call a function, f ‘*invariant*’ under a group-action g iff $f(\mathbf{x}) = f(S_g\mathbf{x})$ and ‘*equivariant*’ iff $T_g f(\mathbf{x}) = f(S_g\mathbf{x})$, where S_g and T_g are linear representations of the group element g [8].

The force field $\mathcal{F}(\cdot)$ is equivariant under E(3) group-actions. However, in practice, classical molecular dynamics simulations do not change parity during simulation, and consequently, our data distribution only contains a single mirror image of molecules.

We extended the PaiNN architecture [9], an E(3)-equivariant message passing neural network (MPNN), making it SE(3) equivariant by breaking its symmetry with respect to parity. We introduced this minor modification as we experienced sporadic parity changes when sampling with a model trained using the PaiNN architecture, and introducing this modification resolved the issue. Briefly, PaiNN embeds a graph $G = (V, E)$, where nodes, V , exchange equivariant messages through edges within a local neighborhood defined as $\mathcal{N}(i) = \{j | \|r_{ij}\| \leq r_{\text{cutoff}}\}$, where r_{ij} is the distance between nodes denoted i and j , and r_{cutoff} is the maximal distance at which nodes are allowed to exchange messages. Messages are pooled and subsequently used to update node features, thereby enabling exchange of equivariant information. We achieve parity symmetry-breaking by constructing the equivariant messages in a manner that depends on cross-products between equivariant node features and direction vectors between interacting nodes. The cross-product is an axial vector (i.e., does not change sign under parity). We combine these vectors with polar vectors (change sign under parity). We refer to this modified PaiNN architecture as ChiroPaiNN (CPaiNN). Further details are in the Appendix D.

Diffusion Models The diffusion model (DM) formalism is a powerful generative modeling framework that learns distributions by modeling a gradual denoising process [4, 10, 11]. In DMs, we pre-specify a *forward diffusion process* (noising process), which gradually transforms the data distribution $p(\mathbf{x}^0)$ to a simple prior distribution $p(\mathbf{x}^T)$, e.g., a standard Gaussian, through a time-inhomogenous Markov process, described by the following SDE (Itô form)

$$d\mathbf{x}^t = f(\mathbf{x}^t, t) dt + g(t) dW. \quad (6)$$

where $0 < t < T$ is the *diffusion time*, f and g are chosen functions, and dW is a Wiener process. We can generate samples from the data distribution $p(\mathbf{x}^0)$ by sampling from $p(\mathbf{x}^T)$ and solving the *backward diffusion process* (denoising process)

$$d\mathbf{x}^t = [f(\mathbf{x}^t, t) - g^2(t)\nabla_{\mathbf{x}^t} \log p(\mathbf{x}^t | t)] dt + g(t) dW \quad (7)$$

by approximating the *score field* $\nabla_{\mathbf{x}^t} \log p(\mathbf{x}^t | t)$ — or equivalently a time-dependent Normal transition kernel [4] — with a deep neural network surrogate $\nabla_{\mathbf{x}^t} \log \hat{p}(\mathbf{x}^t | t, \theta)$. We can use the learned score field to define a neural ordinary differential equation (ODE) [12, 13], or probability flow ODE [14] — eq. 7 less the term $g(t)dW$ and scaling $g^2(t)$ by $1/2$ — which we can leverage for efficient sampling and sample likelihood evaluation.

Here, we are concerned with building equivariant probability density functions under SE(3) group actions. Consequently, we parameterize the DM using a learned Normal transition kernel of a time-inhomogenous diffusion process. By restricting the transition kernels $p(\mathbf{x}^{t+1} | \mathbf{x}^t)$ to be equivariant under SE(3) group-actions, the marginal of \mathbf{x}^{t+1} is always invariant [15]. Combining the equivariant transition kernel with an invariant prior density [16] ensures the whole Markov process is invariant to SE(3) group actions. Consequently, combining an isotropic mean-free Gaussian as prior with ChiroPaiNN-parameterized transition kernels, we can construct an SE(3) equivariant diffusion model.

3 Implicit Transfer Operator

Molecular simulations are Markovian with transition density (e.g. eq. 2) and Normal, however, the latter only for very small *physical* time-steps τ . Here, we aim to approximate the long-time step transition probability $p_{N\tau}(\mathbf{x}_{N\tau} | \mathbf{x}_0)$ to allow for one-step sampling of long-time-scale dynamics.

As data we consider simulation trajectories. The trajectories are generated by explicit simulation which corresponds to sampling ancestrally from the small time-step transition density: $\mathbf{X} = \{\mathbf{x}_\tau, \dots, \mathbf{x}_{N\tau}\} \sim p(\mathbf{x}_{n\tau} | \mathbf{x}_{(n-1)\tau})$, with $n = \{1, \dots, N\}$. In general, the state variable \mathbf{x} , contains both position and velocity information of the particles, along with other details such as box dimensions, depending on the simulation scheme and target ensemble. Throughout this study, we only consider the position information.

We build a surrogate of the conditional transition probability distribution — $p_{N\tau}(\mathbf{x}_{N\tau} | \mathbf{x}_0)$ — from MD data. In practice, we learn a generative model $\mathbf{x}_{t+N\tau} \sim p_\theta(\mathbf{x}_{t+N\tau} | \mathbf{x}_t, N)$ with a conditional denoising diffusion probabilistic model (cDDPM) of the form

$$p(\mathbf{x}_{t+N\tau}^0 | \mathbf{x}_t, N) \triangleq \int p(\mathbf{x}_{t+N\tau}^{0:T} | \mathbf{x}_t, N) d\mathbf{x}^{1:T} \quad (8)$$

where $\mathbf{x}^{1:T}$ are *latent variables* of the same dimension as our output, and follow a joint density describing the backward diffusion process (eq. 7) and $\mathbf{x}^T \sim \mathcal{N}(0, \mathbb{I})$. We define a conditional sample likelihood as

$$\ell(\mathbf{I}; \theta) \triangleq \prod_{i \in \mathbf{I}} p_\theta(\mathbf{x}_{t_i+N_i\tau}^0 | \mathbf{x}_{t_i}, N_i) \quad (9)$$

where \mathbf{I} is a list of generated indices i specifying a time t_i and a time-lag (τ) integer multiple N_i , associating two time-points in the trajectory, \mathbf{X} . Following Ho et al., we train the cDDPM by

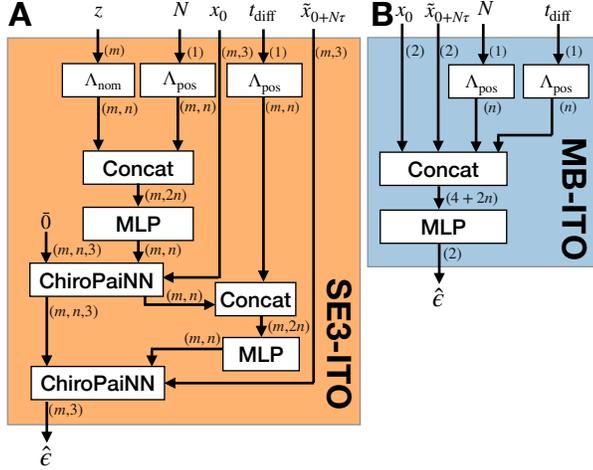


Figure 2: **ITO $\hat{\epsilon}$ networks** (A) SE3-ITO used for molecular application (B) MB-ITO, used for experiments with the Müller-Brown potential. Λ_{pos} and Λ_{nom} are positional and nominal embedding respectively, Concat is a concatenation, and MLP is a multi-layer perceptron. Arrows are annotated with input and output shapes.

optimizing a simplified form of the variational bound of the log-likelihood [4],

$$\mathcal{L}(\theta) = \mathbb{E}_{i \sim \mathbf{I}, \epsilon \sim \mathcal{N}(0, \mathbb{I}), t_{\text{diff}} \sim \mathcal{U}(0, T)} [\|\epsilon - \hat{\epsilon}_\theta(\tilde{\mathbf{x}}_{t_i+N_i\tau}^{t_{\text{diff}}}, \mathbf{x}_{t_i}, N_i, t_{\text{diff}})\|_2], \quad (10)$$

where $\tilde{\mathbf{x}}_t^{t_{\text{diff}}} \triangleq \sqrt{\bar{\alpha}^{t_{\text{diff}}}} \mathbf{x}_t + \sqrt{1 - \bar{\alpha}^{t_{\text{diff}}}} \epsilon$, with $\bar{\alpha}^{t_{\text{diff}}} = \prod_i^{t_{\text{diff}}} (1 - \beta_i)$ and β_i is the variance of the forward diffusion process at diffusion time, i . $\hat{\epsilon}_\theta(\cdot)$ is one of the two ITO neural network model architectures shown in Fig. 2 and is directly related to the score [4].

If the data used to train the conditional transition density is generated by MD simulation with time-invariant potential energy (drift), we can express the generating transition probability as a decomposition of time-variant and -invariant parts (Proof, see Sec. A.2)

$$p(x_{N\tau} | x_0) = \sum_{i=1}^{\infty} \underbrace{\lambda_i^N(\tau)}_{\text{time-variant}} \underbrace{\alpha_i(x_{N\tau})\beta_i(x_0)}_{\text{time-invariant}} \quad (11)$$

where α_i and β_i are *time-invariant* projection coefficients of the state variables on-to the left and right eigenfunctions ϕ_i and ψ_i , of the *Transfer operator* $T_\Omega(\tau)$ [5] and $|\lambda_i(\tau)| \leq 1$ is its i 'th eigenvalue. Consequently, we call our surrogate modeling approach *implicit transfer operator learning*.

As outlined in Algorithm 1, we generate the indices i , e.g. the tuples $(x_{t_i}, x_{t_i+N_i\tau}, N_i)$, in a manner such that the model is exposed to multiple time-lags, sampled uniformly across orders of magnitude, used for gradient-based optimization with Adam [17]. As a result, as illustrated in eq. 11, the model will be exposed to multiple different linear combinations of the eigenfunctions of $T_\Omega(\tau)$ in each batch during training. We conjecture that this data augmentation procedure will enable better learning of implicit representations of these eigenfunctions and, consequently, better generalization across time scales and yield more stable sampling.

3.1 ITO Architectures

We present two architectures for learning cDDPMs encoding ITO models, one for molecular applications SE3-ITO and one for the Müller-Brown benchmark system (Fig. 2). The SE3-ITO architecture uses our new SE(3) equivariant MPNN (ChiroPaiNN, described in sec. 2) to encode \mathbf{x}_t , N , and atom-types, z , to invariant features, s , and equivariant features, v . We concatenate s with an encoding of the diffusion-time t_{diff} and process them through a MLP (multi-layer perceptron). The output from the MLP are passed along with v and $\tilde{\mathbf{x}}_t^{t_{\text{diff}}}$ as input to a second ChiroPaiNN module which predicts $\hat{\epsilon}$. More details on the architecture and hyperparameters are available in Appendices D and E.

Algorithm 1 Training. DisExp is defined in Appendix E

Input: n MD-trajectories; $\mathcal{X} = \{\mathbf{x}_0^j, \dots, \mathbf{x}_{t_j}^j\}_{j=0}^n$, ITO score-model; $\hat{\epsilon}_\theta$, max lag; N_{max}
 $\mathcal{X}' = \text{Concatenate}(\{\mathbf{x}_0^j, \dots, \mathbf{x}_{t_j-N_{\text{max}}}^j\}_{j=0}^n)$
while not converged **do**
 $\mathbf{x}_t \sim \text{Choice}(\mathcal{X}')$
 $N \sim \text{DisExp}(N_{\text{max}})$
 $t_{\text{diff}} \sim \text{Uniform}(0, T)$
 Take gradient step on:
 $\nabla_\theta [\|\epsilon - \hat{\epsilon}_\theta(\tilde{\mathbf{x}}_{t+N\tau}^{t_{\text{diff}}}, \mathbf{x}_t, N, t_{\text{diff}})\|_2]$
end while
return $\hat{\epsilon}_\theta$

Algorithm 2 Ancestral sampling. Sampling from p_θ is defined in Appendix E, Algorithm 4

Input: initial condition \mathbf{x}_0 , lag N , ancestral steps n .
Allocate $\mathcal{T} \in \mathbb{R}^{(n+1) \times \dim(\mathbf{x}_0)}$
 $\mathcal{T}[0] = \mathbf{x}_0$
for $i = 1 \dots n$ **do**
 $\mathbf{x}_i \sim \hat{p}_\theta(\mathcal{T}[i-1], N)$
 $\mathcal{T}[i] = \mathbf{x}_i$
end for
return \mathcal{T}

4 Long time-step stochastic dynamics with Implicit Transfer Operators

4.1 Datasets and test-systems

To evaluate how robustly ITO models can model long time-scale dynamics, we conducted three classes of experiments, ranging from fully observed, high time-resolution, to sparsely observed and

Table 1: **VAMP2 score-gaps.** Difference in VAMP2-scores of ancestral sampling from ITO models with fixed lag and stochastic lags, compared to baseline Langevin simulations. Perfect match is 0, negative and positive values correspond to under and over estimation of meta-stability, respectively. Standard deviations on last decimal place are given in parentheses.

system \ lag	10	100
Müller-Brown (fixed)	-0.0351 (5)	-0.1189 (2)
Müller-Brown (stochastic)	-0.0312 (4)	-0.0970 (5)

low time resolution. Details on training and computational resources are available in Appendices E and F respectively.

Müller-Brown is a 2D potential commonly used for benchmarking molecular dynamics sampling methods. We generate a training data-set by integrating eq. 1 with the Müller-Brown potential energy as $U(\mathbf{x})$ (For details, see Appendix B.1). This dataset corresponds to a fully observed case.

Alanine dipeptide We use publicly available data from MDshare [18]. Simulation is performed with 2 fs integration time-steps and data is saved at 1 ps intervals. The simulations are performed in explicit solvation, but we only model the 22 atoms of the solute, without considering velocities. Consequently, this dataset is only partially observed.

Fast-folding proteins We use molecular dynamics data previously reported by Lindorff-Larsen et al. on the fast-folding proteins Chignolin, Trp-Cage, BBA, and Villin [19]. The data is proprietary but available upon request for research purposes. The simulations were performed in explicit solvent with a 2.5 fs time-step and the positions was saved at 200 ps intervals. We coarse-grain the simulation by representing each amino-acid by the Euclidean coordinate of their $C\alpha$ atom as done previously [20], leading to 10, 20, 28, and 35 particles in each system respectively. Consequently, these data correspond to a mostly unobserved case.

4.2 Stochastic lag improves meta-stability prediction

In sec. 2 we conjecture that exposing an ITO model to multiple lag times during training leads to better and more robust models. To test this, we trained a set of models on the Müller-Brown dataset with fixed constant lags $N = \{10, 100, 1000\}$ (fixed lag) and a single model with $N \sim \text{DisExp}(1000)$ (stochastic lag) using the MB-ITO model (Fig. 2).

We find that the model trained with a stochastic lag systematically outperforms models trained with fixed lag (Table 1). We gauge the agreement by comparing Variational Approach to Markov Processes (VAMP)-2 scores [21] (for details, see Appendix G), between model samples and training data and find that both models tend to underestimate meta-stability compared to training data slightly. However, the model trained with stochastic lag is marginally closer to the reference values. We note that the difference in the ability of fixed and stochastic lag ITO models to capture long-time-scale dynamics is also reflected in the learned transition densities (Fig. 5). Together, these results suggest that lag-time augmentation during training leads to better implicit learning of the Transfer operator’s eigenfunctions than training with a fixed lag.

To test whether this phenomena extends in cases where we do not have full observability and to molecular systems we followed the VAMP2-gaps of alanine di-peptide as a function of epoch for models trained with a fixed lag and a stochastic lag (Fig. 6). We find that the VAMP-2 gaps for stochastic lag and fixed lag in this case are statistically indistinguishable across all epochs. These results suggest that we can without compromising on accuracy build multiple time-scale surrogates by training with stochastic lag-times.

4.3 Efficient and accurate self-consistent long time-scale dynamics

We evaluate the ITO models trained with stochastic lags to capture long time-scale dynamics in a self-consistent manner, in the Chapman-Kolmogorov sense, i.e., $p(\mathbf{x}_{\Delta t} | \mathbf{x}_0) \triangleq p(\mathbf{x}_{N\tau} | \mathbf{x}_0) = \prod_{i=1}^N p(\mathbf{x}_{i\tau} | \mathbf{x}_{(i-1)\tau})$, or if samples generated by direct sampling with time-step $\Delta t = N\tau$ are

distributed similarly to samples generated by performing ancestral sampling N times, each with time-step τ . In direct sampling, we draw samples for a desired time-step $\Delta t = N\tau$ from $\hat{p}_\theta(\mathbf{x}_0, N)$ and in ancestral sampling we draw n samples with time-step $\Delta t = N\tau$ from $\hat{p}_\theta(\mathbf{x}_0, N)$ in an ancestral manner, e.g. $x_{i+1} \sim \hat{p}_\theta(\mathbf{x}_i, N)$, where $i = \{0, \dots, n-1\}$.

For the fully-observed Müller-Brown case, we find that the ITO model is self-consistent by the strong overlap in transition densities sampled in a direct and ancestral manner (Algorithm 2). These results generalize to molecular systems and partially observed systems. Sampling an SE3-ITO model (Fig. 2) trained with alanine dipeptide data, we find strong agreement between the ancestrally and directly sampled transition densities (Fig. 3) and we again have a strong consistency with corresponding transition densities computed from molecular dynamics simulations. Note here, that the time-step of the ITO-sampled transition densities varies from 10^4 to 10^6 times the MD integration time-step. The transition densities for alanine di-peptide (Fig. 3) are calculated using 15'000 trajectories. For direct sampling, this means that we draw 15'000 samples in total. In the case of ancestral sampling, we sample 4, 64, or 512 steps for 15'000 trajectories, to match Δt .

Next, we consider four fast-folding proteins [19] where we coarse-grain the proteins by representing them only with their $C\alpha$ atoms. In this sparsely observed case (CG-SE3-ITO), we find strong model self-consistency, as shown by the comparison between conditional densities from the folded and unfolded states (Fig. 4) projected onto a linear subspace determined using *time-lagged independent component analysis* (time-lagged independent components, tIC) [22] (see Appendix B.3). Further, the long time-scale transition density gradually converges to the data distribution as expected.

Finally, by ancestral sampling (Algorithm 2), we perform a simulation of Chignolin with the same length as the training trajectory ($106 \mu s$), using a CG-SE3-ITO model, and compare with MD. The CG-SE3-ITO simulation is 2120 steps with $\Delta t = 5$ ns. Running in parallel, on a single Titan X GPU we can simulate the CG-SE3-ITO model at a rate of $363 \text{ ns}/(s_w M^2)$ where s_w denotes seconds wall-time (Appendix C.1). Remarkably, these trajectories are virtually indistinguishable in the slowly relaxing TICA coordinates, illustrating stability of ITO. These conclusions extend to the proteins Villin, BBA, and Trp-Cage (See Appendix, Figs. 8[9] and 10[0])

Together these results suggest that ITO models accurately and self-consistently capture the slow dynamics of molecular systems and are robust to situations where the system is only partially observed. In general, we expect robustness to sparsely observed representations as long as the input representations are sufficient to span the eigenfunctions of T_Ω [23, 24]. Approximation errors will translate into systematic under-estimation of relaxation time-scales [7], consistent with our slight under-estimation of VAMP-2 scores (Table 1). In future work, combining the learning of SE3-ITO models with a systematic scheme for coarse-graining [25, 26], could be an avenue for scaling to large-scale molecular systems at a low computational cost.

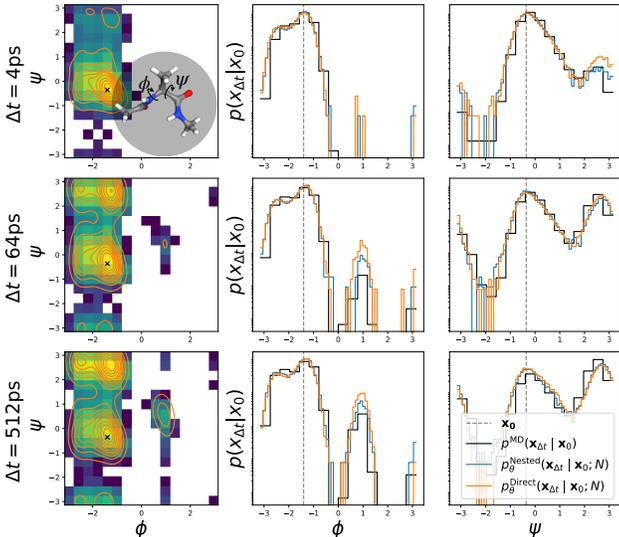


Figure 3: **Transition probability densities of alanine dipeptide dynamics with SE3-ITO model**; Rows of increasing time-lag (from top to bottom). Contours are samples from SE3-ITO model, and 2D histograms show estimates from MD data. The first column shows conditional transition densities projected onto the torsion angles ϕ and ψ (inset). The black cross indicates the initial condition. The second and third columns show marginal distributions of ϕ and ψ , respectively, with direct sampling in orange, ancestral sampling in blue, and MD data in black.

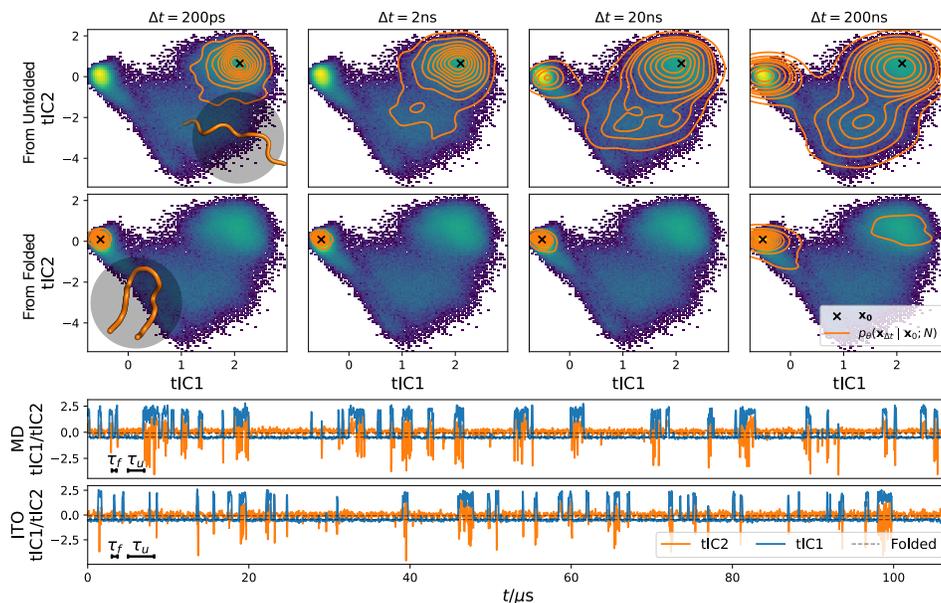


Figure 4: **Reversible protein folding-unfolding of Chignolin with CG-SE3-ITO** Conditional probability densities (orange contours) starting from unfolded (upper panels) and folded (lower panels) protein states, at increasing time-lag (left to right), shown on top of data distribution. Below: time-traces of 106 microsecond MD simulations and ITO simulations on tICs 1 and 2. The two dashed lines correspond to the folded state value in tIC 1 (lower line) and tIC 2 (higher line). Contour lines are based on 10^4 trajectories, generated with ancestral sampling with the length, Δt and time-step 200 ps. For $\Delta t = 200$ ns this corresponds to 1000 ancestral samples.

5 Prediction of dynamic and stationary observables of using CG-SE3-ITO

As outlined in section 2 an important aim of MD simulations is to compute stationary and dynamic observables, which involves intractable integrals typically approximated via Monte Carlo estimators. Using the trained ITO models we can efficiently sample i.i.d. from the transition density needed for computing dynamic observables, and by choosing a time-step which is sufficiently large we can also sample i.i.d. from the Boltzmann distribution μ , the latter akin to *Boltzmann generators* [27] (See Appendix A.1). We note that, the ITO models are surrogates and as such without reweighing we cannot expect unbiased samples from the Boltzmann and dynamic transition densities. Nevertheless, we gauge how accurately ITO models we can compute these observables of interest in the context of protein folding without reweighing:

- **Free Energy of Folding**, $\Delta G = -\log \left[\frac{p_f}{1-p_f} \right]$
- **Mean first passage time, folding**, $\langle \tau_f \rangle = \int_{x_0 \in -f} \int_0^\infty \delta(x_t \in f) p(x_t | x_0, t) dt dx_0$
- **Mean first passage time, unfolding**, $\langle \tau_u \rangle = \int_{x_0 \in f} \int_0^\infty \delta(x_t \in -f) p(x_t | x_0, t) dt dx_0$

where $\{f, -f\} \subset \Omega$ are disjoint subsets corresponding to the folded and unfolded states of a protein, $p_f = \int_{x \in f} \mu(x) dx$, is the folded state probability and $\delta(\cdot)$ is the Dirac delta.

We compute these observables using the reference molecular simulation data [19] and sample statistics from the CG-SE3-ITO models of each of the four fast-folding proteins (details in Appendix B.6). Strikingly, the observables computed using CG-SE3-ITO models agree well with those computed from long all-atom MD simulations (Table 2).

Finally, we analyzed the robustness, convergence, and consistency of these observables (Fig. 7). For Chignolin, we trained five models independently and analyzed model checkpoints when the training loss had stabilized. For each checkpoint and each model, we computed the observables. The values predicted are statistically indistinguishable, suggesting consistency, robustness, and convergence. The

Table 2: **Molecular observables** Standard deviations on last decimal place are given in parentheses. Stationary and dynamic observables are denoted **s** and **d**, respectively.

	$\Delta G_{\text{fold}}/kT$ (s)	$\langle \tau_f \rangle / \mu s$ (d)	$\langle \tau_u \rangle / \mu s$ (d)
Chignolin (MD/ITO)	-1.28(1)/-0.64(33)	0.565(4)/1.02(24)	2.01(2)/2.12(34)
Trp-Cage (MD/ITO)	1.47(6)/2.84(6)	13.6(4)/37(2)	3.4(2)/2.85(9)
BBA (MD/ITO)	0.97(3)/1.52(3)	11.7(2)/8.6(2)	5.1(1)/1.75(4)
Villin (MD/ITO)	1.21(2)/2.22(3)	2.41(3)/3.27(7)	0.68(1)/0.354(5)

average predictions closely match the reference values. Nevertheless, we note that the fluctuations in these values are noticeable.

We implemented all experiments using PyTorch [28], PyTorch Lightning [29], JAX [30], and used DPM-Solver [31] for probability flow ODE Sampling.

6 Related Work

Molecular sampling Sampling molecular configurations is a broad field and can broadly be divided into two main areas: physically motivated sampling of the Boltzmann distribution and conformer generation. The first area includes algorithmic approaches to sample the Boltzmann distribution including Molecular Dynamics simulations [2], Markov Chain Monte Carlo, extended ensemble methods [32, 33, 34, 35], including analysis methods involving deep generative nets [36], and surrogate models which directly approximate the Boltzmann distribution and allow for recovery of unbiased statistics, including Boltzmann generators [37, 16, 38, 39]. Conformer generation concerns generating physically plausible conformers without explicitly trying to follow the Boltzmann distribution. The latter approaches can be split into ML [15, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49] and chemoinformatic [50, 51] approaches. Finally, speeding up molecular simulations by reducing the effective number of particles to simulate through coarse-graining with special purpose forcefield models [52] including machine learned variants [53, 54, 20, 55] and learned coarse-graining maps [25, 26] is an orthogonal approach to sample conformation space. Further, several methods to recover all-atom models from coarse-grained representations through ML [56, 57] and rule-based approaches [58] are available.

Transfer Operator surrogates Building transfer operator surrogates is commonly used in molecular modeling including (Deep Generative) Markov state models (MSM) [59, 7, 60, 61], also including experimental data [62, 63] dynamic graphical models, [64] VAMPnets [65, 21], observable operator models [66], however, primarily for analysis of molecular dynamics data. Markov state models are time-space discrete approximations of the transfer operator and Deep Generative MSM [67] and VAMPnets [65] are deep learning infused versions, where state discretization is learned by deep nets. Dynamic graphical models reparameterize MSMs as kinetic Markov random fields allowing for scaling to larger systems [64]. Klein et al. recently introduced *timewarp* which is a flow-based generative model to simulate molecular systems with a large (up to 0.5 ns), fixed, time-lag, [68] providing asymptotically unbiased equilibrium samples through a Metropolis-Hastings correction [69]. While *timewarp* generates conformers with realistic local structure, it has limitations in capturing long time-scale dynamics, which is reflected in the predicted transition probability densities. In contrast, our approach captures long time-scale dynamics efficiently allowing for accurate prediction of dynamic observables. However, currently, neither the code nor the data from *timewarp* is publicly available precluding direct comparisons on the benchmark tasks established in their paper.

7 Limitations

Surrogate model Implicit Transfer Operators are surrogate models of stochastic dynamics’ conditional transition probability densities. We cannot guarantee unbiased sampling of dynamics and the stationary distribution due to aleatoric (e.g., finite data) and epistemic (e.g., model misspecification) uncertainty. We can overcome the latter by reweighing against a Markov Chain Monte Carlo acceptance criterion as proposed previously [68], to ensure unbiased dynamics path-reweighing is necessary, which in turn requires closed-form expressions for the target path probabilities [70].

Transferability and scalability Currently, ITO does not generalize across chemical space and thermodynamic variables. In future work, we anticipate that generalization across chemical space limitations can be overcome by appropriate data set curation and parameter-sharing schemes [71]. Generalization across thermodynamic variables such as temperature and pressure would require using a surrogate model which is steerable under these changes, e.g., temperature steerable flows [72] or thermodynamic maps [73]. Currently, we assume a fully connected graph that scales $\mathcal{O}(M^2)$ in system size, which limits what systems are practically accessible. Devising new surrogate models which use mean-field approximation approaches from e.g., computational physics [74] or chemistry to truncate the graphs and treat long-range as an additive term [75] could yield more favorable scaling [76].

8 Conclusions

This paper introduces Implicit Transfer Operators (ITO), an approach to building multiple time-scale surrogate models of stochastic molecular dynamics. We implement ITO models with a conditional DDPM using a new time-augmentation scheme and show how ITO models capture fast and slow dynamics on benchmarks and molecular systems. We show ITO models are self-consistent over multiple time scales and highly robust to the marginalization of degrees of freedom in the system, which are unimportant to capture the long-time-scale dynamics. Combined with a SE(3) variant of the PaiNN architecture [9] (ChiroPaiNN), we further show strong empirical evidence of scaling to molecular applications, such as the folding of coarse-grained proteins. As such, we are confident that ITO is a stepping-stone toward general-purpose surrogates of molecular dynamics.

Acknowledgments and Disclosure of Funding

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and The Novo Nordisk Foundation (SURE, NNF19OC0057822). OW's work was funded in part by the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606) and by the Pioneer Centre for AI, DNRG grant number P1. Preliminary results were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at Alvis (project: NAISS 2023/22-392) partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

References

- [1] Bernt Øksendal. *Stochastic Differential Equations*. Springer Berlin Heidelberg, 2003. DOI: [10.1007/978-3-642-14394-6](https://doi.org/10.1007/978-3-642-14394-6), URL: <https://doi.org/10.1007/978-3-642-14394-6>.
- [2] Daan Frenkel and Berend Smit. *Understanding Molecular Simulation*. Elsevier, 2002. DOI: [10.1016/b978-0-12-267351-1.x5000-7](https://doi.org/10.1016/b978-0-12-267351-1.x5000-7), URL: <https://doi.org/10.1016/b978-0-12-267351-1.x5000-7>.
- [3] Paul Langevin. "Sur la théorie du mouvement brownien". In: *C. R. Acad. Sci. (Paris)* 146 (1908), 530—533.
- [4] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denosing Diffusion Probabilistic Models*. 2020. DOI: [10.48550/ARXIV.2006.11239](https://arxiv.org/abs/2006.11239), URL: <https://arxiv.org/abs/2006.11239>.
- [5] David Ruelle. *Thermodynamic Formalism*. en. Encyclopedia of mathematics and its applications. Harlow, England: Longman Higher Education, Nov. 1978.
- [6] Christof Schütte et al. "Conformation dynamics". In: *Proc. Int. Congr. ICIAM* (2009), pp. 297–336.
- [7] Jan-Hendrik Prinz et al. "Markov models of molecular kinetics: Generation and validation". In: *The Journal of Chemical Physics* 134.17 (May 2011), p. 174105. DOI: [10.1063/1.3565032](https://doi.org/10.1063/1.3565032), URL: <https://doi.org/10.1063/1.3565032>.
- [8] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Springer New York, 1977. DOI: [10.1007/978-1-4684-9458-7](https://doi.org/10.1007/978-1-4684-9458-7), URL: <https://doi.org/10.1007/978-1-4684-9458-7>.

- [9] Kristof T. Schütt, Oliver T. Unke, and Michael Gastegger. *Equivariant message passing for the prediction of tensorial properties and molecular spectra*. 2021. eprint: [arXiv:2102.03150](https://arxiv.org/abs/2102.03150).
- [10] Yang Song et al. *Score-Based Generative Modeling through Stochastic Differential Equations*. 2020. DOI: [10.48550/ARXIV.2011.13456](https://doi.org/10.48550/ARXIV.2011.13456) URL: <https://arxiv.org/abs/2011.13456>.
- [11] Jascha Sohl-Dickstein et al. “Deep Unsupervised Learning using Nonequilibrium Thermodynamics”. In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by Francis Bach and David Blei. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, 2015, pp. 2256–2265. URL: <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- [12] Ricky T. Q. Chen et al. “Neural Ordinary Differential Equations”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018. URL: https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf.
- [13] Will Grathwohl et al. “Scalable Reversible Generative Models with Free-form Continuous Dynamics”. In: *International Conference on Learning Representations*. 2019. URL: <https://openreview.net/forum?id=rJxgknCck7>.
- [14] Yang Song et al. “Score-Based Generative Modeling through Stochastic Differential Equations”. In: *International Conference on Learning Representations*. 2021. URL: <https://openreview.net/forum?id=PxtIG12RRHS>.
- [15] Minkai Xu et al. *GeoDiff: a Geometric Diffusion Model for Molecular Conformation Generation*. 2022. eprint: [arXiv:2203.02923](https://arxiv.org/abs/2203.02923).
- [16] Jonas Köhler, Leon Klein, and Frank Noe. “Equivariant Flows: Exact Likelihood Generative Learning for Symmetric Densities”. In: *Proceedings of the 37th International Conference on Machine Learning*. Ed. by Hal Daumé III and Aarti Singh. Vol. 119. Proceedings of Machine Learning Research. PMLR, 2020, pp. 5361–5370. URL: <https://proceedings.mlr.press/v119/kohler20a.html>.
- [17] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2014. eprint: [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [18] URL: <https://markovmodel.github.io/mdshare/ALA2/#alanine-dipeptide>.
- [19] K. Lindorff-Larsen et al. “How Fast-Folding Proteins Fold”. In: *Science* 334.6055 (Oct. 2011), pp. 517–520. DOI: [10.1126/science.1208351](https://doi.org/10.1126/science.1208351) URL: <https://doi.org/10.1126/science.1208351>.
- [20] Marloes Arts et al. *Two for One: Diffusion Models and Force Fields for Coarse-Grained Molecular Dynamics*. 2023. eprint: [arXiv:2302.00600](https://arxiv.org/abs/2302.00600).
- [21] Hao Wu and Frank Noé. “Variational Approach for Learning Markov Processes from Time Series Data”. In: *Journal of Nonlinear Science* 30.1 (2019), pp. 23–66. DOI: [10.1007/s00332-019-09567-y](https://doi.org/10.1007/s00332-019-09567-y) URL: <https://doi.org/10.1007/s00332-019-09567-y>.
- [22] Guillermo Pérez-Hernández et al. “Identification of slow molecular order parameters for Markov model construction”. In: *The Journal of Chemical Physics* 139.1 (July 2013), p. 015102. DOI: [10.1063/1.4811489](https://doi.org/10.1063/1.4811489) URL: <https://doi.org/10.1063/1.4811489>.
- [23] Natasa Djurdjevac, Marco Sarich, and Christof Schütte. “Estimating the Eigenvalue Error of Markov State Models”. In: *Multiscale Modeling & Simulation* 10.1 (Jan. 2012), pp. 61–81. DOI: [10.1137/100798910](https://doi.org/10.1137/100798910) URL: <https://doi.org/10.1137/100798910>.
- [24] Marco Sarich, Frank Noé, and Christof Schütte. “On the Approximation Quality of Markov State Models”. In: *Multiscale Modeling & Simulation* 8.4 (Jan. 2010), pp. 1154–1177. DOI: [10.1137/090764049](https://doi.org/10.1137/090764049) URL: <https://doi.org/10.1137/090764049>.
- [25] Andreas Krämer et al. “Statistically Optimal Force Aggregation for Coarse-Graining Molecular Dynamics”. In: *The Journal of Physical Chemistry Letters* 14.17 (Apr. 2023), pp. 3970–3979. DOI: [10.1021/acs.jpcllett.3c00444](https://doi.org/10.1021/acs.jpcllett.3c00444) URL: <https://doi.org/10.1021/acs.jpcllett.3c00444>.
- [26] Wangfei Yang et al. “Slicing and Dicing: Optimal Coarse-Grained Representation to Preserve Molecular Kinetics”. In: *ACS Central Science* 9.2 (Jan. 2023), pp. 186–196. DOI: [10.1021/acscentsci.2c01200](https://doi.org/10.1021/acscentsci.2c01200) URL: <https://doi.org/10.1021/acscentsci.2c01200>.

- [27] Frank Noé et al. “Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations”. In: *Proceedings of the National Academy of Sciences* 106.45 (Nov. 2009), pp. 19011–19016. DOI: [10.1073/pnas.0905466106](https://doi.org/10.1073/pnas.0905466106), URL: <https://doi.org/10.1073/pnas.0905466106>.
- [28] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: 32 (2019). Ed. by H Wallach et al. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- [29] William Falcon and The PyTorch Lightning team. *PyTorch Lightning*. Version 1.4. Mar. 2019. DOI: [10.5281/zenodo.3828935](https://zenodo.org/record/3828935), URL: <https://github.com/Lightning-AI/lightning>.
- [30] James Bradbury et al. *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. 2018. URL: <http://github.com/google/jax>.
- [31] Cheng Lu et al. “DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps”. In: *arXiv preprint arXiv:2206.00927* (2022).
- [32] Jes Frellsen et al. “Bayesian Generalised Ensemble Markov Chain Monte Carlo”. In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*. Ed. by Arthur Gretton and Christian C. Robert. Vol. 51. Proceedings of Machine Learning Research. Cadiz, Spain: PMLR, 2016, pp. 408–416. URL: <http://proceedings.mlr.press/v51/frellsen16.html>.
- [33] Alessandro Laio and Michele Parrinello. “Escaping free-energy minima”. In: *Proceedings of the National Academy of Sciences* 99.20 (2002), pp. 12562–12566. ISSN: 0027-8424. DOI: [10.1073/pnas.202427399](https://www.pnas.org/content/99/20/12562.full.pdf), eprint: <https://www.pnas.org/content/99/20/12562.full.pdf>, URL: <https://www.pnas.org/content/99/20/12562>.
- [34] D. P. Landau, Shan-Ho Tsai, and M. Exler. “A new approach to Monte Carlo simulations in statistical physics: Wang-Landau sampling”. In: *American Journal of Physics* 72.10 (2004), pp. 1294–1302. DOI: [10.1119/1.1707017](https://doi.org/10.1119/1.1707017), eprint: <https://doi.org/10.1119/1.1707017>, URL: <https://doi.org/10.1119/1.1707017>.
- [35] Amey P. Pasarkar et al. “Vendi sampling for molecular simulations: Diversity as a force for faster convergence and better exploration”. In: *The Journal of Chemical Physics* 159.14 (Oct. 2023). DOI: [10.1063/5.0166172](https://doi.org/10.1063/5.0166172), URL: <https://doi.org/10.1063/5.0166172>.
- [36] Yihang Wang, Lukas Herron, and Pratyush Tiwary. “From data to noise to data for mixing physics across temperatures with generative artificial intelligence”. In: *Proceedings of the National Academy of Sciences* 119.32 (Aug. 2022). DOI: [10.1073/pnas.2203656119](https://doi.org/10.1073/pnas.2203656119), URL: <https://doi.org/10.1073/pnas.2203656119>.
- [37] Frank Noé et al. “Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning”. In: *Science* 365.6457 (Sept. 2019). DOI: [10.1126/science.aaw1147](https://doi.org/10.1126/science.aaw1147), URL: <https://doi.org/10.1126/science.aaw1147>.
- [38] Leon Klein, Andreas Krämer, and Frank Noé. *Equivariant flow matching*. 2023. eprint: [arXiv:2306.15030](https://arxiv.org/abs/2306.15030).
- [39] Juan Viguera Diez et al. “Generation of conformational ensembles of small molecules via Surrogate Model-Assisted Molecular Dynamics”. In: (Nov. 2023). DOI: [10.26434/chemrxiv-2023-sx61w](https://doi.org/10.26434/chemrxiv-2023-sx61w), URL: <http://dx.doi.org/10.26434/chemrxiv-2023-sx61w>.
- [40] Bowen Jing et al. *Torsional Diffusion for Molecular Conformer Generation*. 2022. eprint: [arXiv:2206.01729](https://arxiv.org/abs/2206.01729).
- [41] Elman Mansimov et al. “Molecular geometry prediction using a deep generative graph neural network”. In: *Scientific Reports* 9.1 (2019), pp. 1–13.
- [42] Chence Shi et al. *Learning Gradient Fields for Molecular Conformation Generation*. 2021. eprint: [arXiv:2105.03902](https://arxiv.org/abs/2105.03902), URL: <http://arxiv.org/abs/2105.03902>.
- [43] Gregor Simm and Jose Miguel Hernandez-Lobato. “A Generative Model for Molecular Distance Geometry”. In: *Proceedings of the 37th International Conference on Machine Learning*. Ed. by Hal Daumé III and Aarti Singh. Vol. 119. Proceedings of Machine Learning Research. PMLR, 2020, pp. 8949–8958. URL: <http://proceedings.mlr.press/v119/simm20a.html>.
- [44] Chence Shi et al. “Learning Gradient Fields for Molecular Conformation Generation”. In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 9558–9568. URL: <https://proceedings.mlr.press/v139/shi21b.html>.

- [45] Tarun Gogineni et al. “TorsionNet: A Reinforcement Learning Approach to Sequential Conformer Search”. In: *CoRR* abs/2006.07078 (2020). arXiv: [2006.07078](https://arxiv.org/abs/2006.07078), URL: <https://arxiv.org/abs/2006.07078>
- [46] Robin Winter, Frank Noé, and Djork-Arné Clevert. *Auto-Encoding Molecular Conformations*. 2021. URL: <http://arxiv.org/abs/2101.01618>
- [47] Victor Garcia Satorras et al. *E(n) Equivariant Normalizing Flows for Molecule Generation in 3D*. 2021. eprint: [2105.09016](https://arxiv.org/abs/2105.09016), URL: <http://arxiv.org/abs/2105.09016>
- [48] Niklas Gebauer, Michael Gastegger, and Kristof Schütt. “Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019. URL: <https://proceedings.neurips.cc/paper/2019/file/a4d8e2a7e0d0c102339f97716d2fdfb6-Paper.pdf>
- [49] Minkai Xu et al. *An End-to-End Framework for Molecular Conformation Generation via Bilevel Programming*. 2021. eprint: [arXiv:2105.07246](https://arxiv.org/abs/2105.07246)
- [50] Sereina Riniker and Gregory A. Landrum. “Better Informed Distance Geometry: Using What We Know To Improve Conformation Generation”. In: *Journal of Chemical Information and Modeling* 55.12 (Nov. 2015), pp. 2562–2574. DOI: [10.1021/acs.jcim.5b00654](https://doi.org/10.1021/acs.jcim.5b00654), URL: <https://doi.org/10.1021/acs.jcim.5b00654>
- [51] Paul C. D. Hawkins et al. “Conformer Generation with OMEGA: Algorithm and Validation Using High Quality Structures from the Protein Databank and Cambridge Structural Database”. In: *Journal of Chemical Information and Modeling* 50.4 (Mar. 2010), pp. 572–584. DOI: [10.1021/ci100031x](https://doi.org/10.1021/ci100031x), URL: <https://doi.org/10.1021/ci100031x>
- [52] Paulo C. T. Souza et al. “Martini 3: a general purpose force field for coarse-grained molecular dynamics”. In: *Nature Methods* 18.4 (Mar. 2021), pp. 382–388. DOI: [10.1038/s41592-021-01098-3](https://doi.org/10.1038/s41592-021-01098-3), URL: <https://doi.org/10.1038/s41592-021-01098-3>
- [53] Jiang Wang et al. “Machine Learning of Coarse-Grained Molecular Dynamics Force Fields”. In: *ACS Central Science* 5.5 (Apr. 2019), pp. 755–767. DOI: [10.1021/acscentsci.8b00913](https://doi.org/10.1021/acscentsci.8b00913), URL: <https://doi.org/10.1021/acscentsci.8b00913>
- [54] Brooke E. Husic et al. “Coarse graining molecular dynamics with graph neural networks”. In: *The Journal of Chemical Physics* 153.19 (Nov. 2020), p. 194101. DOI: [10.1063/5.0026133](https://doi.org/10.1063/5.0026133), URL: <https://doi.org/10.1063/5.0026133>
- [55] Jonas Köhler et al. “Flow-Matching: Efficient Coarse-Graining of Molecular Dynamics without Forces”. In: *Journal of Chemical Theory and Computation* 19.3 (Jan. 2023), pp. 942–952. DOI: [10.1021/acs.jctc.3c00016](https://doi.org/10.1021/acs.jctc.3c00016), URL: <https://doi.org/10.1021/acs.jctc.3c00016>
- [56] Soojung Yang and Rafael Gómez-Bombarelli. *Chemically Transferable Generative Backmapping of Coarse-Grained Proteins*. 2023. eprint: [arXiv:2303.01569](https://arxiv.org/abs/2303.01569)
- [57] Yaxin An and Sanket A. Deshmukh. “Machine learning approach for accurate backmapping of coarse-grained models to all-atom models”. In: *Chemical Communications* 56.65 (2020), pp. 9312–9315. DOI: [10.1039/d0cc02651d](https://doi.org/10.1039/d0cc02651d), URL: <https://doi.org/10.1039/d0cc02651d>
- [58] Leandro E. Lombardi, Marcelo A. Martí, and Luciana Capece. “CG2AA: backmapping protein coarse-grained structures”. In: *Bioinformatics* 32.8 (Dec. 2015), pp. 1235–1237. DOI: [10.1093/bioinformatics/btv740](https://doi.org/10.1093/bioinformatics/btv740), URL: <https://doi.org/10.1093/bioinformatics/btv740>
- [59] Christof Schütte et al. *A Hybrid Monte Carlo Method for Essential Molecular Dynamics*. eng. Tech. rep. SC-98-04. Takustr. 7, 14195 Berlin: ZIB, 1998.
- [60] William C. Swope, Jed W. Pitera, and Frank Suits. “Describing Protein Folding Kinetics by Molecular Dynamics Simulations. 1. Theory”. In: *The Journal of Physical Chemistry B* 108.21 (Apr. 2004), pp. 6571–6581. DOI: [10.1021/jp037421y](https://doi.org/10.1021/jp037421y), URL: <https://doi.org/10.1021/jp037421y>
- [61] Brooke E. Husic and Vijay S. Pande. “Markov State Models: From an Art to a Science”. In: *Journal of the American Chemical Society* 140.7 (Feb. 2018), pp. 2386–2396. DOI: [10.1021/jacs.7b12191](https://doi.org/10.1021/jacs.7b12191), URL: <https://doi.org/10.1021/jacs.7b12191>
- [62] Christopher Kolloff and Simon Olsson. “Rescuing off-equilibrium simulation data through dynamic experimental data with dynAMMo”. In: *Machine Learning: Science and Technology* 4.4 (Dec. 2023), p. 045050. ISSN: 2632-2153. DOI: [10.1088/2632-2153/ad10ce](https://doi.org/10.1088/2632-2153/ad10ce), URL: <http://dx.doi.org/10.1088/2632-2153/ad10ce>

- [63] Simon Olsson et al. “Combining experimental and simulation data of molecular processes via augmented Markov models”. In: *Proceedings of the National Academy of Sciences* 114.31 (July 2017), 8265–8270. ISSN: 1091-6490. DOI: [10.1073/pnas.1704803114](https://doi.org/10.1073/pnas.1704803114). URL: <http://dx.doi.org/10.1073/pnas.1704803114>
- [64] Simon Olsson and Frank Noé. “Dynamic graphical models of molecular kinetics”. In: *Proceedings of the National Academy of Sciences* 116.30 (July 2019), pp. 15001–15006. DOI: [10.1073/pnas.1901692116](https://doi.org/10.1073/pnas.1901692116). URL: <https://doi.org/10.1073/pnas.1901692116>
- [65] Andreas Mardt et al. “VAMPnets for deep learning of molecular kinetics”. In: *Nature Communications* 9.1 (2018). DOI: [10.1038/s41467-017-02388-1](https://doi.org/10.1038/s41467-017-02388-1). URL: <https://doi.org/10.1038/s41467-017-02388-1>
- [66] Hao Wu, Jan-Hendrik Prinz, and Frank Noé. “Projected metastable Markov processes and their estimation with observable operator models”. In: *The Journal of Chemical Physics* 143.14 (Oct. 2015), p. 144101. DOI: [10.1063/1.4932406](https://doi.org/10.1063/1.4932406). URL: <https://doi.org/10.1063/1.4932406>
- [67] Hao Wu et al. *Deep Generative Markov State Models*. 2018. eprint: [arXiv:1805.07601](https://arxiv.org/abs/1805.07601)
- [68] Leon Klein et al. *Timewarp: Transferable Acceleration of Molecular Dynamics by Learning Time-Coarsened Dynamics*. 2023. eprint: [arXiv:2302.01170](https://arxiv.org/abs/2302.01170)
- [69] W. K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1 (Apr. 1970), pp. 97–109. DOI: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97). URL: <https://doi.org/10.1093/biomet/57.1.97>
- [70] S. Kieninger and B. G. Keller. “Path probability ratios for Langevin dynamics—Exact and approximate”. In: *The Journal of Chemical Physics* 154.9 (Mar. 2021), p. 094102. DOI: [10.1063/5.0038408](https://doi.org/10.1063/5.0038408). URL: <https://doi.org/10.1063/5.0038408>
- [71] Geemi P Wellawatte, Glen M Hocky, and Andrew D White. “Neural potentials of proteins extrapolate beyond training data”. In: (2023).
- [72] Manuel Dibak et al. “Temperature steerable flows and Boltzmann generators”. In: *Physical Review Research* 4.4 (Oct. 2022). DOI: [10.1103/physrevresearch.4.1042005](https://doi.org/10.1103/physrevresearch.4.1042005). URL: <https://doi.org/10.1103/physrevresearch.4.1042005>
- [73] Lukas Herron et al. *Inferring phase transitions and critical exponents from limited observations with Thermodynamic Maps*. 2023. eprint: [arXiv:2308.14885](https://arxiv.org/abs/2308.14885)
- [74] V Rokhlin. “Rapid solution of integral equations of classical potential theory”. In: *Journal of Computational Physics* 60.2 (Sept. 1985), pp. 187–207. DOI: [10.1016/0021-9991\(85\)90002-6](https://doi.org/10.1016/0021-9991(85)90002-6). URL: [https://doi.org/10.1016/0021-9991\(85\)90002-6](https://doi.org/10.1016/0021-9991(85)90002-6)
- [75] P. P. Ewald. “Die Berechnung optischer und elektrostatischer Gitterpotentiale”. In: *Annalen der Physik* 369.3 (1921), pp. 253–287. DOI: [10.1002/andp.19213690304](https://doi.org/10.1002/andp.19213690304). URL: <https://doi.org/10.1002/andp.19213690304>
- [76] Arthur Kosmala et al. *Ewald-based Long-Range Message Passing for Molecular Graphs*. 2023. eprint: [arXiv:2303.04791](https://arxiv.org/abs/2303.04791)
- [77] Benjamin Trendelkamp-Schroer et al. “Estimation and uncertainty of reversible Markov models”. In: *The Journal of Chemical Physics* 143.17 (Nov. 2015), p. 174101. DOI: [10.1063/1.4934536](https://doi.org/10.1063/1.4934536). URL: <https://doi.org/10.1063/1.4934536>
- [78] Susanna Röblitz and Marcus Weber. “Fuzzy spectral clustering by PCCA+ application to Markov state models and data classification”. In: *Advances in Data Analysis and Classification* 7.2 (May 2013), pp. 147–179. DOI: [10.1007/s11634-013-0134-6](https://doi.org/10.1007/s11634-013-0134-6). URL: <https://doi.org/10.1007/s11634-013-0134-6>
- [79] Igor Mezić. “Spectral Properties of Dynamical Systems, Model Reduction and Decompositions”. In: *Nonlinear Dynamics* 41.1-3 (2005), pp. 309–325. DOI: [10.1007/s11071-005-2824-x](https://doi.org/10.1007/s11071-005-2824-x). URL: <https://doi.org/10.1007/s11071-005-2824-x>
- [80] Moritz Hoffmann et al. “Deeptime: a Python library for machine learning dynamical models from time series data”. In: *Machine Learning: Science and Technology* 3.1 (2021), p. 015009. DOI: [10.1088/2632-2153/ac3de0](https://doi.org/10.1088/2632-2153/ac3de0). URL: <https://doi.org/10.1088/2632-2153/ac3de0>