RESEARCH ARTICLE | NOVEMBER 16 2020

Coarse graining molecular dynamics with graph neural networks **I**

Brooke E. Husic (); Nicholas E. Charron (); Dominik Lemm (); Jiang Wang (); Adrià Pérez; Maciej Majewski 💿 ; Andreas Krämer 💿 ; Yaoyi Chen 💿 ; Simon Olsson; Gianni de Fabritiis; Frank Noé 💿 ; Cecilia Clementi 🔽 💿

() Check for updates

J. Chem. Phys. 153, 194101 (2020) https://doi.org/10.1063/5.0026133 **D** CHORUS



CrossMark

The Journal of Chemical Physics

Special Topic: Festschrift in honor of Yuen-Ron Shen

Submit Today





Coarse graining molecular dynamics with graph neural networks **I**

Cite as: J. Chem. Phys. 153, 194101 (2020); doi: 10.1063/5.0026133 Submitted: 21 August 2020 • Accepted: 27 October 2020 • Published Online: 16 November 2020



Brooke E. Husic,^{1,2,a)} Dicholas E. Charron,^{3,4} Dominik Lemm,⁵ Jiang Wang,^{4,6} Adrià Pérez,⁵ Maciej Majewski,⁵ Andreas Krämer,¹ Yaoyi Chen,^{1,7} Simon Olsson,¹ Gianni de Fabritiis,^{5,8,b} Frank Noé,^{1,4,6,9,c)} And Cecilia Clementi^{3,4,6,9,d}

AFFILIATIONS

¹Department of Mathematics and Computer Science, Freie Universität, Berlin, Germany

- ²Department of Chemistry, Stanford University, Stanford, California 94305, USA
- ³Department of Physics, Rice University, Houston, Texas 77005, USA
- ⁴Center for Theoretical Biological Physics, Rice University, Houston, Texas 77005, USA
- ⁵Computational Science Laboratory, Universitat Pompeu Fabra, PRBB, C/Dr. Aiguader 88, Barcelona, Spain
- ⁶Department of Chemistry, Rice University, Houston, Texas 77005, USA
- ⁷IMPRS-BAC, Max Planck Institute for Molecular Genetics, Berlin, Germany
- ^eInstitucio Catalana de Recerca i Estudis Avanats (ICREA), Passeig Lluis Companys 23, Barcelona, Spain
- ⁹Department of Physics, Freie Universität, Berlin, Germany

^{a)}Electronic mail: bhusic@fu-berlin.de

- ^{b)}Electronic mail: gianni.defabritiis@upf.edu
- ^{c)}Electronic mail: frank.noe@fu-berlin.de
- ^d Author to whom correspondence should be addressed: cecilia.clementi@fu-berlin.de

ABSTRACT

Coarse graining enables the investigation of molecular dynamics for larger systems and at longer timescales than is possible at an atomic resolution. However, a coarse graining model must be formulated such that the conclusions we draw from it are consistent with the conclusions we would draw from a model at a finer level of detail. It has been proved that a force matching scheme defines a thermodynamically consistent coarse-grained model for an atomistic system in the variational limit. Wang *et al.* [ACS Cent. Sci. 5, 755 (2019)] demonstrated that the existence of such a variational limit enables the use of a supervised machine learning framework to generate a coarse-grained force field, which can then be used for simulation in the coarse-grained space. Their framework, however, requires the manual input of molecular features to machine learn the force field. In the present contribution, we build upon the advance of Wang *et al.* and introduce a hybrid architecture for the machine learning of coarse-grained force fields that learn their own features via a subnetwork that leverages continuous filter convolutions on a graph neural network architecture. We demonstrate that this framework succeeds at reproducing the thermodynamics for small biomolecular systems. Since the learned molecular representations are inherently transferable, the architecture presented here sets the stage for the development of machine-learned, coarse-grained force fields that are transferable across molecular systems.

Published under license by AIP Publishing. https://doi.org/10.1063/5.0026133

I. INTRODUCTION

Technologies facilitating molecular dynamics (MD) simulations, such as distributed computing $^{\rm l-3}$ and bespoke hardware, $^{\rm 4}$

have made great strides in terms of the time- and length-scales accessible *in silico*. However, even the longest protein simulations still fail to reach total times exceeding milliseconds, and dedicated analysis methods are required to infer dynamics at longer timescales.^{5,6} In

the context of such limitations at full atomistic resolution, coarse graining provides a crucial methodology to more efficiently simulate and analyze biomolecular systems. In addition to the practical advantages that arise from more efficient sampling, coarse graining can also elucidate the physical components that play key roles in molecular processes.

Coarse graining is especially useful for analyzing structures and processes that reach beyond the length and timescales accessible to all-atom MD. Important examples include protein folding, protein structure prediction, and protein interactions.⁷ Some of the most-used coarse-grained models for such studies are structurebased models,⁸ MARTINI,^{9,10} CABS,¹¹ AWSEM,¹² and Rosetta.¹³ These models differ with respect to their potential energy function, parameterization approaches, and resolution, which, in combination, determine their efficiency, accuracy, and transferability. In the past decade, coarse-grained models have become increasingly powerful due to an unprecedented wealth of experimental reference data and computational capabilities. In this context, the development of more realistic architectures and modeling approaches is of prime importance.

In the field of computer science, advances in hardware and autodifferentiation software have enabled enormous progress in machine learning algorithms, including at the intersection of computation and the physical sciences.^{14,15} Crucial to the use of neural networks in the physical sciences is a consideration for the form the training data takes before it is input into the network. One strategy for representing molecules mathematically is through the use of graphs, whose nodes and edges intuitively correspond to atoms and bonds of (or interatomic distances within) a molecule, respectively. By performing multiple convolution operations on a graph, each node can influence other increasingly distant nodes. The use of graph neural networks^{16,17} in the molecular sciences is therefore a promising direction in a variety of applications, and graph convolutional architectures have been used to predict molecular¹ and material²² properties as well as atomic energies^{23,2} and forces.25

In this work, we combine the use of graph representations of molecules with a supervised neural network architecture in the coarse graining context. We consider coarse graining to be the process of reducing structural degrees of freedom to facilitate more efficient simulation with specific goals in mind (e.g., reproducing system thermodynamics). Coarse graining can be implemented with a "top down" or "bottom up" approach, although other categories can be determined and strategies can be combined.²⁶ In a "top down" scheme, coarse graining frameworks are explicitly designed to reproduce certain macroscale emergent properties.²⁶ In a "bottom up" framework, which we consider here, implementations focus instead on reproducing specific features from a more detailed model.

The latter involves (i) a mapping from the entities in a finegrained (e.g., atomistic) representation to a smaller set of interaction sites, often called "beads," and (ii) a physical model (i.e., Hamiltonian function) for the coarse-grained system comprising those beads. Good choices for the mapping and model will lead to more efficient simulation while preserving the biophysical properties of interest to the researcher. Modern machine learning techniques have been recently employed to learn both the mapping^{27,28} and the model^{25,29–32} components of bottom up coarse graining. In the present contribution, we focus on the coarse graining model and employ a bottom up "force matching" scheme formulated as a supervised machine learning problem to reproduce the thermodynamics of small biomolecular systems. Particularly, we modify the architecture of the recently introduced CGnet framework³¹ such that the molecular features it requires are *learned* via graph convolutional neural networks instead of hand-selected as in the original formulation. By leveraging the inherently transferable SchNet scheme^{24,33} to learn features, we render the entire CGnet framework transferable across molecular systems.

Our goal in this paper is to present the theory underlying CGSchNet—our new transferable coarse graining architecture—and demonstrate its success on learning the thermodynamics of individual biomolecular systems. We find that our new protocol produces more accurate free energy surfaces in comparison to the use of handselected features, is more robust to hyperparameter choices, and requires less regularization. Presented alongside a machine learning software package that implements the methods introduced, the current contribution sets out a framework for the machine learning of transferable, coarse-grained molecular force fields and demonstrates its application to a small peptide system and the miniprotein chignolin.³⁴ The practical application of the methods described herein to larger protein systems, particularly those characterized by meaning-ful tertiary structure, remains an open challenge that will be explored in the future work.

II. THEORY

Force matching was pioneered in the atomistic setting in which forces obtained from an inexpensive calculation are matched to forces computed at a more computationally expensive level of theory (i.e., quantum) via an optimization scheme.³⁵ The method was later adapted by the coarse graining community; in that context, coarse-grained representations are sought such that the forces computed from the coarse-grained energy function for a given configuration match the average forces on corresponding atomistic representations.³⁶

Because coarse graining away degrees of freedom entails that multiple atomistic structures will correspond to the same coarsegrained configuration, it is impossible to obtain zero error during force matching in the coarse graining context. However, it can be proved that the coarse graining model that matches the *mean* forces yields the correct thermodynamics and that the objective is variationally bounded from below by a value that necessarily exceeds zero.

In Sec. II, we overview the major advances that enable the present contribution. The practically inclined reader may proceed directly to Sec. III, where we discuss the CGnet architecture and introduce this work's methodological contribution: namely, the incorporation of learnable molecular features into CGnet via the use of continuous filter convolutions on a graph neural network (i.e., SchNet^{24,33}). We will see in Sec. III that the scheme we introduce here enables, at least in principle, a coarse graining architecture that is transferable across system size and sequence. The practical use of this architecture to learn a force field in a transferable context will be addressed in the future work.

A. Force matching

Consider an all-atom dataset of coordinates and corresponding forces, which we have obtained using a high level calculation (e.g., *ab initio*). We denote each three-dimensional structure $\mathbf{r}_i \in \mathbb{R}^{3N}$, i = 1, ..., M, and the forces $\mathbf{F}(\mathbf{r}_i) \in \mathbb{R}^{3N}$, where N is the number of atoms in the system. Now consider a trial energy function $\hat{V}(\mathbf{r}_i; \Theta)$, which takes as arguments an atomistic configuration \mathbf{r}_i and any parameters Θ . We would like to use \hat{V} to predict the forces on every \mathbf{r}_i —presumably in a more efficient way—by taking its negative derivative. We can write the "force matching" problem of comparing the two sets of atomistic forces as

$$L(\mathbf{R}; \mathbf{\Theta}) = \frac{1}{3MN} \sum_{i=1}^{M} \|\underbrace{\mathbf{F}(\mathbf{r}_i)}_{\text{"True"}} + \underbrace{\nabla_{\mathbf{r}_i} \hat{V}(\mathbf{r}_i; \mathbf{\Theta})}_{\text{predicted}} \|^2, \quad (1)$$

where \mathbf{R} is the set of all M sampled atomistic configurations.

The objective (1) was introduced by Ercolessi and Adams to analyze *ab initio* simulations of elemental aluminum.³⁵ The authors highlight the method's need to accommodate invariant properties of the system and discuss the requirement of a variety of geometries, physical conditions, and system identities in **R** if the learned potential is to be transferable across conformation, thermodynamic, or chemical space, respectively. Subsequent work has derived analytical approaches to this scheme in the context of liquids.^{37,38}

A decade later, Izvekov and Voth introduced the multiscale coarse graining (MS–CG) method, a groundbreaking advance that adapts force matching to the coarse graining context.^{36,39} The MS–CG framework involves two steps: first, atoms are aggregated into "interaction sites" according to a linear mapping from N atoms to n interaction sites (henceforth "beads"),

$$\mathbf{x}_i = \mathbf{\Xi} \mathbf{r}_i \in \mathbb{R}^{3n},\tag{2}$$

where \mathbf{x}_i is the coarse-grained representation with n < N beads and the matrix $\Xi \in \mathbb{R}^{3n \times 3N}$ effectively performs a clustering from the original atoms to the beads. Then, force matching is performed between a transformation of the atomistic forces and a set of predicted coarse-grained forces. This procedure thereby creates a "multiscale" link between the all-atom and coarse-grained representations.³⁶

Consider a *coarse-grained* energy function $U(\mathbf{x}; \boldsymbol{\Theta})$. Let us say that we have a set of M coarse-grained configurations that we have obtained by applying (2) to every configuration $\mathbf{r}_i \in \mathbf{R}$. To calculate the forces on the beads, we then take the negative derivative of U with respect to the reduced coordinates; in other words, we evaluate

$$-\nabla U(\mathbf{\Xi}\mathbf{r}_i;\mathbf{\Theta}) = -\nabla_{\mathbf{x}_i} U(\mathbf{x}_i;\mathbf{\Theta}) \in \mathbb{R}^{3n}$$

for each configuration *i*. From here, we have all the ingredients to write down the adaptation of (1) to the MS–CG method,

$$L(\mathbf{R}; \boldsymbol{\Theta}) = \frac{1}{3Mn} \sum_{i=1}^{M} \| \underbrace{\Xi_{\mathbf{F}} \mathbf{F}(\mathbf{r}_i)}_{\text{Atomistic forces}} + \underbrace{\nabla U(\Xi \mathbf{r}_i; \boldsymbol{\Theta})}_{\text{predicted from}} \|^2, \qquad (3)$$

where $\Xi_F F$ is the *instantaneous* coarse-grained force (also called the *local mean force*); that is, the projection of the atomistic force into the coarse-grained space. A general expression for the force projection⁴⁰ is $\Xi_F = (\Xi\Xi^{\top})^{-1}\Xi$. Other choices for the mapping Ξ_F are possible and used for coarse graining.⁴¹

In principle, the coarse-grained energy $U(\mathbf{x})$ that is *exactly* thermodynamically consistent with the atomistic energy $V(\mathbf{r})$ can be expressed analytically as

$$U(\mathbf{x}) = -k_B T \ln p^{CG}(\mathbf{x}) + \text{Constant}, \qquad (4)$$

where k_B is Boltzmann's constant and *T* is the absolute temperature. The function p^{CG} is the marginal probability density,

$$p^{\rm CG}(\mathbf{x}') = \frac{\int_{\mathcal{R}} \exp\left(-\frac{V(\mathbf{r})}{k_B T}\right) \delta(\mathbf{x}' - \mathbf{\Xi}\mathbf{r}) d\mathbf{r}}{\int_{\mathcal{R}} \exp\left(-\frac{V(\mathbf{r})}{k_B T}\right) d\mathbf{r}},$$
(5)

where \mathcal{R} is the set of all possible atomistic configurations. Since we are concerned with all theoretically accessible structures and (thus) employ an integral formulation, we have dropped the subscripts *i* with the understanding that \mathbf{x}' and \mathbf{r} now refer to infinitesimally small regions of their respective configuration spaces. \mathbf{x}' is distinguished from \mathbf{x} to emphasize that (5) is substituted into (4) as a function, not a number.

The coarse-grained energy function (4) is called the potential of mean force (PMF) and is an analog of the atomistic potential energy function. Via (5), it is a function of weighted averages of energies of atomistic configurations. For a given coarsegrained structure \mathbf{x}' , in (5), we evaluate whether every possible $\mathbf{r} \in \mathcal{R}$ maps to \mathbf{x}' . We expect multiple atomistic configurations **r** to map to \mathbf{x}' due to the reduction in degrees of freedom that results from structural coarse graining (n.b., this means the PMF is, in fact, a free energy, as it contains entropic information^{2b}). In these cases, the Dirac delta function in (5) returns one, and the contribution of that atomistic configuration to the marginal probability distribution is a function of its Boltzmann factor. If r does not map to \mathbf{x}' , then the evaluation of the delta function, and thus the contribution of that atomistic structure to the free energy of \mathbf{x}' , is zero. The denominator of the right-hand side of (5) is the all-atom partition function, which serves as a normalization factor

To calculate the forces on our coarse-grained beads, we must take the gradient of (4). However, since we cannot exhaustively sample \mathcal{R} , (5) is intractable, and we must approximate U instead. One way to approximate U is to employ force matching—that is, by minimizing expression (3)—as will be described in Sec. II B. Another method, which we do not discuss in this report, is through relative entropy,⁴² whose objective is related to that of force matching.^{26,43}

B. Coarse graining as a supervised machine learning problem

In 2008, Noid *et al.*⁴¹ formalized the notion of thermodynamic consistency and established the conditions under which it is guaranteed by the MS–CG approach: specifically, thermodynamic consistency is achieved when the coarse-grained coordinates are a linear combination of the all-atom coordinates [cf. (2)] and that the equilibrium distribution of the coarse-grained configurations is equal to the one implied by the equilibrium distribution of the atomic configurations [cf. (4)]. Noid *et al.* then proved that, under certain restrictions of the coarse-grained mapping, the coarse-grained potential that achieves thermodynamic consistency at a given temperature is unique [up to an additive constant, cf. (5)].⁴¹

The authors define an error functional that is (uniquely) minimized for the thermodynamically consistent coarse-grained force field.^{41,103} This framework provides the variational principle underlying the MS–CG method. Consequently, a variational approach can be used to search for the consistent coarse-grained force field.⁴⁴ The variational principle entails that we can refer to (1) and (3) as "loss functions" because they return a scalar that assumes a minimum value on the optimal model. In the recent reports from both Wang *et al.*³¹ and Wang and Gómez-Bombarelli,²⁶ this fact is leveraged to formulate coarse graining via force matching as a supervised or semi-supervised machine learning problem, respectively. Here, we build upon on the supervised learning case introduced in Ref. 31 as *CGnet*.

In their study, Wang *et al.* presented several crucial contributions.³¹ First, they decompose the error term implied by (3) into three physically meaningful components; namely, bias, variance, and noise. Second, the authors introduce *CGnet*: a neural network architecture designed to minimize the loss in (3). Once a CGnet is trained, it can be used as a force field for new data points in the coarse-grained space while enforcing known properties of the system such as symmetries and equivariances (see Sec. III C). Third, Wang *et al.* augmented their initial framework to introduce regularized CGnets.³¹ Regularized CGnets avoid catastrophically wrong predictions observed in their "unregularized" counterparts by introducing the calculation of prior energy terms before training. This adjustment means that, instead of learning the forces directly, the neural network learns a *correction* to the prior terms in order to match the atomistic forces.

Using regularized CGnets (henceforth, we assume all CGnets are regularized) on two peptide systems, the authors demonstrated effective learning of coarse-grained force fields that could not be obtained with a few-body model approach.³¹ It is from this base-line that we present CGSchNet, an augmentation of the CGnet methodology.

III. METHODS

In the quantum community, supervised machine learning has been used to predict energies on small molecules through a variety of approaches.^{23,33,45–59} In particular, the SchNet architecture is based on the use of continuous filter convolutions and a graph neutral network.^{17,24,33} SchNet is a scalable, transferable framework that employs representation learning to predict the properties and behavior of small organic molecules. In the vein of the original force matching procedure of Ercolessi and Adams,³⁵ SchNet has also been used to predict forces on atomic data from a quantum mechanical gold standard.³³

In Sec. III A, we briefly overview the CGnet scheme upon which we base the method introduced in this work. Then, in Sec. III B, we describe SchNet and introduce our adaptation of SchNet to the coarse graining problem by incorporating it into a CGnet to create a hybrid "CGSchNet" architecture. The original implementation of CGnet is not transferable across different systems due to its reliance on hand-selected structural features.³¹ We recognized that SchNet could be leveraged as a subcomponent of CGnet in order to *learn* the features, thereby converting CGnet—i.e., force matching via supervised machine learning—to a transferable framework for the first time.

A. Original CGnet architecture

For both CGnet and CGSchNet, our training data comprise an MD simulation that has already been performed and for which the atomistic forces have been retained or recomputed. Both the configurations and the forces are in \mathbb{R}^{3N} space for *N* atoms. We then determine our mapping matrix Ξ and use it to prepare our input data (coarse-grained structures) and labels (atomistic forces mapped to the coarse-grained space), both will be in \mathbb{R}^{3n} for *n* beads [recall (2)].

While the mapping is permitted to be more general, in our work, we restrict it to the special case where the matrix Ξ contains zeroes and ones only. With this choice of mapping, the projection of the forces in (3) becomes simply $\Xi_F = \Xi$. Our mapping thus "slices" the original atomic configuration such that the corresponding coarse-grained representation comprises a subset of the original atoms. For example, a useful mapping might retain only protein backbone atoms or α -carbons.

To construct a CGnet, the structural data are preprocessed such that it is represented by features with the desired properties. Wang *et al.* used a set of distances, planar angles, and torsional angles.³¹ In the present work, on the other hand, instead of using hand-selected structural features, we require only distances and bead identities from which features are *learned*; this is described in Sec. III B.

For their regularized implementation, Wang *et al.* used up to two types of prior terms in CGnets.³¹ The first is a harmonic prior on selected distances (i.e., bonds or pseudobonds) and angles. The second is a repulsion prior that can be used on nonbonded distances. Correspondingly, these priors are defined as follows for a given feature f_i calculated from the data (e.g., a particular distance):

$$U_i^{\text{harmonic}}(f_i) = \frac{k_B T}{2 \text{Var}[f_i]} (f_i - \mathbb{E}[f_i])^2,$$
(6a)

$$U_i^{\text{repulsion}}(f_i) = \left(\frac{\sigma}{f_i}\right)^c.$$
 (6b)

The constants in (6b) can be determined through cross-validated hyperparameter optimization as in Ref. 31. The prior energy is the sum of each prior term for all relevant features f_i . In principle, any

scalar function of protein coordinates can be used to construct a prior energy term.

The original CGnet uses a fully connected network to learn corrections to the prior energy.³¹ Crucially, the last layer of the network returns a scalar output. Because of this single node bot-tleneck structure, the resulting coarse-grained force field will be curl-free and is therefore guaranteed to conserve energy.^{31,50} Since all the steps described are differentiable, we can use an autod-ifferentiation framework such as PyTorch⁶⁰ to take the derivative of the energy with respect to the original (coarse-grained) spatial coordinates via backpropagation. This derivative corresponds to the predicted forces on the coarse-grained beads in \mathbb{R}^{3n} , which can then be compared to the known forces on the training coordinates.

B. Replacing structural features with graph neural networks

Wang *et al.* show that CGnets constructed upon hand-selected structural features produce machine-learned force fields yielding accurate free energy surfaces.³¹ The model architecture is found to be somewhat sensitive to various hyperparameters and required individual tuning for each system (see Fig. 5 in the work of Wang *et al.*³¹). Furthermore, a new system will, in general, require retraining because the feature size is fixed according to the system geometry.

In the present contribution, we replace the fixed structural features employed in the original CGnet formulation (i.e., distances, angles, and torsions)³¹ with *learned* features computed using continuous filter convolutions on a graph neural network (SchNet^{24,33}). The SchNet architecture thereby becomes a subunit of CGnet with its own, separate neural network scheme; we refer to this hybrid architecture as *CGSchNet*.

The term graph neural network was introduced by Battaglia et al.¹⁷ as a generalization for networks operating on graph structures, including, but not limited to, graph convolutional networks¹⁶ and message passing networks.²⁰ These networks have in common that they have a notion of n nodes V connected by edges \mathcal{E} in a graph structure. In each neural network layer, information is passed between nodes and representations of the nodes and/or edges are updated. The various types of graph neural networks differ according to whether there are node updates, edge updates, or both, as well as by how functions are shared across the network and how the network output is generated. A fairly general formulation of a graph neural network with node updates is as follows: each node *i* is associated with an initial node representation $\mathbf{h}_{i}^{(0)}$; in other words, $\mathbf{h}_{i}^{(0)}$ is a vector that represents the type or identity of the node. In each layer of the neural network, the node representations are updated according to

$$\mathbf{h}_{i}^{(t+1)} = f[(\mathbf{h}_{i}^{(t)}), (\mathbf{e}_{ij})], \tag{7}$$

where i, j = 1, ..., n, e_{ij} are edge features defined for all edges in \mathcal{E} and f is a trainable neural network function. After T such layers, an output is generated,

$$\mathbf{o} = g[(\mathbf{h}_i^{(T)}), (\mathbf{e}_{ij})]. \tag{8}$$

In the present contribution, we make the following choices:

- Because multi-body interactions are important for the coarse graining problem, edges are defined between all beads (or all beads within a specified cutoff radius).
- 3. The edge features e_{ij} are taken to be the distances between beads, implying translational and rotational invariance of the network output.
- 4. The update function f in (7) is chosen to be a continuous convolution update as in SchNet.²⁴
- 5. The entire trainable part of a CGnet³¹—in this case a beadwise multilayer perceptron/dense neural network—becomes the output function g in (8). Because this output is beadwise, the learnable coarse grain energy is invariant with respect to permutations of identical beads.
- 6. The output **o** is a scalar; namely, the coarse-grained energy before the addition of the prior energy term.

Below, we describe the SchNet updates in more detail and how to incorporate SchNet into CGnet to create CGSchNet.

1. Learning molecular representations with SchNet

One key motivating factor for the original development of SchNet is that, unlike the images and videos that comprise the datasets for much of modern machine learning, molecular structures are not restricted to a regular grid. Therefore, Schütt *et al.* introduced continuous-filter convolutions to analyze the structures of small molecules with the goal of predicting energies and forces according to a quantum mechanical gold standard.²⁴ This development builds upon previous work in predicting atomic properties directly from structural coordinates.^{20,23,50}

SchNet is a graph neural network where the nodes correspond to particles embedded in three-dimensional space and the convolutional filters depend on interparticle distances, which preserves invariances expected in the system.^{17,24} While SchNet was originally used to predict quantum-chemical energies from atomistic representations of small molecules, here, we employ it to learn a feature representation that replaces the hand-selected features in a CGnet for the purpose of predicting the coarse-grained energy on the coarse-grained bead coordinates \mathbf{x}_i .

As in other graph neural networks, SchNet learns feature vectors on the nodes (here, coarse-grained beads). The initial node features at the input are called node or bead embeddings $\mathbf{h}_i^{(0)}$, which are given by trainable, shared vectors with d_h dimensions ("Embeddings" in Fig. 1),

$$\mathbf{h}_i^{(0)} = \mathbf{a}_{k(i)}.$$

Here, k(i) is a lookup table that maps the bead index *i* to its type *k*. In the present applications, we use nuclear charges (capped alanine) or amino acid identities (chignolin) as bead types. The bead embeddings are shared among beads of the same type and are optimized during training. Crucially, this entails that SchNet *learns* a molecular representation, which avoids the common paradigm of fixed, heuristic feature representations.

Next, we describe how bead representations are updated ("Interaction block" and "cfconf" in Fig. 1). In our current architecture, these updates are implemented as in the original SchNet unless noted otherwise (see Refs. 24 and 33 for details).

ARTICLE





In each interaction layer, we perform a continuous convolution between beads. For this, the inter-bead distances $|\mathbf{x}_j - \mathbf{x}_i|$ are featurized using radial basis functions **e**, e.g., one-dimensional Gaussians centered at different distances. These featurized distances serve as the input to a filter-generating neural network w that maps the featurized distance input $\mathbf{e}(|\mathbf{x}_j - \mathbf{x}_i|)$ to a d_h -dimensional filter. This filter is applied to the bead representations \mathbf{h}_i as follows ("cfconf" in Fig. 1):

$$\mathbf{z}_{i}^{(t)} = \sum_{j} w^{(t)} \left[\mathbf{e} (|\mathbf{x}_{j} - \mathbf{x}_{i}|) \right] \cdot b^{(t)} (\mathbf{h}_{i}^{t}).$$
(10)

Here, w and b are trainable functions and \cdot is element-wise multiplication. As in the original SchNet implementation,²⁴ w is a dense neural network and b is a beadwise linear layer. The sum in (10) is taken over every bead j within the neighborhood of bead i, which can be all other beads in the system or a subset thereof if a finite neighborhood is specified. Even when interactions are limited to particles within a cutoff radius, a sequence of multiple interaction layers will eventually allow all particles to be interacting and therefore be able to express complex multi-body interactions.

In each layer, the bead representations are updated in interaction blocks, each of which comprises a residual update of the bead representation via a nonlinear function of the continuous convolution outputs $\mathbf{z}_{i}^{(t)}$ ("interaction block" in Fig. 1),

$$\mathbf{h}_{i}^{(t+1)} = \mathbf{h}_{i}^{(t)} + g^{(t)}(\mathbf{z}_{i}^{(t)}).$$
(11)

The residual update step is an "additive refinement" that prevents gradient annihilation in deep networks.⁶¹ As described by Schütt *et al.*,^{24,33} the trainable function *g* involves beadwise linear layers and a nonlinearity. Instead of the softplus nonlinearity used in the original SchNet,²⁴ here, we use the hyperbolic tangent.

Following the last interaction layer, we must choose an output function (8). As in the original SchNet implementation, the output of the final SchNet interaction block is input into a beadwise CGnet multilayer perceptron/dense network. An important feature of transferability is permutation invariance of beads with identical type. In the context of coarse graining, this means the contribution of a bead to the coarse-grained energy should depend on its location in the molecular graph, but not at the index this bead is positioned in the input. SchNet layers are permutationequivariant, i.e., any exchange of the input representations $\mathbf{h}_{i}^{(0)}$ will correspond to the same exchange of the learned representations $\mathbf{h}_{i}^{(T)}$. In order to obtain permutation-invariant energies (as in the original SchNet publications^{24,33}), the beadwise output CGnet network contracts down to a scalar energy prediction that is then summed over all beads to yield the total learnable part of the coarsegrained energy. It is important to note that the models used in this study employ priors that are not permutation invariant, and so, the non-learnable part of the coarse-grained energy (i.e., the prior terms) breaks permutation invariance in the model overall. The development of permutation invariant priors is left for future work.

In the present paper, we do not use CGSchNet in a transferable manner, rather demonstrate its capabilities when trained on individual molecular systems as a foundation for future work. For this reason, here we use a (regularized) beadwise CGnet as the output function; i.e., a beadwise multilayer perceptron/dense neural network at whose output the learned part of the coarsegrained energy is predicted. In doing so, the SchNet interaction layers learn the input representation for a beadwise CGnet, and the beadwise CGnet "fine-tunes" the bead energies predicted by SchNet.

2. CCSchNet: A transferable architecture for coarse graining

CGnet as originally presented is incapable of learning a transferable coarse-grained force field due to its reliance upon systemspecific structural features.³¹ Since SchNet is inherently a transferable framework, learning CGnet features using SchNet enables the transferability of the entire CGnet architecture across molecular systems. Here, we present the advance of incorporating SchNet^{24,33} into CGnet to replace hand-selected features with machine-learned ones. In CGSchNet, instead of predetermined structural features i.e., distances, angles, and torsions—a SchNet is used instead, enabling the model to learn the representation itself (see Fig. 1). By replacing fixed-size geometric features with SchNet, we obtain a more flexible representation that both scales better with system size and is amenable to a transferable architecture.³³ While angles and torsions may still be included in the prior energy terms, they are no longer propagated through any neural networks.

The use of SchNet requires us to provide not only structural coordinates but also a type for every bead. In the original (i.e., non-coarse graining) implementation for systems at atomic resolution, the types are atomic numbers.^{23,24,33} In the new context presented here (i.e., leveraging SchNet for coarse graining), we may specify coarse-grained bead types—effectively, chemical environments; however, we deem appropriate for the system under study; for example, amino acid identities may be used.

One can view the typing requirement of the SchNet architecture as the mechanism for incorporating physical or chemical intuition about the system into the model, as opposed to fixed structural features. The benefit of the SchNet choice is that it enables an architecture that is transferable across size and sequence space because a set of embeddings can apply to multiple systems with the same components (e.g., atoms as in SchNet²⁴ or amino acid types in proteins), whereas hand-selected structural features are not transferable across different systems.

Finally, we note that we train CGSchNet with the coarsegrained force matching loss (3), which compares our predicted forces to the known forces from the training set. Unlike in the original SchNet formulation,²⁴ we cannot straightforwardly incorporate an additional "energy matching" term into the coarse graining framework. This is because we do not have labels for the coarsegrained free energies: these energies are defined by an integral over all microscopic configurations associated with the same coarsegrained configuration [cf. (4)], and this integral cannot be solved exactly.

C. Coarse-grained simulations

A trained CGSchNet can be used as a force field to simulate the system in the coarse-grained space. Specifically, Langevin dynamics^{62,63} are employed to propagate coarse-grained coordinates \mathbf{x}_t forward in time according to

$$\frac{\partial^2 \mathbf{x}_t}{\partial t^2} = -\mathbf{M}^{-1} \nabla U(\mathbf{x}_t) - \gamma \frac{\partial \mathbf{x}_t}{\partial t} + \sqrt{2k_b T \gamma} \mathbf{M}^{-\frac{1}{2}} \mathbf{W}(t), \qquad (12)$$

where the diagonal matrix **M** contains the bead masses, γ is a collision rate with units ps⁻¹, and **W**(*t*) is a stationary Gaussian process with $\langle W(t) \rangle = 0$ and $\langle W(t)W(t') \rangle = \delta(t - t')$, where $\langle \cdot \rangle$ is the mean. In practice, we integrate (12) using a "BAOAB" Langevin integrator,⁶⁴ and the integral of W(t) is a Wiener process. A special case of Langevin dynamics is the so-called "overdamped" Langevin dynamics, also referred to as Brownian dynamics. Overdamped Langevin dynamics lack inertia. After setting the acceleration to zero, dividing both sides by γ and rearranging terms, in the overdamped case, (12) becomes

$$\frac{\partial \mathbf{x}_t}{\partial t} = -\frac{\mathbf{D}}{k_B T} \nabla U(\mathbf{x}_t) + \mathbf{D}^{\frac{1}{2}} \sqrt{2} \mathbf{W}(t), \qquad (13)$$

where the diffusion matrix $\mathbf{D} \equiv \mathbf{M}^{-1}k_B T/\gamma$. Although \mathbf{D} contains a notion of mass, we note that propagating coarse-grained dynamics via (13) does not actually require bead masses, since the product $\mathbf{M}\gamma$ can be considered without separating its factors. Wang *et al.*³¹ used exclusively (13) with the Euler method to simulate dynamics from CGnets, with a constant diffusion matrix proportional to the identity matrix. In both formulations, the noise term is intended to indirectly model collisions—e.g., from and among solvent particles—that are not present in the coarse-grained coordinate space. Since Langevin dynamics depend only on the coordinates (and, unless overdamped, velocities) of the previous time step, these simulations can easily be run in parallel from a set of initial coordinates. The resulting coarse-grained simulation dataset can then be used for further analysis, as will be shown in Sec. IV.

IV. RESULTS

A. Capped alanine

Capped alanine—often referred to as alanine dipeptide for its two peptide bonds—is a common benchmark for MD methods development because the heavy-atom dynamics of the central alanine are completely described by the dihedral (torsional) angles ϕ and ψ (see Fig. 2). We performed a single 1- μ s all-atom, explicit solvent MD simulation for capped alanine and saved the forces to use for CGSchNet training (see Ref. 31 and Sec. A of the supplementary material). We can visualize the occupancies of backbone angle conformations by creating a histogram of the data on $\phi \times \psi$ space and visualizing the populations of the histogram bins. This is called a Ramachandran map and is depicted in Fig. 4(a) for the atomistic simulation using a 60 × 60 regular spatial discretization.

As an initial benchmark of the CGSchNet method, we aim to learn a force field for a coarse-grained representation of capped alanine such that we can reproduce its heavy-atom dynamics using a trained CGSchNet instead of a more expensive explicit solvent allatom MD simulation. For our coarse-grained mapping, we select the backbone heavy atoms C– $[N-C_{\alpha}-C]_{Ala}$ –N as well as the alanine C_{β} for a total of six beads.⁶⁵ We use atomic numbers for the bead embeddings as in the original SchNet formulation.²⁴ A CGSchNet



FIG. 2. Capped alanine in water. The six shaded atoms are the ones preserved in the coarse-grained representation. The ϕ and ψ dihedral angles completely describe the central alanine's heavy-atom dynamics.

is trained on the coordinates and forces of the all-atom simulation depicted in Fig. 4(a). The learning procedure involves a hyperparameter selection routine and the training of multiple models under fivefold cross-validation for each hyperparameter set (see Sec. B of the supplementary material).

Once a final architecture has been selected, the trained model can serve as a force field in the coarse-grained space, i.e., by predicting the forces on a set of input coarse-grained coordinates. Along with an integrator, predicted forces can be used to propagate coarsegrained coordinates forward in time (recall Sec. III C). This procedure (i.e., force prediction with CGSchNet followed by propagation with an integrator) is iterated until a simulation dataset of the desired duration has been obtained. Since we employ fivefold crossvalidation during the model training procedure, we have five trained CGSchNet models with a common architecture at hand. To perform our coarse-grained simulation, we simultaneously predict the forces on each set of input coordinates from all five trained networks, and the mean force vector is used to propagate Langevin dynamics according to (12).

To facilitate sampling, 100 coarse-grained simulations of length 200 ns each are performed in parallel from various starting positions in Ramachandran space (see Sec. C and Fig. S3 of the supplementary material). The time series of the ϕ and ψ values for two of the trajectories that feature transitions among the major basins are plotted in Fig. 3. The same trajectories are also overlaid on the two-dimensional energy surface in Fig. S4 of the supplementary material.

Free energy surfaces resulting from the coarse-grained simulation dataset are presented in Fig. 4(b). We can see qualitatively that the two-dimensional free energy surface from the CGSchNet simulation captures the same basins as the surface calculated from the baseline all-atom simulation. In the one-dimensional free energy surfaces, we see that the barriers are well-approximated by the CGSchNet simulation data.



FIG. 3. Two 100-ns trajectories simulated using a CGSchNet trained on atomistic data of capped alanine. The orange and magenta lines represent the value of the dihedral angles ϕ and ψ , respectively, over the course of each simulation. Relatively steep changes in the *y*-direction indicate transitions among basins; one can see that both trajectories feature multiple transitions in both reaction coordinates. A moving average of 250 simulation frames is used to smooth the darker curves.

To calibrate our understanding of the CGSchNet simulation dataset's relationship to the baseline atomistic simulation dataset, we create a set of new systems by perturbing the Cartesian coordinates of the latter with noise distributed as $\mathcal{N}(0, \sigma^2)$ for $\sigma \in \{0, 0.01, 0.02, \ldots, 0.30\}$ Å. From the perturbed Cartesian coordinates, the new ϕ and ψ dihedrals are calculated and assigned to the same 60×60 regularly spaced bins in Ramachandran space. Examples of the perturbed free energy surfaces are shown in Figs. 4(c)-4(e) for $\sigma = 0.1$ Å, 0.2 Å, and 0.3 Å, respectively. We see that the surfaces become smeared and the free energy barriers are reduced with increasing noise.

This ensemble of perturbed simulation datasets enables us to understand the CGSchNet-produced simulation in the context of the baseline atomistic simulation. To quantify the relationship between two distributions, we can use the Kullback–Leibler (KL) divergence⁶⁶ and a mean squared error (MSE) formulation. The KL divergence is defined for discrete distributions as

$$-\sum_{i}^{m} p_{i} \ln \frac{q_{i}}{p_{i}}, \forall p_{i} \ge 0,$$
(14)

where *p* and *q* are the "reference" and "trial" distributions, respectively, and *m* is the number of bins in each discrete distribution. In this case, *p* and *q* represent the normalized bin *counts*. The index *i* returns the normalized count from the *i*th bin of a 60 × 60 regular discretization of $\phi \times \psi$ space. The distribution obtained from the baseline atomistic simulation always serves as the reference. The mean squared error used here is

$$\frac{1}{m'} \sum_{i}^{m'} (P_i - Q_i)^2 \ni p_i q_i > 0,$$
(15)

where p_i and q_i remain the normalized bin *counts* and P_i and Q_i represent the corresponding discrete distributions of bin *energies* calculated as $P_i = k_B T \log p_i$ for Boltzmann's constant k_B and absolute temperature T. When no count is recorded for a bin in either p_i or q_i , those bins are omitted from the mean. m' represents the number of bins in which $p_i q_i > 0$ (i.e., both have finite energies).⁶⁷ 50 different trials are performed at different random seeds for the full set of noise scales [i.e., at each noise scale for a given trial, values are drawn from $\mathcal{N}(0, \sigma^2) \in \mathbb{R}^{M \times 3n}$, where M is the length of the trajectory dataset and n is the number of coarse-grained beads]. Within each trial, at each noise scale value σ , the KL divergence and MSE are calculated. The results are presented in the left-hand side plot of Fig. 5.

We see in Fig. 5 that as the noise increases, both divergence metrics also increase. The dashed lines in Fig. 5 show us that the error on the CGSchNet simulation dataset is approximately comparable to the corresponding error on the perturbed dataset with noise scale $\sigma = 0.1$ Å [Fig. 4(c)].⁶⁸ Upon qualitative comparison of the free energy surfaces, however, the former has more visual fidelity to the baseline surface in Fig. 4(a) than to the broader spread seen (and expected) in the latter. We know that coarse graining can result in increased population in transition regions that are rarely visited in an all-atom model; this is what we observe in Fig. 4(b). As a corollary, we do *not* expect coarse graining to result in the absence of states known to exist in the baseline system.



FIG. 4. Two- and one-dimensional free energy surfaces for five capped alanine datasets. From left to right, datasets are the baseline all-atom capped alanine simulation (a), the coarse-grained CGSchNet simulation produced for analysis (b), and datasets generated from perturbations of the original Cartesian coordinates of the baseline dataset drawn from noise distributed as $\mathcal{N}(0, \sigma^2)$ for $\sigma = 0.1$ Å (c), 0.2 Å (d), and 0.3 Å (e). To create each two-dimensional surface, the ϕ and ψ Ramachandran angles are calculated from the spatial coordinates and discretized into 60×60 regularly spaced square bins. The bin counts are converted to free energies by taking the natural log of the counts and multiplying by $-k_BT$; the color scale is the same in all five two-dimensional surfaces and darker color represents lower free energy (i.e., greater stability). To obtain the one-dimensional ϕ and ψ landscapes, free energies are calculated for 60 regularly spaced bins along the reaction coordinate. The shaded region always represents the baseline dataset, and the bold line represents the dataset indicated in the inset title.

B. Chignolin

The CLN025 variant of chignolin is a 10-amino acid miniprotein³⁴ featuring a β -hairpin turn in its folded state (Fig. 6). Due to its fast folding, its kinetics have been investigated in several MD studies.⁶⁹⁻⁷⁴ Our training data are obtained from an atomistic



FIG. 5. The Kullback–Leibler divergence (left) and mean squared error (right) are calculated between the baseline capped alanine dataset and the datasets obtained from perturbations to the baseline simulation at noise scale values of $\sigma \in \{0, 0.01, 0.02, \ldots, 0.30\}$ where the former is the reference distribution and the latter is the trial distribution. This procedure is performed 50 times with different random seeds; both plots show the superposition of those 50 lines. The colored horizontal dashed line shows the value of the metric when comparing the CGSchNet simulation to the baseline, and the black vertical dashed line indicates the noise scale σ that return the closest value for that metric. The Kullback–Leibler (KL) divergence is computed for normalized bin counts in $\phi \times \psi$ space, and the MSE is computed for the energies of those bins as described in the matrix.



FIG. 6. The miniprotein chignolin. The α -carbon backbone is visualized in opaque black, and these ten atoms are the only beads preserved in the coarse-grained representation. The atomistic system is also solvated, although the water molecules are not shown here.

simulation of chignolin in explicit solvent for which we stored the forces (see Ref. 31 and Sec. A of the supplementary material). To build our CGSchNet model, we retain only the ten α -carbons for our coarse-grained beads. For the SchNet embeddings, we assign each amino acid type its own environment with a separate designation for the two terminal tyrosines. After determining hyperparameters for our CGSchNet model, we simulate chignolin in the coarse-grained space using Langevin dynamics (12) as in Sec. IV A. The procedures for CGSchNet training and simulation are similar to those used for capped alanine and are described in Secs. B and C of the supplementary material.

Given our CGSchNet simulation data, we are interested not only in performing a similar analysis to the one in Sec. IV A for alanine dipeptide (i.e., comparison to the baseline dataset with and without noise added) but also to simulation data obtained from a CGnet trained according to the protocol in Ref. 31 for the same system. PWe thus also construct a CGnet according to the parameters selected in Ref. 31 (i.e., using fixed geometric features as described in Sec. III A; see also Sec. B of the supplementary material). Then, we create a simulation dataset using the protocol described in Sec. IV A and Sec. C of the supplementary material. Finally, we employ a similar protocol to Sec. IV A by perturbing the raw Cartesian coordinates of the all-atom chignolin simulation dataset with noise distributed as $\mathcal{N}(0, \sigma^2)$ for $\sigma \in \{0, 0.03, 0.06, \ldots, 0.90\}$ Å.

For each type of system (i.e., baseline, CGSchNet, CGnet, and baseline with noise perturbation), we build Markov state models



FIG. 8. The Kullback–Leibler divergence (left) and mean squared error (right) are calculated between the baseline chignolin dataset and the datasets obtained from perturbations to the baseline simulation at noise scale values of $\sigma \in \{0, 0.03, 0.06, \ldots, 0.90\}$, where the former is the reference distribution and the latter is the trial distribution. This procedure is performed 50 times with different random seeds; both plots show the superposition of those 50 lines. The colored horizontal dashed lines show the values of the metric when comparing the CGSchNet (purple) and the CGnet (blue) simulations to the baseline. The black vertical dashed line indicates the noise scale σ that returns the closest value for that metric. The KL divergence is computed for normalized bin counts in reweighted TIC 1 × TIC 2 space, and the MSE is computed for energies, as described in the main text.

(MSMs)⁷⁵ (see Refs. 83 and 84 for recent overviews). First, the data are "featurized" from Cartesian coordinates into the set of 45 distances between pairs of α -carbons. From these distances, time-lagged independent component analysis (TICA)^{85,86} is performed to



FIG. 7. Two- and one-dimensional free energy surfaces for four chignolin datasets. From left to right, datasets are the baseline chignolin simulation (a), the coarse-grained CGSchNet simulation (b), and the coarse-grained CGnet simulation (c) and the dataset generated from the perturbation of the original Cartesian coordinates of the baseline dataset drawn from noise distributed as $\mathcal{N}(0, \sigma^2)$ for $\sigma = 0.3$ Å (d). Each two-dimensional surface is obtained from a 120 × 120 histogram on TIC 1 × TIC 2 space with weights determined from the MSM built for each system (see the main text). The one-dimensional surfaces are similarly obtained from 120-bin histograms on a single TIC. For each reweighted histogram bin, the free energy is obtained by taking the natural log of the counts and multiplying by $-k_BT$; the color scale is the same in all four two-dimensional surfaces and darker color represents lower free energy (i.e., greater stability). The shaded region always represents the baseline dataset and the bold line represents the dataset indicated in the inset title.

yield four slow reaction coordinates for the system. These four reaction coordinates are clustered into 150 discrete, disjoint states using the *k*-means algorithm. An MSM is then estimated from the cluster assignments. The MSM for the baseline simulation dataset is constructed first; then, the other simulation datasets are projected onto the space defined by the former. MSM essentials are presented in Sec. D of the supplementary material from a theoretical standpoint, and the specific protocols used for the MSM analysis in this section are given in Sec. E of the supplementary material.

The stationary distribution of each MSM is then used to reweight the TICA coordinates used for its own construction. Histograms of the first two TICA coordinates are presented in the top row of Fig. 7 for the baseline, CGSchNet, and CGnet simulation datasets as well as the baseline dataset for $\sigma = 0.3$. The first two reweighted TICA coordinates are also individually binned into one-dimensional free energy surfaces, which are depicted in the second and third rows of Fig. 7. We see that the free energy barriers along these reaction coordinates are reasonably approximated by the CGSchNet simulation.

Figure 8 shows the same divergence metrics calculated in Fig. 5 in Sec. IV A. Again, we see that both the KL divergence and the MSE increase monotonically with the magnitude of the noise. In this case, we can assess the equivalent noise value for both the CGSchNet and CGnet simulation datasets. For both divergences measured, we see



FIG. 9. Two-dimensional free energy surfaces (a) and sample folded (b), unfolded (c), and misfolded (d) conformations from the baseline atomistic simulation of chignolin (left column) and the CGSchNet simulation (right column). The free energy surfaces are built from 150-state MSMs that were constructed from the slowest two TICA coordinates in contact distance space. The color scale is the same for both surfaces and darker color represents lower free energy (i.e., greater stability). Each set of ten sampled structures corresponds to the MSM state represented by the star on the free energy surface of the same color (one of the ten structures is opaque for clarity).

that the CGSchNet simulation corresponds to a lesser value of added noise than the CGnet simulation.

We can also obtain free energy surfaces from the MSMs constructed for the systems; the surfaces for the baseline and CGSchNet simulation datasets of chignolin are presented in Fig. 9(a) on the left and right, respectively. We see that the three major basins observed in the atomistic data are captured by CGSchNet. These basins represent folded, unfolded, and misfolded ensembles and are indicated in Fig. 9(a) with blue, green, and yellow stars, respectively. Each star represents one of the 150 MSM states and was manually selected from the MSM states near the relevant basin (see Fig. S8 of the supplementary material for a visualization of all 150 MSM states). To verify that the protein conformations are similar in each of the states, we sample ten structures from each starred state per simulation dataset. The structures are visualized in Figs. 9(b)-9(d), and the similarity of the structures on the left-hand side (baseline simulation) to those on the right-hand side (CGSchNet simulation) from corresponding MSM states is apparent.

The analysis of the CGSchNet and CGnet simulation datasets so far used TICA reaction coordinates that were obtained by projecting the simulation data onto coordinates defined by a TICA model built for the baseline atomistic data (see Sec. E of the supplementary material). This was done in order to compare simulation results using the same reaction coordinates. We can also construct TICA models from the simulation data without projection to determine the scaling factor for the coarse-grained timescale. For this analysis, we build two further (independent) TICA models for the CGSchNet and CGnet datasets at a lag time long enough for the TICA timescales to have leveled off (see Fig. S9 of the supplementary material). A 100-round bootstrapping analysis of the longest TICA timescale from the CGSchNet simulation data yields a time scaling factor of 2.2 with a standard deviation of 0.4. From this time rescaling, we determine that the effective collision rate (i.e., friction) of the coarse-grained simulations is 180 ps^{-1} -260 ps⁻¹. This value is four orders of magnitudes larger than the friction constant in the all-atom model (0.1 ps^{-1}) ,³¹ which we expect because we have coarse-grained out the solvent dynamics. The same analysis for the CGnet simulation data yields a scaling factor of 2.2 \pm 0.3 and a corresponding effective collision rate of 190 ps⁻¹ -250 ps^{-1} .

Although we use MSMs and TICA models to obtain thermodynamics and effective friction constants, we do not attempt a kinetic analysis in the present work because the scope of force matching is limited to thermodynamic consistency.^{36,41} The matching of dynamics in addition to thermodynamics is an open challenge that has been the subject of recent work.⁸⁷ Given coarsegrained dynamics, analytical methods have been derived that enable their rescaling to the dynamics of the system's all-atom counterpart.⁸⁸

V. DISCUSSION

Coarse graining holds the promise of simulating larger systems at longer timescales than are currently possible at the atomistic level. However, mathematical frameworks must be developed in order to ensure that the results obtained from a coarsegrained model are faithful to those that would be obtained from an atomistic simulation or experimental measurement. Force matching^{35,36} is one such framework that, when certain restrictions are applied, guarantees thermodynamic consistency with atomistic data in the variational limit.⁴¹ Such a variational framework enables the formulation of the force matching problem as a supervised machine learning task, which is presented in Ref. 31 as CGnet.

A key limitation of the original CGnet is that it is not transferable across different systems: a new network must be trained for each individual molecular system under study because the molecular features from which it learns the force field must be chosen by hand. Here, we replace manually determined features with a learnable representation. This representation is enabled by the use of continuous filter convolutions on a graph neutral network (i.e., SchNet^{24,33}). SchNet is an inherently transferable architecture originally designed to match energies and forces to quantum calculations for small organic molecules. By leveraging SchNet in the *coarse graining* context—i.e., to learn the molecular features input into a CGnet, we render the hybrid CGnet architecture (i.e., CGSchNet) transferable across molecular systems of different sizes and sequences.

Our aim in the present contribution is threefold: to summarize the variational framework enabling a supervised learning approach to force matching, to provide an accompanying software package implementing the methods discussed herein (see the Appendix), and to demonstrate that CGSchNet produces results on individual systems that are superior to those obtained from bespoke features. The advances presented in this work prepare us to address the ultimate challenge of machine learning a coarse-grained force field that is transferable across molecular systems.

In our computational experiments performed on capped alanine and the miniprotein chignolin, we find that CGSchNet's performance exceeds that of CGnet in three ways. First, the free energy surface obtained from CGSchNet simulations of chignolin is more accurate than the free energy surface presented for the same systems in Ref. 31. Second, CGSchNet is more robust to network hyperparameters than its predecessor. In fact, for the CGSchNet hyperparameters varied during model training (see Sec. B of the supplementary material), the same selections are used for both systems presented in Sec. IV. Third, CGSchNet employs less regularization; particularly, it does not require the extra step of enforcing a Lipschitz constraint⁸⁹ on its network's weight matrices as was found to be necessary for CGnet.³¹

While our current protocol has demonstrated success for a capped monopeptide and a 10-amino acid miniprotein, adapting the CGSchNet pipeline to produce accurate coarse-grained force fields for larger protein systems remains an open challenge. Addressing this challenge may require specific sampling strategies when obtaining training data, the incorporation of new priors that inform tertiary structure formation, or modifications to the CGSchNet architecture itself such as regularization. Successfully modeling the thermodynamics of protein folding or conformational change via a transferable, machine-learned force field would signify a major success for the union of artificial intelligence and the computational molecular sciences.

The method introduced herein enables us to reproduce the thermodynamics of small protein systems using an architecture

that is transferable across system size and sequence. However, CGSchNet is not readily transferable across thermodynamic states. Related work leveraging the same variational principle in a semisupervised learning context allows the learning of coarse-grained representations over multiple thermodynamic states, enabling transferability across different temperatures.^{25,28} This method has been demonstrated for ionic liquids for which nonequilibrium transport properties are of prime interest. Finally, the reproduction of kinetics is an open research problem, and methods for the so-called "spectral matching" problem have recently been introduced.⁸⁷ Ideally, both force and spectral matching could be pursued in conjunction to match both thermodynamics and kinetics simultaneously.

Structural, bottom up coarse graining consists of two aspects: the model resolution and the force field. Here, we assume that the resolution is set and focus on the force field, but the choice of an optimal model resolution is itself a significant challenge that is interconnected to the goal of force field optimization. How to choose a resolution for coarse graining—and the interplay of this choice with transferable force field architectures—remains an open question. Recent work has employed machine learning and data-driven approaches to pursue an optimal resolution using various objectives.^{27,28}

Altogether, the methodology we introduce in the present contribution establishes a transferable architecture for the machine learning of coarse-grained force fields, and we expect our accompanying software to facilitate progress not only in that realm but also toward the outstanding challenges of learning coarse-grained dynamics and optimizing a model's resolution.

SUPPLEMENTARY MATERIAL

See the supplementary material for further specifics on simulation, model training, and MSM construction.

AUTHORS' CONTRIBUTIONS

B.E.H., N.E.C., and D.L. contributed equally to this work.

ACKNOWLEDGMENTS

B.E.H. is immeasurably grateful to Moritz Hoffmann for his wisdom and support. The authors are grateful to Dr. Ankit Patel for discussions on model debugging and machine learning techniques, to Iryna Zaporozhets for discussions concerning model priors, to Dr. Stefan Doerr for software help, to Dr. Jan Hermann for discussions regarding SchNet, and to the reviewers for their constructive comments.

We acknowledge funding from the European Commission (Grant No. ERC CoG 772230 "ScaleCell") to B.E.H., A.K., Y.C., and F.N. from MATH+, the Berlin Mathematics Center, to B.E.H. (Grant No. EF1-2), S.O. (Grant No. AA1-6), and F.N. (Grant Nos. EF1-2 and AA1-6), from the Natural Science Foundation (Grant Nos. CHE-1738990, CHE-1900374, and PHY-1427654) to N.E.C., J.W., and C.C., from the Welch Foundation (Grant No. C-1570) to N.E.C., J.W., and C.C., from the NLM Training Program in Biomedical Informatics and Data Science (Grant No. 5T15LM007093-27) to

N.E.C., from the Einstein Foundation Berlin to C.C., and from the Deutsche Forschungsgemeinschaft (SFB1114 Project Nos. A04 and C03) to F.N. D.L. was supported by a FPI fellowship from the Spanish Ministry of Science and Innovation (Grant Nos. MICINN and PRE2018-085169). G.d.-F. acknowledges support from MINECO (Unidad de Excelencia Mariá de Maeztu AEI; Grant Nos. CEX2018-000782-M and Grant No. BIO2017-82628-P) and FEDER. This project received funding from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 823712 (CompBioMed2 Project). We thank the GPUGRID donors for their computational time.

Simulations were performed on the computer clusters of the Center for Research Computing at Rice University, supported, in part, by the Big-Data Private-Cloud Research Cyberinfrastructure MRI-award (NSF Grant No. CNS-1338099), and on the clusters of the Department of Mathematics and Computer Science at Freie Universität, Berlin.

Part of this research was performed while B.E.H., N.E.C., D.L., J.W., G.d.-F., C.C., and F.N. were visiting the Institute for Pure and Applied Mathematics (IPAM) at the University of California, Los Angeles, for the Long Program "Machine Learning for Physics and the Physics of Learning." IPAM is supported by the National Science Foundation (Grant No. DMS-1440415).

APPENDIX: SOFTWARE

The cgnet software package is available at https://github.com/ coarse-graining/cgnet under the BSD-3-Clause license. cgnet requires NumPy,⁹⁰ SciPy,⁹¹ and PyTorch,⁶⁰ and optional functionalities further depend on pandas,⁹² MDTraj,⁹³ and Scikit-learn.⁹⁴ The examples are provided in Jupyter notebooks, ⁹⁵ which also require Matplotlib.⁹⁶ The SchNet part of the code is inspired by SchNetPack,⁹⁷ and the Langevin dynamics simulation code is adapted from OpenMM.⁹⁸ In addition to cgnet and the packages already mentioned, visualization was aided by Seaborn⁹⁹ and visual molecular dynamics (VMD).¹⁰⁰ Analysis was facilitated by PyEMMA.^{101,102}

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

¹M. Shirts and V. S. Pande, "Screen savers of the world unite!," Science **290**, 1903–1904 (2000).

²F. Allen, G. Almasi, W. Andreoni, D. Beece, B. J. Berne, A. Bright, J. Brunheroto, C. Cascaval, J. Castanos, P. Coteus, P. Crumley, A. Curioni, M. Denneau, W. Donath, M. Eleftheriou, B. Flitch, B. Fleischer, C. J. Georgiou, R. Germain, M. Giampapa, D. Gresh, M. Gupta, R. Haring, H. Ho, P. Hochschild, S. Hummel, T. Jonas, D. Lieber, G. Martyna, K. Maturu, J. Moreira, D. Newns, M. Newton, R. Philhower, T. Picunko, J. Pitera, M. Pitman, R. Rand, A. Royyuru, V. Salapura, A. Sanomiya, R. Shah, Y. Sham, S. Singh, M. Snir, F. Suits, R. Swetz, W. C. Swope, N. Vishnumurthy, T. J. C. Ward, H. Warren, and R. Zhou, "Blue gene: A vision for protein science using a petaflop supercomputer," IBM Syst. J. 40, 310–327 (2001).

³I. Buch, M. J. Harvey, T. Giorgino, D. P. Anderson, and G. De Fabritiis, "High-throughput all-atom molecular dynamics simulations using distributed computing," J. Chem. Inf. Model. **50**, 397–403 (2010).

⁴D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, R. H. Larson, J. K. Salmon, C. Young, B. Batson, K. J. Bowers, J. C. Chao, M. P. Eastwood, J. Gagliardo, J. P. Grossman, C. R. Ho, D. J. Ierardi, I. Kolossváry, J. L. Klepeis, T. Layman, C. McLeavey, M. A. Moraes, R. Mueller, E. C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S. C. Wang, "Anton, a special-purpose machine for molecular dynamics simulation," Commun. ACM 51, 91-97 (2008).

⁵T. J. Lane, D. Shukla, K. A. Beauchamp, and V. S. Pande, "To milliseconds and beyond: Challenges in the simulation of protein folding," Curr. Opin. Struct. Biol. 23, 58-65 (2013).

⁶N. Plattner, S. Doerr, G. De Fabritiis, and F. Noé, "Complete protein-protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling," Nat. Chem. 9, 1005 (2017).

⁷S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, and A. Kolinski, "Coarse-grained protein models and their applications," Chem. Rev. 116, 7898-7936 (2016).

⁸C. Clementi, H. Nymeyer, and J. N. Onuchic, "Topological and energetic factors: What determines the structural details of the transition state ensemble and "en-route" intermediates for protein folding? An investigation for small globular proteins," J. Mol. Biol. 298, 937-953 (2000).

⁹S. J. Marrink, H. J. Risselada, S. Yefimov, D. P. Tieleman, and A. H. De Vries, "The MARTINI force field: Coarse grained model for biomolecular simulations," . Phys. Chem. B 111, 7812-7824 (2007).

¹⁰L. Monticelli, S. K. Kandasamy, X. Periole, R. G. Larson, D. P. Tieleman, and S.-J. Marrink, "The MARTINI coarse-grained force field: Extension to proteins," J. Chem. Theory Comput. 4, 819-834 (2008).

¹¹A. Kolinski et al., "Protein modeling and structure prediction with a reduced representation," Acta Biochim. Pol. 51, 349-371 (2004).

¹²A. Davtyan, N. P. Schafer, W. Zheng, C. Clementi, P. G. Wolynes, and G. A. Papoian, "AWSEM-MD: Protein structure prediction using coarse-grained physical potentials and bioinformatically based local structure biasing," J. Phys. Chem. B 116, 8494-8503 (2012).

¹³R. Das and D. Baker, "Macromolecular modeling with rosetta," Annu. Rev. Biochem. 77, 363-382 (2008).

¹⁴F. Noé, A. Tkatchenko, K.-R. Müller, and C. Clementi, "Machine learning for molecular simulation," Annu. Rev. Phys. Chem. 71, 361-390 (2020).

¹⁵R. Gómez-Bombarelli and A. Aspuru-Guzik, "Machine learning and big-data in computational chemistry," in Handbook of Materials Modeling: Methods: Theory and Modeling (Springer, Cham, 2020), pp. 1939-1962.

¹⁶T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in 5th International Conference on Learning Representations, Toulon, France, April 24-26, 2017.

¹⁷P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner et al., "Relational inductive biases, deep learning, and graph networks," arXiv:1806.01261 (2018).

¹⁸D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, "Convolutional networks on graphs for learning molecular fingerprints," in Advances in Neural Information Processing Systems 28, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., 2015), pp. 2224-2232.

¹⁹S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley, "Molecular graph" convolutions: Moving beyond fingerprints," J. Comput.-Aided Mol. Des. 30, 595-608 (2016).

²⁰J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in Proceedings of the 34th International Conference on Machine Learning (JMLR.org, 2017), Vol. 70, pp. 1263-1272.

²¹E. N. Feinberg, D. Sur, Z. Wu, B. E. Husic, H. Mai, Y. Li, S. Sun, J. Yang, B. Ramsundar, and V. S. Pande, "Potential net for molecular property prediction," ACS Cent. Sci. 4, 1520-1530 (2018).

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

²³K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, "Quantum-chemical insights from deep tensor neural networks," Nat. Commun. 8, 13890 (2017).

²⁴K. Schütt, P.-J. Kindermans, H. E. Sauceda Felix, S. Chmiela, A. Tkatchenko, and K.-R. Müller, "SchNet: A continuous-filter convolutional neural network for modeling quantum interactions," in Advances in Neural Information Processing Systems 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017), pp. 991–1001.

²⁵J. Ruza, W. Wang, D. Schwalbe-Koda, S. Axelrod, W. H. Harris, and R. Gomez-Bombarelli, "Temperature-transferable coarse-graining of ionic liquids with dual graph convolutional neural networks," arXiv:2007.14144 (2020). ²⁶W. G. Noid, "Perspective: Coarse-grained models for biomolecular systems,"

J. Chem. Phys. 139, 090901 (2013).

²⁷L. Boninsegna, R. Banisch, and C. Clementi, "A data-driven perspective on the hierarchical assembly of molecular structures," J. Chem. Theory Comput. 14, 453-460 (2018).

²⁸W. Wang and R. Gómez-Bombarelli, "Coarse-graining auto-encoders for molecular dynamics," npj Comput. Mater. 5, 125 (2019).

²⁹S. T. John and G. Csányi, "Many-body coarse-grained interactions using Gaussian approximation potentials," J. Phys. Chem. B 121, 10934-10949 (2017).

³⁰L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, "DeePCG: Constructing coarse-grained models via deep neural networks," J. Chem. Phys. 149, 034101 (2018).

³¹J. Wang, S. Olsson, C. Wehmeyer, A. Pérez, N. E. Charron, G. De Fabritiis, F. Noé, and C. Clementi, "Machine learning of coarse-grained molecular dynamics force fields," ACS Cent. Sci. 5, 755-767 (2019).

32 J. Wang, S. Chmiela, K.-R. Müller, F. Noé, and C. Clementi, "Ensemble learning of coarse-grained molecular dynamics force fields with a kernel approach," J. Chem. Phys. 152, 194106 (2020).

³³K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller, "SchNet-A deep learning architecture for molecules and materials," J. Chem. Phys. 148, 241722 (2018).

³⁴S. Honda, T. Akiba, Y. S. Kato, Y. Sawada, M. Sekijima, M. Ishimura, A. Ooishi, H. Watanabe, T. Odahara, and K. Harata, "Crystal structure of a ten-amino acid protein," J. Am. Chem. Soc. 130, 15327-15331 (2008).

³⁵F. Ercolessi and J. B. Adams, "Interatomic potentials from first-principles calculations: The force-matching method," Europhys. Lett. 26, 583 (1994).

³⁶S. Izvekov and G. A. Voth, "A multiscale coarse-graining method for biomolecular systems," J. Phys. Chem. B 109, 2469-2473 (2005).

37S. Izvekov, M. Parrinello, C. J. Burnham, and G. A. Voth, "Effective force fields for condensed phase systems from ab initio molecular dynamics simulation: A new method for force-matching," J. Chem. Phys. 120, 10896-10913 (2004).

³⁸M. G. Guenza, M. Dinpajooh, J. McCarty, and I. Y. Lyubimov, "Accuracy, transferability, and efficiency of coarse-grained models of molecular liquids," J. Phys. Chem. B 122, 10257-10278 (2018).

³⁹S. Izvekov and G. A. Voth, "Multiscale coarse graining of liquid-state systems," Chem. Phys. 123, 134105 (2005).

⁴⁰G. Ciccotti, T. Lelièvre, and E. Vanden-Eijnden, "Projection of diffusions on submanifolds: Application to mean force computation," Commun. Pure Appl. Math. 61, 371-408 (2008).

⁴¹W. G. Noid, J.-W. Chu, G. S. Ayton, V. Krishna, S. Izvekov, G. A. Voth, A. Das, and H. C. Andersen, "The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models," J. Chem. Phys. 128, 244114 (2008).

⁴²M. S. Shell, "The relative entropy is fundamental to multiscale and inverse thermodynamic problems," J. Chem. Phys. 129, 144108 (2008).

⁴³J. F. Rudzinski and W. G. Noid, "Coarse-graining entropy, forces, and structures," J. Chem. Phys. 135, 214101 (2011).

⁴⁴In practice, such a search is limited by the basis of trial force fields as well as the finite simulation data used.¹⁰

⁴⁵J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," Phys. Rev. Lett. 98, 146401 (2007).

⁴⁶A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, "Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons," Phys. Rev. Lett. 104, 136403 (2010).

⁴⁷ M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. Von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," Phys. Rev. Lett. **108**, 058301 (2012).

⁴⁸A. P. Bartók, M. J. Gillan, F. R. Manby, and G. Csányi, "Machine-learning approach for one-and two-body corrections to density functional theory: Applications to molecular and condensed water," Phys. Rev. B 88, 054104 (2013).

⁴⁹J. S. Smith, O. Isayev, and A. E. Roitberg, "ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost," Chem. Sci. 8, 3192–3203 (2017).

⁵⁰S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, "Machine learning of accurate energy-conserving molecular force fields," Sci. Adv. 3, e1603015 (2017).

⁵¹ A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. Kermode, G. Csányi, and M. Ceriotti, "Machine learning unifies the modeling of materials and molecules," Sci. Adv. 3, e1701816 (2017).

⁵²J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg, "Less is more: Sampling chemical space with active learning," J. Chem. Phys. **148**, 241733 (2018).

⁵³A. Grisafi, D. M. Wilkins, G. Csányi, and M. Ceriotti, "Symmetry-adapted machine learning for tensorial properties of atomistic systems," Phys. Rev. Lett. 120, 036002 (2018).

⁵⁴G. Imbalzano, A. Anelli, D. Giofré, S. Klees, J. Behler, and M. Ceriotti, "Automatic selection of atomic fingerprints and reference configurations for machinelearning potentials," J. Chem. Phys. **148**, 241730 (2018).

⁵⁵T. T. Nguyen, E. Székely, G. Imbalzano, J. Behler, G. Csányi, M. Ceriotti, A. W. Götz, and F. Paesani, "Comparison of permutationally invariant polynomials, neural networks, and Gaussian approximation potentials in representing water interactions through many-body expansions," J. Chem. Phys. **148**, 241725 (2018).

⁵⁶L. Zhang, J. Han, H. Wang, W. Saidi, R. Car, and W. E, "End-to-end symmetry preserving inter-atomic potential energy model for finite and extended systems," in *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018), pp. 4436–4446.

⁵⁷L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, "Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics," Phys. Rev. Lett. **120**, 143001 (2018).

⁵⁸T. Bereau, R. A. DiStasio, Jr., A. Tkatchenko, and O. A. Von Lilienfeld, "Noncovalent interactions across organic and biological subsets of chemical space: Physics-based potentials parametrized from machine learning," J. Chem. Phys. **148**, 241706 (2018).

⁵⁹H. Wang and W. Yang, "Toward building protein force fields by residue-based systematic molecular fragmentation and neural network," J. Chem. Theory Phys. 15, 1409–1417 (2018).

⁶⁰ A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), pp. 8024–8035.

⁶¹K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2016), pp. 770–778.

⁶²T. Schneider and E. Stoll, "Molecular-dynamics study of a three-dimensional one-component model for distortive phase transitions," Phys. Rev. B 17, 1302 (1978).

⁶³J. Fass, D. A. Sivak, G. E. Crooks, K. A. Beauchamp, B. Leimkuhler, and J. D. Chodera, "Quantifying configuration-sampling error in Langevin simulations of complex molecular systems," Entropy **20**, 318 (2018).

⁶⁴B. Leimkuhler and C. Matthews, "Rational construction of stochastic numerical methods for molecular sampling," Appl. Math. Res. Express **2013**, 34–56.

⁶⁵As in Ref. 32, we require the β -carbon in order to break the symmetry of the system (i.e., to enforce chirality). In their demonstration of CGnet, Wang *et al.*³¹

used only the five backbone heavy atoms as beads because chirality is enforced through dihedral features, which we do not use here.

⁶⁶S. Kullback and R. A. Leibler, "On information and sufficiency," Ann. Math. Stat. **22**, 79–86 (1951).

⁶⁷We could alternatively compute the mean squared error between discrete distributions of counts without omitting any bins.

⁶⁸Similar results were obtained for the two-dimensional Wasserstein distance.

⁶⁹K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, "How fast-folding proteins fold," Science **334**, 517–520 (2011).

⁷⁰K. A. Beauchamp, R. McGibbon, Y.-S. Lin, and V. S. Pande, "Simple fewstate models reveal hidden complexity in protein folding," Proc. Natl. Acad. Sci. U. S. A. **109**, 17807–17813 (2012).

⁷¹ B. E. Husic, R. T. McGibbon, M. M. Sultan, and V. S. Pande, "Optimized parameter selection reveals trends in Markov state models for protein folding," J. Chem. Phys. **145**, 194103 (2016).

⁷²K. A. McKiernan, B. E. Husic, and V. S. Pande, "Modeling the mechanism of CLN025 beta-hairpin formation," J. Chem. Phys. 147, 104107 (2017).

⁷³ M. M. Sultan and V. S. Pande, "Automated design of collective variables using supervised machine learning," J. Chem. Phys. **149**, 094106 (2018).

⁷⁴ M. K. Scherer, B. E. Husic, M. Hoffmann, F. Paul, H. Wu, and F. Noé, "Variational selection of features for molecular kinetics," J. Chem. Phys. 150, 194108 (2019).

⁷⁵R. Zwanzig, "From classical dynamics to continuous time random walks," J. Stat. Phys. **30**, 255–262 (1983).

⁷⁶C. Schütte, A. Fischer, W. Huisinga, and P. Deuflhard, "A direct approach to conformational dynamics based on hybrid Monte Carlo," J. Comput. Phys. 151, 146–168 (1999).

⁷⁷W. C. Swope, J. W. Pitera, and F. Suits, "Describing protein folding kinetics by molecular dynamics simulations. 1. Theory," J. Phys. Chem. B 108, 6571–6581 (2004).

⁷⁸N. Singhal, C. D. Snow, and V. S. Pande, "Using path sampling to build better markovian state models: Predicting the folding rate and mechanism of a tryptophan zipper beta hairpin," J. Chem. Phys. **121**, 415–425 (2004).

⁷⁹J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill, and W. C. Swope, "Automatic discovery of metastable states for the construction of Markov models of macromolecular conformational dynamics," J. Chem. Phys. **126**, 155101 (2007).

⁸⁰ F. Noé, I. Horenko, C. Schütte, and J. C. Smith, "Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states," J. Chem. Phys. **126**, 155102 (2007).

⁸¹N.-V. Buchete and G. Hummer, "Coarse master equations for peptide folding dynamics," J. Phys. Chem. B **112**, 6057–6069 (2008).

⁸²J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, "Markov models of molecular kinetics: Generation and validation," J. Chem. Phys. **134**, 174105 (2011).

⁸³B. E. Husic and V. S. Pande, "Markov state models: From an art to a science,"
 J. Am. Chem. Soc. 140, 2386–2396 (2018).

⁸⁴F. Noé, "Machine learning for molecular dynamics on long timescales," in *Machine Learning Meets Quantum Physics* (Springer, Cambridge, 2020), pp. 331–372.

⁸⁵G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, "Identification of slow molecular order parameters for Markov model construction," J. Chem. Phys. **139**, 015102 (2013).

⁸⁶C. R. Schwantes and V. S. Pande, "Improvements in Markov state model construction reveal many non-native interactions in the folding of NTL9," J. Chem. Theory Comput. 9, 2000–2009 (2013).

⁸⁷ F. Nüske, L. Boninsegna, and C. Clementi, "Coarse-graining molecular systems by spectral matching," J. Chem. Phys. **151**, 044116 (2019).

⁸⁸I. Lyubimov and M. Guenza, "First-principle approach to rescale the dynamics of simulated coarse-grained macromolecular liquids," Phys. Rev. E 84, 031801 (2011).

⁸⁹H. Gouk, E. Frank, B. Pfahringer, and M. Cree, "Regularisation of neural networks by enforcing Lipschitz continuity," arXiv:1804.04368 (2018). ⁹⁰C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," Nature 585, 357–362 (2020).

⁹¹ E. Jones, T. Oliphant, P. Peterson *et al.*, SciPy: Open source scientific tools for Python, 2001.

⁹²W. McKinney et al., "Data structures for statistical computing in python," in Proceedings of the 9th Python in Science Conference (Austin, TX, 2010), Vol. 445, pp. 51–56.

pp. 51–56. ⁹³ R. T. McGibbon, K. A. Beauchamp, M. P. Harrigan, C. Klein, J. M. Swails, C. X. Hernández, C. R. Schwantes, L.-P. Wang, T. J. Lane, and V. S. Pande, "MDTraj: A modern open library for the analysis of molecular dynamics trajectories," Biophys. J. **109**, 1528–1532 (2015).

⁹⁴F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res. **12**, 2825–2830 (2011).

⁹⁵T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, "Jupyter notebooks—A publishing format for reproducible computational workflows," in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, edited by F. Loizides and B. Schmidt (IOS Press, 2016), pp. 87–90.

pp. 87-90. ⁹⁶J. D. Hunter, "Matplotlib: A 2D graphics environment," Comput. Sci. Eng. 9, 90-95 (2007). ⁹⁷K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko, and K.-R. Müller, "SchNetPack: A deep learning toolbox for atomistic systems," J. Chem. Theory Comput. 15, 448–455 (2018).

⁹⁸ P. Eastman, J. Swails, J. D. Chodera, R. T. McGibbon, Y. Zhao, K. A. Beauchamp, L.-P. Wang, A. C. Simmonett, M. P. Harrigan, C. D. Stern *et al.*, "OpenMM 7: Rapid development of high performance algorithms for molecular dynamics," PLoS Comput. Biol. **13**, e1005659 (2017).

⁹⁹M. Waskom, O. Botvinnik, D. O'Kane, P. Hobson, S. Lukauskas, D. C. Gemperline, T. Augspurger, Y. Halchenko, J. B. Cole, J. Warmenhoven, J. de Ruiter, C. Pye, S. Hoyer, J. Vanderplas, S. Villalba, G. Kunter, E. Quintero, P. Bachant, M. Martin, K. Meyer, A. Miles, Y. Ram, T. Yarkoni, M. L. Williams, C. Evans, C. Fitzgerald, Brian, C. Fonnesbeck, A. Lee, and A. Qalieh, Mwaskom/Seaborn: V0.8.1 (September, 2017), 2017.

¹⁰⁰W. Humphrey, A. Dalke, and K. Schulten, "VMD—Visual molecular dynamics," J. Mol. Graph. 14, 33–38 (1996).

¹⁰¹M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Pérez-Hernández, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz, and F. Noé, "PyEMMA 2: A software package for estimation, validation, and analysis of Markov models," J. Chem. Theory Comput. **11**, 5525–5542 (2015).

¹⁰²C. Wehmeyer, M. K. Scherer, T. Hempel, B. E. Husic, S. Olsson, and F. Noé, "Introduction to Markov state modeling with the PyEMMA software—v1. 0," LiveCoMS 1, 5965 (2018).

¹⁰³W. G. Noid, P. Liu, Y. Wang, J.-W. Chu, G. S. Ayton, S. Izvekov, H. C. Andersen, and G. A. Voth, "The multi-scale coarse-graining method. II. Numerical implementation for coarse-grained molecular models," J. Chem. Phys. **128**, 244115 (2008).