

Thermodynamic Transferability in Coarse-Grained Force Fields Using Graph Neural Networks

Published as part of *Journal of Chemical Theory and Computation* special issue "Machine Learning and Statistical Mechanics: Shared Synergies for Next Generation of Chemical Theory and Computation".

Emily Shinkle,* Aleksandra Pachaliev, Riti Bahl, Sakib Matin, Brendan Gifford, Galen T. Craven, and Nicholas Lubbers*



Cite This: *J. Chem. Theory Comput.* 2024, 20, 10524–10539



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: Coarse-graining is a molecular modeling technique in which an atomistic system is represented in a simplified fashion that retains the most significant system features that contribute to a target output while removing the degrees of freedom that are less relevant. This reduction in model complexity allows coarse-grained molecular simulations to reach increased spatial and temporal scales compared with corresponding all-atom models. A core challenge in coarse-graining is to construct a force field that represents the interactions in the new representation in a way that preserves the atomistic-level properties. Many approaches to building coarse-grained force fields have limited transferability between different thermodynamic conditions as a result of averaging over internal fluctuations at a specific thermodynamic state point. Here, we use a graph-convolutional neural network architecture, the Hierarchically Interacting Particle Neural Network with Tensor Sensitivity (HIP-NN-TS), to develop a highly automated training pipeline for coarse-grained force fields, which allows for studying the transferability of coarse-grained models based on the force-matching approach. We show that this approach yields not only highly accurate force fields but also that these force fields are more transferable through a variety of thermodynamic conditions. These results illustrate the potential of machine learning techniques, such as graph neural networks, to improve the construction of transferable coarse-grained force fields.



INTRODUCTION

Molecular simulations elucidate the microscopic physical processes that give rise to a physical system's function and behavior. One of the principal components that determines the accuracy of a molecular simulation is the force field, a mathematical model that calculates the forces acting on the particles in the system as a function of their positions, i.e., the interatomic forces. Force fields are typically constructed and calibrated by a combination of top-down parametrization techniques (so that a simulation reproduces known properties of the target system such as structural, thermodynamic, and dynamical properties measured in experiments) and bottom-up techniques (by fitting to forces generated using first-principles calculations^{1–6}). Molecular models with well-parametrized force fields enable the determination of key physical and chemical properties needed by researchers in a variety of domains such as chemistry, materials science, and biophysics.^{3,7–9} These properties are generated by extracting observable quantities from the results of a simulation using a variety of sampling techniques.^{10,11}

Even with a force field in hand, performing molecular simulations can incur a significant computational cost, primarily because large numbers of atoms are often required in order to understand a system's collective behaviors and statistical properties, and time scales associated with processes of interest can be far longer than the time scales on which atoms themselves evolve. Many techniques have been developed to bridge this gap between atomistic and experimental scales, for example, Accelerated Molecular Dynamics,¹² enhanced sampling methods such as metadynamics,¹³ and correlation function approaches to thermodynamic information such as using the Bogoliubov–Born–Green–Kirkwood–Yvon (BBGKY) hierarchy.¹⁴ To ameliorate the computational cost associated with simulating a system at

Received: June 17, 2024

Revised: November 12, 2024

Accepted: November 15, 2024

Published: November 23, 2024






Figure 1. Illustration of the workflow used to create and analyze the ML CG models.

an all-atom (AA) level, the system can be redefined by using a coarse-grained (CG) representation in which some of the atomistic degrees of freedom are removed to reduce the overall model complexity. Specifically, coarse-graining is a molecular modeling technique in which collections of atoms with highly interrelated behavior are reduced to single particles.^{1,15–26} By neglecting and removing the internal dynamics of the atoms within each group, CG models increase computational efficiency and allow simulations to be performed on increased spatial and temporal scales. Coarse-graining has been shown to be capable of reproducing the structural and thermodynamic properties of a broad class of systems including molecular liquids,^{1,27–30} polymers,^{31–34} and proteins.^{9,15,16,18}

A limitation of traditional coarse-graining techniques and the corresponding force fields used to perform the CG simulations is transferability; that is, CG models are optimized at a specific thermodynamic state point but perform poorly outside of those conditions.³⁵ The transferability problem arises because the effects of the removed degrees of freedom are themselves a function of thermodynamic conditions.³⁶ Researchers are attempting to tackle the transferability problem in coarse-graining using a variety of techniques, including inverse Monte Carlo methods,³⁷ extended ensemble approaches,^{38,39} and integral equation methods,⁴⁰ among others. A related approach is to use machine learning (ML) methods that target thermodynamic consistency^{41–43} to build data-driven CG models that are transferable across thermodynamic conditions. These CG models aim to retain statistical consistency across scales and satisfy known relations between thermodynamic variables and their derivative quantities. Transferability in CG models has been explored, for example, in the context of matching AA configurations across phases, interfaces, and thermodynamic conditions by accounting for both energetic and entropic effects in the CG space.^{44–47} Developing transferable CG models using ML is currently an open and active research area. As applications of ML in coarse-graining expand, it is important to assess if CG models developed using data-driven approaches can improve current methods that seek accurate force field transferability through, for example, the inclusion of correction terms or fitting to properties other than a single structural observable. Using ML to mitigate the need to develop CG force fields that rely on quantities averaged across thermodynamic state points is a potential advantage. The flexibility of ML methods can also potentially improve force field generalizability so that a CG model developed on

one class of system/molecule can be applied to a different class.

The development of AA and CG force fields is a complex task that involves selection of functional forms, creation and curation of data, optimization of parameters, evaluation of preliminary models, and tuning of hyperparameters (such as cost functions weights) associated with fitting.⁴⁸ Recently, approaches using ML methods have been applied to build many-body models of atomistic^{3,7,8,49,50} and coarse-grained^{34,51–53} forces with increased flexibility in comparison to traditional methods. These approaches allow for tight matching with reference data but introduce even more hyperparameters to the model and fitting procedure, further complexifying the automation of force field development. Machine learning techniques such as neural networks^{54–57} and active learning⁵⁸ have been applied to increase the accuracy and reduce the complexity of developing CG force fields. Machine learning applications in coarse-graining include implicit solvation models⁵⁹ and the prediction of solvation energies.⁶⁰ While including many-body interactions and effects of those interactions in CG force fields may improve the accuracy of a CG model, it is more difficult to optimize a complex many-body functional form with simple fitting approaches, e.g., using serial updates to single parameters at a time. Machine learning, however, can provide a computationally tractable approach to quickly fit many-body CG forces when sufficient data are available. The inclusion of many-body forces often results in a better reproduction of the AA configuration distributions.^{60–63}

One common goal in coarse-graining is to maintain thermodynamic consistency between length scales while at the same time realizing the transferability of the CG model across thermodynamic states. This multiobjective optimization is often difficult to realize in practice, however, principally because the multiscale nature of the CG procedure gives rise to nontrivial operations in the construction of force fields and associated mappings between length scales that maintain thermodynamic consistency. Machine learning models are often able to accurately describe a system's behavior over a wide range of thermodynamic states due to the flexible form of the interpolating functions used in ML, for example, neural networks. Advances in ML applied to coarse-graining have been made in parallel with advances in AA force field development using ML-based methods.^{3,7,8} However, the application of ML to develop transferable CG force fields is currently a complex task with many open questions.

In this paper, we present an ML CG workflow to construct force fields based on the force-matching (or multiscale coarse-graining) approach⁶⁴ using the Hierarchically Interacting Particle Neural Network with Tensor Sensitivity (HIP-NN-TS)^{65,66} architecture, which has previously only been applied to AA systems. We show that this workflow is robust in that it is able to consistently build a large number of accurate CG models for a variety of chemical physics systems across many thermodynamic state points in a single phase on the phase diagram. We then use these models to study transferability of the many-body ML CG approach. We compare these results to two-body effective potentials using the recently developed OpenMSCG⁶⁷ software. We find that the ML CG approach is more consistently accurate, can take advantage of additional training data, and produces more transferable models across varying temperature. We furthermore study several molecules across temperature and density variations, comparing single-state-point models with models produced by training to all available data and find that both types are surprisingly transferable even as systems undergo large changes in their structural ordering.

METHODS

In this section, we describe training data generation, the ML model architecture, the training procedure, and evaluation methods used in this work. An overview of the workflow is shown in Figure 1.

Coarse-Graining. We use a typical coarse-graining framework in which collections of atoms in an AA system are mapped to CG *beads*, and internal degrees of freedom in each collection of atoms are removed from the CG model. With this aim, a mapping from the AA configuration space to the CG resolution is selected. We use a bottom-up approach to construct the corresponding CG force field, which is defined to preserve the configuration probability density from the AA space in the CG space.⁶⁸ The relation between the configuration probabilities in each representation is

$$P_{\text{CG}}(\mathbf{R}) = \int P_{\text{AA}}(\mathbf{r}) \delta[\mathcal{M}(\mathbf{r}) - \mathbf{R}] d\mathbf{r} \quad (1)$$

where $\mathbf{r} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\}$ represents the positions $\mathbf{r}_i$ for atoms $1 \leq i \leq n$ in the AA space, $\mathbf{R} = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$ represents the positions $\mathbf{R}_I$ for CG beads $1 \leq I \leq N$, $\mathcal{M}$ represents the mapping function from $\mathbf{r}$ to $\mathbf{R}$, and P_{CG} and P_{AA} represent the respective CG and AA configuration probability densities. For the canonical ensemble generated by an AA energy E_{AA} at inverse temperature β , this equation can be cast in terms of a free energy function E_{CG} as

$$e^{-\beta E_{\text{CG}}(\mathbf{R})} \propto \int e^{-\beta E_{\text{AA}}(\mathbf{r})} \delta[\mathcal{M}(\mathbf{r}) - \mathbf{R}] d\mathbf{r} \quad (2)$$

This free energy function is a *potential of mean force* (PMF), in that the CG forces derived from it should match the force averages across all possible atomistic configurations corresponding to each CG configuration, based on the selected mapping $\mathcal{M}$.

Herein, each CG bead represents the atoms that compose a single molecule and we invoke the common center-of-mass mapping to construct the CG representation. For a discussion of other popular mapping choices, see ref 14. We define the mass M_I of a bead I as a sum of the masses m_i for each atom i in the corresponding molecule as

$$M_I = \sum_{i=1}^n m_i \Delta_{i,I} \quad (3)$$

where Δ is an indicator for which atoms correspond to which bead,

$$\Delta_{i,I} = \begin{cases} 1 & \text{atom } i \text{ is in molecule } I \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The CG bead coordinates $\mathbf{R}_I$ are calculated from the AA configurations using the center of mass position,

$$\mathbf{R}_I = \mathcal{M}(\mathbf{r})_I = \frac{\sum_{i=1}^n m_i \mathbf{r}_i \Delta_{i,I}}{M_I} \quad (5)$$

Also required is a mapping $\mathcal{B}$ from AA forces $\mathbf{f} = \{\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_n\}$ to CG forces $\mathbf{F} = \{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_N\}$. Ciccotti, Kapral, and Vanden-Eijnden⁶⁹ derived a set of criteria for $\mathcal{B}$ that guarantees that the corresponding free energy function E_{CG} is consistent with eq 1. Summarily, while $\mathcal{B}$ may be nonlinear in $\mathbf{r}$, it must be linear in $\mathbf{f}$, and $\mathcal{B}$ must serve as an inverse of the AA gradient of the coordinate mapping, $\nabla_{\mathbf{r}} \mathcal{M}$, when contracted over the AA indices. A simple choice consistent with this criteria is to define $\mathcal{B}$ so that $\mathbf{F}_I$ is the sum of forces $\mathbf{f}_i$ for atoms i in bead I , i.e.,

$$\mathbf{F}_I = \mathcal{B}(\mathbf{f})_I = \sum_{i=1}^n \mathbf{f}_i \Delta_{i,I} \quad (6)$$

Using these CG mappings, for each frame of data, $\mathbf{F}$ provides an unbiased estimator of the negative derivative of the CG free energy E_{CG} with respect to the coordinates $\mathbf{R}$.

Model Architecture. To build the CG free energy function E_{CG} , we applied the Hierarchically Interacting Particle Neural Network with Tensor Sensitivity (HIP-NN-TS).⁶⁶ HIP-NN-TS is a graph-convolutional neural network (GCNN); the convolutions, implemented in an *interaction layer*, make the model invariant under rotations, translations, and permutations of the atoms in the simulation. The tensor-sensitivity component builds upon the original model⁶⁵ by introducing many-body features into the individual neurons of the interaction layer.

We adapted the open-source hippynn⁷⁰ codebase that implements HIP-NN-TS to take bead positions as inputs rather than atom positions. The models then predict the free energy of the system based on the bead configurations. The force on each bead is calculated as the negative gradient of the predicted free energy with respect to the bead positions using automatic differentiation.⁷¹ Using the language of ref 66, the ML CG models developed here use tensor order $l = 2$ and contain $n_{\text{int}} = 1$ interaction layer, $n_{\text{atom}} = 3$ atomic environment layers, $n_{\nu} = 20$ sensitivity functions, and $n_{\text{feature}} = 32$ atomic features per layer. These hyperparameters were determined through trial and error during a preliminary phase of the study based on values, which had shown success in previous HIP-NN networks used for atomistic simulations.^{65,66} A posthoc analysis, described further in the Supporting Information section *Neural network hyperparameter investigation*, indicates a good deal of flexibility in selecting these hyperparameters without significantly diminishing the quality of the results.

In addition to the neural network component of the free energy, a short-range pairwise-repulsive potential was added to the models. The data used to train the models comes from equilibrium simulations in which there is an effective lower bound on intermolecular distance, r . Without the addition of

the repulsive potential, the models can generate unphysical, untrained predictions for the forces between pairs of beads whose distance is less than r , due to the lack of data. Without the repulsive term, the GCNN potential would have a significant uncertainty in the small- r region. This is because those configurations are not seen in the training data and, without the repulsive term, small interparticle distances would constitute significant extrapolation for the potential. By including the repulsive term, we add a physics-based term (in contrast to the ML-based component) that alleviates the sampling problem. This is a common technique to employ in such a situation, although the exact forms of the short-range potentials vary.^{54,73} The repulsive potential improves the stability of the model because extrapolation of the potential to small interbead distances is no longer performed. The repulsive potential ensures that any two beads separated by less than r are repelled. The repulsive potential is of the form

$$E_{\text{rep}}(r) = E_0 e^{-ar} \quad (7)$$

where r is the distance between the pair of beads, and E_0 , $a > 0$ are parameters set based on the specific system. Importantly, because the pairwise potential pertains to the lack of data, it is necessary to set E_0 and a before training the neural network. The procedure used for identifying E_0 and a is detailed in the [Supporting Information](#) section *Repulsive potential parametrization details*. We emphasize that the repulsive pair potential significantly improves the stability of learned ML CG models during simulations, resulting in a highly automated workflow.

Example scripts for training such a model and for using the resulting model to run coarse-grained MD are included in the open-source hippynn⁷⁰ repository.

Training. For the methanol comparison study, we follow the data generation scheme described in an OpenMSCG tutorial⁷⁴ for building CG models of methanol. The molecular dynamics suite GROMACS⁷⁵ is used to generate 100,000 time steps of 1 fs each of a box with 1728 methanol molecules under periodic boundary conditions. Every 100th step is saved, resulting in 1000 AA frames. The topology and initial coordinates are downloaded from the GROMACS webpage and the OPLS-AA⁷⁶ force field is used. The simulation is run in the canonical ensemble using a Nosé–Hoover thermostat.

For the cross-molecular study, we use the LAMMPS⁷⁷ software package with the all-atom (i.e., nonconstrained) GROMOS-54A⁷⁸ force field to simulate three molecular fluids: methanol, benzene, and methane across a variety of temperatures and densities above the critical temperature for each molecule. The molecular topologies were obtained using the Automated Topology Builder (ATB) and Repository.^{79,80} PACKMOL⁸¹ was used to generate initial coordinates, and Moltemplate⁸² was used to generate LAMMPS input files. We performed MD simulations in the canonical ensemble using a Nosé–Hoover thermostat. Each simulation contained 1024 molecules, and the density was controlled using cubic boxes of various size with periodic boundary conditions. Following an equilibration procedure, each MD simulation consisted of 50,000 time steps of 1 fs. Every 50th frame was recorded, totalling 1000 frames.

For both studies, after generating the AA data, the mappings M of eq 5 and B of eq 6 were applied to each frame to create the training data for the ML CG models. For the methanol comparison study and the cross-molecular study, we trained *single-state* models using data from a single-state point (i.e., for each temperature and density combination). These single-state

models are then applied using MD at the same state point to which they were trained to establish a baseline. They are also applied with MD at other temperatures and densities to test their transferability. For the cross-molecular study, we also trained a *multistate* model for each molecule using combined data from every state point and tested it using MD at each of those state points. Figure 2 shows an illustration of these three different test schemes.




Figure 2. Illustration of the three types of tests performed on the ML CG method.

The multistate data sets are the same size as the single-state data sets, built by evenly subselecting the frames from each state point. In each case, 1000 frames of training data are split randomly into 800 training frames, 100 validation frames, and 100 testing frames. The loss function used to train the ML CG models was the sum of the root-mean-square error (RMSE) and mean absolute error (MAE) for the model predictions versus the training data values of the forces for each configuration in the data set. Both the MAE and RMSE were included in the loss function because the combination was found to be successful in prior works.^{65,66} A posthoc investigation, described in the [Supporting Information](#) section *Neural network hyperparameter investigation*, suggests that using RMSE but not MAE (as prescribed by the MS-CG formalism and corresponding PMF) to train a model has minimal impact on the radial distribution function (RDF) and angular distribution function (ADF) error presented later.

Performance Evaluation. We measure the accuracy of CG models by how well they reproduce statistics from the AA models in accordance with eq 1. In particular, the RDFs measured during MD simulations using the CG models were compared to the RDFs generated from the AA MD simulations. The RDF, denoted by $g(r)$, describes the statistics for finding a particle at a distance r from a randomly chosen particle. It is normalized so that an RDF is asymptotically one for large radii by defining the RDF as the ratio between the local density variations and the bulk density.⁸³ The RDF is a




Figure 3. Subfigures (a, b) show a comparison of methanol RDFs generated using (1) a reference AA simulation, (2) the MS-CG technique, and (3) the single-state ML CG models. These RDFs were generated at (a) 200 K, density 0.77 g/cm³ and (b) 400 K, density 0.77 g/cm³. Subfigure (c) summarizes the corresponding results across 11 temperatures. Subfigure (d) shows the transferability of a single-state model for each method, each trained with 300 K data.

key structural metric used in characterizing the degree of local ordering in a system. It can be determined using theoretical, computational, or experimental approaches and can be used to derive much thermodynamic information about a system.^{83–87} In this work, AA RDFs are computed using molecular centers of mass, and CG RDFs, correspondingly, are computed using bead positions. As a ratio between densities, the RDF has no units; we denote this as (–) in the figures.

In order to construct the CG RDFs, we performed CG MD simulations in the canonical ensemble. The initial positions for these simulations were taken from a random frame of the model training data at the appropriate temperature and density. After equilibrating the CG system, 50,000 time steps of 1 fs each were run, and each 50th frame was recorded. The RDFs were computed from these 1000 frames of data. The AA RDFs were taken from the AA training data, which was also performed using canonical ensemble MD as detailed in the **Training** subsection.

To quantify the difference between the AA and CG RDFs, we use the *total absolute error* (TAE). The TAE is the total area between two curves, i.e.,

$$\text{TAE}(g_1, g_2) = \int_0^\infty |g_2 - g_1| dr \quad (8)$$

For the calculations in this paper, we used a finite sum approximation of TAE. Specifically, we used

$$\text{TAE}(g_1, g_2) \approx \sum_{i=0}^n |g_2(r_i) - g_1(r_i)| \Delta r \quad (9)$$

where $\{r_i\}$ represents a sequence of distance values, evenly spaced between $r_0 = 0$ and an appropriately chosen $r_n = r_{\max} > 0$, with width Δr between each pair of successive values. We chose to compute the TAE between RDFs rather than the MAE between RDFs because the RDF TAE is not strongly sensitive to the choice of upper cutoff $r_{\max}$ of the RDF radius, whereas the RDF MAE is.

Additionally, we compare the ADFs from the ML CG MD runs to those from the AA MD simulations. The ADF of a trajectory describes the distribution of angles between triplets of points in each frame of a trajectory. This higher-order structural function provides further insight into the arrangements of particles in the trajectory beyond what is contained in the RDF. More precisely, the ADF is the probability distribution of angles in the range $[0^\circ, 180^\circ]$ where each angle is formed by a central particle to two neighbors, where both arms of the angle from the central particle have length less than a set cutoff distance $r_{\max}$. In general, higher values of $r_{\max}$ will produce less structured ADFs, as the behavior of particles at greater distances from one another is less correlated. As with the RDF, the ADF is a ratio and hence has no units, which we denote as (–) in the figures. Unlike the RDF, for the ADF, either TAE or MAE can be appropriate to quantify the difference between two histograms, as the range of values for which the distribution is defined is fixed at $[0^\circ, 180^\circ]$. We use the ADF MAE here, again approximated by a finite sum, which is dimensionless.

A usual method for evaluating ML models is to compare predicted model outputs (forces, in our case) with the corresponding true values from test data withheld during




Figure 4. Comparison of methanol ADFs generated using (1) a reference AA simulation, (2) the MS-CG technique, and (3) the single-state ML CG models. Subfigures (a), (c), and (e) show ADFs for 200 K, density 0.77 g/cm³ with ADF cutoff values r_{max} of 4, 5, and 6 Å respectively. Subfigures (b), (d), and (f) show ADFs for 400 K, density 0.77 g/cm³ with the same respective cutoff values.

training. However, for the CG problem here, the training data $\mathbf{F}$ is a statistical distribution of forces, whereas the model predicts the mean of this distribution, $\langle \mathbf{F} \rangle$; as a result, there is intrinsic noise in the loss function.⁸⁸ For this reason, conventional metrics are not easily usable for assessing the quality of the ML CG model. Nevertheless, several conventional metrics for the performance of the models are reported in the [Supporting Information](#) section, *Training statistics*.

RESULTS

In this section, the performance of the ML CG models is explored. Both *single-state* models, those trained using data at a specific temperature and density, and *multistate* models, those trained using data from a range of temperatures and densities, are discussed.

Methanol Comparison Study. To establish a baseline for the ML CG approach, we compared it to the recently released OpenMSCG software for coarse-graining methanol.⁶⁷ OpenMSCG is based upon methods proposed by Izvekov and Voth^{1,89} and later further developed and generalized by others including Noid et al.^{64,68,90} The OpenMSCG software provides a set of *force-matching* routines that implement a

bottom-up coarse-graining method that calculates the effective CG interactions by minimizing the difference between CG forces and reference AA forces; this is quite similar to our workflow and uses the same coordinate and force mappings. The main difference is that OpenMSCG uses a pairwise force/energy function between beads instead of a GCNN.

The resulting RDFs for this comparison are given in [Figure 3](#). Both methods perform well, with the ML CG method exhibiting significantly lower error at lower temperatures, as shown in [Figure 3a](#) for 200 K. These results illustrate that, in general, we expect the neural network-based methodology to result in CG force fields that perform as well as or better than those constructed using force-matching when applied to data at a specific state point. [Figure 3b](#) shows the results for methanol at 400 K. In this case, the OpenMSCG and ML CG methods give very similar and excellent overall accuracy. As such, in this case, there is only a limited potential for improvement over the OpenMSCG model. We performed this study for a range of temperatures from 200 to 400 K, with the overall RDF TAEs presented in [Figure 3c](#). At lower temperatures, the ML CG models perform substantially better than the OpenMSCG models. As the temperature is increased,




Figure 5. Summary of ADF comparisons for the MS-CG technique and the ML CG technique against a reference AA simulation. Subfigures (a), (c), and (e) show results for the single-state baseline models using cutoff values $r_{\max}$ of (a) 4, (c) 5, and (e) 6 Å. Subfigures (b), (d), and (f) show results for the single-state transferability test using the 300 K model for each method with ADF cutoff values $r_{\max}$ of (b) 4, (d) 5, and (f) 6 Å.

the difference between the two methods is less pronounced. In the temperature range 340 K–380 K, the OpenMSCG method is slightly more accurate, although the difference between the methods is relatively small. The models shown in Figure 3c were trained by using 1000 frames of training data. In the Supporting Information, Figure S2, we show similar results constructed using 10 and 100 frames of training data. With far fewer data, the advantage of ML CG over OpenMSCG is less pronounced; we find the performance of OpenMSCG saturates more quickly as the data set size is increased.

When comparing the ADFs generated with the OpenMSCG method and the ML CG method against the AA reference data, as shown in Figure 4 and Figure 5 parts (a), (c), and (e), we see again that the ML CG models perform better or similarly to the OpenMSCG models in every case. The difference is most stark for the ADFs generated with a cutoff value $r_{\max} = 4$ Å, where the ADFs are most structured. The MAE between the OpenMSCG ADFs with the reference AA ADFs is around twice that for the ML CG models at every temperature tested. When the ADF cutoff value $r_{\max}$ is increased to 5 or 6 Å, this difference largely disappears, and both methods reproduce the ADFs very closely.

To understand the transferability of the models, we applied the single-state model learned at 300 K for each method to a range of temperatures and recorded the TAEs of the resulting RDFs in Figure 3d. In both cases, the model produces lower RDF errors in the vicinity of the training state point, and the TAE rises smoothly at higher and lower temperatures. However, the ML CG model shows a significantly lower RDF TAE compared to OpenMSCG when applied at temperatures further from the training state point. At temperatures lower than 300 K, the TAE for the ML CG models is approximately a factor 2 lower in comparison to OpenMSCG. The corresponding ADF errors are shown in Figure 5 parts (b), (d), and (f). For cutoff value $r_{\max} = 4$ Å (subfigure (b)), the MAEs for the ML CG ADF are approximately half the MAE for the OpenMSCG model at each temperature. For cutoff value $r_{\max} = 5$ Å (subfigure (d)), this difference is lessened but the ML CG model still slightly outperforms the OpenMSCG model in every case. Finally, with cutoff value $r_{\max} = 6$ Å (subfigure (f)), the MAE error values for the OpenMSCG ADFs are nearly identical to the ML CG error values for temperatures between 300 and 400 K, while




Figure 6. Baseline comparison of reference RDFs generated from AA MD against RDFs generated via an ML CG single-state model for (a) methanol, (b) benzene, and (c) methane.

the OpenMSCG errors are closer to double that of the ML CG models at temperatures from 200 to 280 K.

The primary difference between the ML CG method and the OpenMSCG method is that the ML CG models incorporate many-body effects, while the OpenMSCG methods are pairwise. These results illustrate how accounting for many-body interactions may be relevant to producing more transferable CG force fields.

Multimolecule Study. To further study thermodynamic transferability in ML CG models, we applied the ML CG method to several molecules (methanol, benzene, and methane) across a range of temperatures and densities, as described in the [Training](#) subsection.

Figure 6 shows examples of the RDFs generated by using single-state models for methanol, benzene, and methane. Each subfigure shows a reference AA RDF computed using molecular centers-of-mass for atomistic MD data generated at the same temperature and density. There is excellent agreement between the AA and CG RDFs in each case, illustrating that the ML CG models accurately capture the shape and magnitude of the RDF peaks for each of the examined molecular fluids.

Figure 7 shows the results of three workflows for all three molecules studied across the nine state points used for each molecule. First, the single-state models are used to provide a baseline of a usual CG approach that trains and tests at the same thermodynamic state point (denoted *single-state (base.)* in the figure). Second, for each molecule, the single-state model trained at the center state point is applied to each of the test state points (denoted *single-state (trans.)*), providing a picture of the transferability of the ML CG model when extrapolating through thermodynamic state space. Finally, the multistate model is applied to each state point (denoted *multistate (trans.)*) to test whether training to all state points

provides improved transferability characteristics (for details, see the [Training](#) section). Each molecular fluid was studied at three densities (shown in three panels) and three temperatures (shown on the vertical axis of each panel) in [Figure 7](#). The error bars show the standard deviation calculated across five trials, where, in each trial, several factors were randomly varied: the train/valid/test split for the data used to create the ML CG model, the initial weights of the HIP-NN-TS network, the initial frame for the ML CG dynamics, and the random number seed used for the thermostat. Of the 405 trajectory comparisons used to create this figure, some of those with the highest TAEs are shown in Supporting Information [Figure S5](#).

The findings for the single-state baseline test were consistent for all three molecules. For methanol ([Figure 7a](#)), we observe that the single-state models generate similar, very low (<0.25 Å) TAE values over all the densities and temperatures studied. For benzene ([Figure 7b](#)), the average TAE is good (<0.45), although somewhat higher than methanol, but again, single-state models produce fairly consistent TAE values across all state points. For methane ([Figure 7c](#)), the single-state models perform best at the lowest temperature and highest density studied. This is an interesting observation because at this state point, we expect the RDF to have the most structure and stronger correlations at larger r values in comparison to, for example, higher temperatures and lower densities. For the high-density state points, the ML CG model for methane performs worse as the temperature is increased. Overall, the results in [Figures 6](#) and [7](#) illustrate that the ML CG methodology developed in this paper can be applied to construct CG free energy functions that generate RDFs in strong agreement with AA results.

For the transferability test of the single-state models, the results show a surprising level of transferability for ML CG models; in general, the TAEs for this test are almost always




Figure 7. Total absolute error (TAE) between RDFs generated by AA MD and those generated by using various ML CG models for (a) methanol, (b) benzene, and (c) methane. The horizontal placement of the markers has been offset slightly for visual clarity. The error bars show the standard deviation calculated by constructing five ML CG models randomly sampled as described in the text.

within a factor of 2 of the baseline single-state performance, with the exception of high-temperature, high-density methanol. Even more surprising, for several state points, the single-state transferability performance is actually superior to the baseline models. Another observation is that significant variance between runs can be observed for high-density benzene. Under more careful examination, these models produced slightly understructured RDFs, and the variance between models is explained by the degree of under-structuring. A similar but less pronounced effect is visible for high-density methanol, where some variance appears for the single-state transferability model. In this case, the RDFs were not definitively under- or overstructured. Despite these fluctuations, the RDF quality is still reasonable.

Like the single-state transferability tests, the multistate models exhibit remarkable transferability, in some instances even performing better than the single-state models. The higher variance observed in the single-state transferability tests for benzene was significantly reduced. More detailed investigation of the underlying RDFs demonstrates that the multistate ML CG models reproduce the AA statistics even through relatively large changes in the structural ordering of the fluid. Figure 8 shows the most extreme state points (high-density, low-temperature, and low-density, high-temperature)

with respect to structural ordering. All of the molecules undergo significant changes in structural ordering across the space of thermodynamic states. We also tested whether the multistate model can be productively applied to state points not present in the training data. We examined this for methanol by generating AA ground truth at two intermediate densities and two intermediate temperatures, yielding four new state points. The results of the multistate model at these previously unseen points are shown in Supporting Information Figure S4 and demonstrate that the performance of the model does not significantly change when it is applied to intermediate state points not seen during training.

Further, when comparing ADFs generated using AA MD against the ADFs generated with the ML CG method, we see strikingly similar results. ADFs from selected multistate transferability tests are shown in Figure 9, with the results pictured in each subfigure coming from the same trial as the results in the corresponding subfigure in Figure 8. Each subfigure shows the ADFs corresponding to several choices of the ADF cutoff value r_{max} . The only pictured case in which there is visible deviation from the reference ADFs is for methanol at 800 K with density 0.85 g/cm³, where the ADFs for cutoff values r_{max} of 4, 5, and 6 Å exhibit slightly less structure than the corresponding reference ADFs. The ADF




Figure 8. Comparison of reference RDFs generated by AA MD against RDFs generated via the ML CG multistate model for (a) and (b) methanol, (c) and (d) benzene, and (e) and (f) methane. The state points represented are the leftmost and rightmost pictured in Figure 7 for each molecule, which have, respectively, the most and least structured RDFs.

results for all molecules, tests, and state points at selected cutoff values are displayed in Figure 10. Several of the ADFs with the highest MAE values among those captured in this figure are provided in Supporting Information Figure S6. Many similarities to the RDF results shown in Figure 7 can be found. For each molecule, the single-state baseline test results were largely consistent for each molecule. For methanol (Figure 10a), the ADF MAE values are around 0.02 for each temperature and density combination. For benzene (Figure 10b), the ADF MAEs are again all near 0.02, with slightly more fluctuation between state points than what was observed for methanol. Finally, for methane (Figure 10c), the ADF MAEs range from around 0.01 to 0.02, with improved performance as the density increases. These results provide further evidence that the single-state ML CG models can closely replicate the structure of the AA trajectories used to train them.

Figure 10 also shows the ADF MAEs for the transferability test of the single-state models. Again, the results further support the claim of the transferability of these single-state ML CG models. In many instances, the ADF error for the transferability test model is nearly equivalent to the ADF error for the single-state baseline model. The largest increases in

ADF error when performing the transferability test were for benzene at 1.06 g/cm³, the highest density tested, followed by methanol at 800 K and 1.27 g/cm³ and at 600 K and 0.85 g/cm³. In all cases, the ADF MAE is no more than three times the baseline result.

Finally, Figure 10 includes the results for the multistate transferability test ADFs. We observe very similar results for the multistate and single-state transferability tests in almost every instance, with the highest error for the multistate transferability test ADFs being for benzene at 1.06 g/cm³, followed by several temperature and density combinations for methanol. Again, in each instance, the ADF MAE is no more than three times the baseline result. And just as observed with the RDF errors, there is fluctuation between whether the single-state or multistate transferability test results are better. This finding reinforces the claim of transferability of both the single-state and multistate ML CG models across temperature and density.

The ML CG workflow produces accurate models for a range of molecules and state points without tuning. The main parameter that might be adjusted is the interaction radius of




Figure 9. Comparison of reference ADFs generated by AA MD against ADFs generated via the ML CG multistate model at several different cutoff values $r_{\max}$ for (a) and (b) methanol, (c) and (d) benzene, and (e) and (f) methane. The cutoff values used are 4, 5, 6, and 7 Å for methanol; 6, 7, 8, and 9 Å for benzene; and 5, 6, 7, and 8 Å for methane. The smaller cutoff values in each case correspond to the more structured ADFs.

the HIP-NN-TS architecture, but in this work, a uniform interaction distance of 12 Å was used for all molecules.

Computational Cost. Computation times for ML CG training and MD are affordable enough to enable high-throughput analysis. In the multimolecule study, for benzene and methane, training the ML CG models and running MD with the trained model each consistently took less than 30 min with a single NVIDIA A100 GPU. For methanol, each took consistently under one h with the same architecture. As a result, in the course of this work, we were able to train more than 150 models and to run more than 400 MD simulations. Scientifically, the most important aspect of automating these simulations was the addition of the repulsive potential determined as a pretraining step.

For the results in the [Methanol Comparison Study](#) subsection above, the ML CG method takes about 45 min to run 100,000 steps on 1728 beads, whereas the pairwise OpenMSCG method takes a few minutes to run 100,000 time steps. Given the extremely optimized implementation of pairwise tabular potentials in LAMMPS, a factor of about 20 between the OpenMSCG example and ML CG potential is better than expected.

We also performed a comparison of our ML CG models to the AA HIP-NN-TS models of ref 66 using a periodic box of 1024 methanol molecules. The AA ML potential could compute the forces in approximately 0.3 s on average, where the CG potential could compute the forces in approximately 0.04 s on average. There is also a factor of 4 between the typical time step for AA MD with an ML potential (0.25 fs) and the time step used here (1 fs), leading to an overall 30 times faster simulation capacity with the ML CG method.

All coarse-grained MD simulations in this manuscript were performed using a time step of 1 fs. A time step of between 0.5 and 2 fs is common for coarse-grained neural network models. For example, ref 91 uses 2 fs and ref 92 uses 2 and 0.5 fs depending on the model. We tested a time step of 5 fs for one MD run of each type (single-state baseline, single-state transferability test, and multistate transferability test) for each of the nine methanol temperature/density combinations used in the [Multimolecule Study](#) subsection. We tested energy conservation using this 5 fs time step, which was satisfactory in 17 of the 27 models. The models that failed to achieve energy conservation were mostly for high-density or high-temperature configurations. Improving the dynamics of the ML CG models




Figure 10. Mean absolute error (MAE) between ADFs generated by AA MD and using various ML CG models for (a) methanol (with a cutoff value of $r_{\max} = 4$ Å), (b) benzene (with a cutoff value of $r_{\max} = 6$ Å), and (c) methane (with a cutoff value of $r_{\max} = 5$ Å). The horizontal placement of the markers has been offset slightly for visual clarity. The error bars show the standard deviation calculated by constructing five ML CG models randomly sampled as described in the text.

could increase their ability to conserve energy with larger time steps at high densities and temperatures, allowing for an even greater speedup when compared with AA CG methods. Overall this indicates that our overall speedup factor over all-atom HIP-NN-TS can thus be estimated as in the range 30x to 150x.

Of course, there are many other factors which will affect this number, variations in the network architecture hyperparameters, the overall number of particles considered (1024 particles in the CG simulation does not typically saturate the GPU) and, perhaps most importantly, variations in the number of atoms which are coarse-grained over. In regards to this last factor, methanol has a 6-to-1 ratio between atomistic and CG representations, but benzene has a 12-to-1 particle ratio between atomistic and CG representations and furthermore has fewer neighbors in the same interaction radius, leading to reduced cost for message-passing operations in the neural network.

CONCLUSIONS

In this work, we built an ML CG workflow based on the bottom-up approach of Multi-Scale Coarse-Graining, for the first time using the Hierarchically Interacting Particle Neural

Network (HIP-NN)^{65,66} with the aim of studying the transferability of coarse-grained models through varying thermodynamic conditions. In order to produce a robust workflow suitable for studying many molecules and state points, it was important to include a pairwise-repulsive potential with parameters set before training. Otherwise, the workflow is very similar to training an AA potential using (fluctuating) forces but not energies. As a result, we were able to build more than 150 models and run over 400 MD simulations with those models, exploring the accuracy of models trained and used at the same state point, models trained and used at different state points, and models trained and used at multiple state points.

Our results show that the ML CG models here produce more consistent structural accuracy (as quantified by the RDF and ADF) than OpenMSCG pairwise CG models built using the same data using liquid methanol at a variety of temperatures. Furthermore, training a model at 300 K and deploying it across the range of 200 to 400 K showed that the ML CG model is also significantly more transferable than the pairwise model. This is intriguing because although the neural network is highly expressive, it is not obvious how it distills many-body contributions to the CG free energy that make the

potential more transferable. This is even more surprising when considering the inherent noise in force-matching, which drastically reduces the model's ability to match its training data precisely. This stands in contrast to the construction of AA potentials, where forces and energies can be matched during training with extreme precision.

We furthermore applied the ML CG workflow to study supercritical methanol, benzene, and methane across a range of densities and temperatures to study transferability from a more broad perspective. These conditions span significant variations in structural ordering in the fluids. We found that while overall model accuracy is superior using single-state models, models transferred across thermodynamic states do not produce extreme levels of error and in some instances appear to paradoxically produce somewhat lower error than single-state models. We also tested building force-matched models trained across multiple state points. Although such a workflow violates the thermodynamic assumptions of coarse-graining (as the CG free energy is a function of temperature and density), there was no difficulty in producing these multistate models, and they performed very well. However, given the relatively surprising level of transferability of single-state ML CG models, the comparative advantage of training to multiple state points was not strong. Multistate training did improve on the variance of the RDF error for higher density methanol and benzene.

Finally, we estimated the speedup factor in the ML CG models here over all-atom HIP-NN-TS potentials to assess the computational leverage of coarse-graining. This analysis leads to an estimate of 30x-150x speedup, with the large variance being due to the time step used in the CG model, which could, in many cases, have been pushed up to 5 fs. This shows that ML CG models do, in fact, accomplish key cost-reduction associated with coarse-graining. Simulation throughput can be drastically improved while many characteristics of the underlying fine-scale (AA) system.

Although the ML CG models are able to reproduce the structure of the AA trajectories with high accuracy, they do not replicate the dynamics correctly. This is a known and widely observed problem with coarse-grained models, which occurs due to the smoothing of the potential energy surface and the removal of friction effects. Techniques to address this issue are currently in development.^{93–105}

Another area for future work is to explore transferability for cases of larger differences in structure, such as through phase transitions, e.g., between crystal and liquid. It might be that in this case, a multistate training procedure shows stronger advantages. However, it is more difficult to automate training data generation across thermodynamic phase changes, and this challenge would need to be addressed. Additionally, there is a wide variety of possible targets for coarse-graining, such as proteins, macromolecules, alloys, and liquid mixtures, which might be explored. Another future work possibility is to explore the potential for nonequilibrium coarse-grained simulations. Given the surprising transferability of these potentials, it may be possible to accurately model near-equilibrium conditions where temperatures and/or pressures evolve, either as a function of time or over space given by some boundary conditions. An exciting possibility is to incorporate additional thermodynamics into the free energy function, essentially applying the concept of thermodynamically consistent learning⁴³ to coarse-graining; such a concept has recently been introduced and explored for the coarse-graining of hexane.¹⁰⁶ With a wide range of recent improvements and ideas, machine

learning-based coarse-graining is poised to enable accurate simulations on large length and time scales across a wider range of thermodynamic conditions.

■ ASSOCIATED CONTENT

Data Availability Statement

The data that support the findings of this study are available from the corresponding authors upon reasonable request and institutional approval.

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.4c00788>.

Further details on neural network hyperparameter selection, repulsive potential parameter selection, model training statistics, a discussion on training set size, results of an interpolation test, and additional RDFs and ADFs (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Emily Shinkle – *Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*; orcid.org/0000-0003-2255-0540; Email: eshinkle@lanl.gov

Nicholas Lubbers – *Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*; orcid.org/0000-0002-9001-9973; Email: nlubbers@lanl.gov

Authors

Aleksandra Pachalieva – *Earth and Environmental Sciences Division and Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*; orcid.org/0000-0003-1246-0410

Riti Bahl – *Department of Mathematics, Emory University, Atlanta, Georgia 30322, United States; Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*

Sakib Matin – *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*

Brendan Gifford – *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*; orcid.org/0000-0002-4116-711X

Galen T. Craven – *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*

Complete contact information is available at:

<https://pubs.acs.org/doi/10.1021/acs.jctc.4c00788>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

We acknowledge support from the Los Alamos National Laboratory (LANL) Directed Research and Development funds (LDRD). This research was performed in part at the Center for Nonlinear Studies (CNLS) at LANL. This research used resources provided by the Darwin testbed at LANL, which is funded by the Computational Systems and Software Environments subprogram of LANL's Advanced Simulation and Computing program. LANL is operated by the Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy (contract no.

89233218NCA000001). We would like to thank David Rosenberger and Kipton Barros for helpful discussions and the referees for their thoughtful reviews.

REFERENCES

- (1) Izvekov, S.; Voth, G. A. Multiscale coarse graining of liquid-state systems. *J. Chem. Phys.* **2005**, *123*, No. 134105.
- (2) Lu, C.; Wu, C.; Ghoreishi, D.; Chen, W.; Wang, L.; Damm, W.; Ross, G. A.; Dahlgren, M. K.; Russell, E.; Von Bargen, C. D.; Abel, R.; Friesner, R. A.; Harder, E. D. OPLS4: Improving Force Field Accuracy on Challenging Regimes of Chemical Space. *J. Chem. Theory Comput.* **2021**, *17*, 4291–4300.
- (3) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chem. Rev.* **2021**, *121*, 10142–10186.
- (4) Polêto, M. D.; Lemkul, J. A. Integration of experimental data and use of automated fitting methods in developing protein force fields. *Commun. Chem.* **2022**, *5*, No. 38.
- (5) Zongo, K.; Béland, L.; Ouellet-Plamondon, C. First-principles database for fitting a machine-learning silicon interatomic force field. *MRS Adv.* **2022**, *7*, 39–47.
- (6) Matin, S.; Allen, A. E. A.; Smith, J.; Lubbers, N.; Jadrich, R. B.; Messerly, R.; Nebgen, B.; Li, Y. W.; Tretiak, S.; Barros, K. Machine Learning Potentials with the Iterative Boltzmann Inversion: Training to Experiment. *J. Chem. Theory Comput.* **2024**, *20*, 1274–1281.
- (7) Botu, V.; Batra, R.; Chapman, J.; Ramprasad, R. Machine Learning Force Fields: Construction, Validation, and Outlook. *J. Phys. Chem. C* **2017**, *121*, 511–522.
- (8) Wu, S.; Yang, X.; Zhao, X.; Li, Z.; Lu, M.; Xie, X.; Yan, J. Applications and Advances in Machine Learning Force Fields. *J. Chem. Inf. Model.* **2023**, *63*, 6972–6985.
- (9) Rosenberger, D.; Smith, J. S.; Garcia, A. E. Modeling of Peptides with Classical and Novel Machine Learning Force Fields: A Comparison. *J. Phys. Chem. B* **2021**, *125*, 3598–3612.
- (10) Frenkel, D.; Smit, B. *Understanding Molecular Simulation: From Algorithms to Application*; Academic Press: NY, 1996.
- (11) Fedik, N.; Zubatyuk, R.; Kulichenko, M.; Lubbers, N.; Smith, J. S.; Nebgen, B.; Messerly, R.; Li, Y. W.; Boldyrev, A. I.; Barros, K.; et al. Extending machine learning beyond interatomic potentials for predicting molecular properties. *Nat. Rev. Chem.* **2022**, *6*, 653–672.
- (12) Perez, D.; Uberuaga, B. P.; Shim, Y.; Amar, J. G.; Voter, A. F. Accelerated molecular dynamics methods: introduction and recent developments. *Annu. Rev. Comput. Chem.* **2009**, *5*, 79–98.
- (13) Hénin, J.; Lelièvre, T.; Shirts, M. R.; Valsson, O.; Delemotte, L. Enhanced Sampling Methods for Molecular Dynamics Simulations [Article v1.0]. *Living J. Comput. Mol. Sci.* **2022**, *4*, 1583.
- (14) Rudzinski, J. F.; Noid, W. G. Investigation of Coarse-Grained Mappings via an Iterative Generalized Yvon–Born–Green Method. *J. Phys. Chem. B* **2014**, *118*, 8295–8312.
- (15) Shih, A. Y.; Arkhipov, A.; Freddolino, P. L.; Schulten, K. Coarse Grained Protein-Lipid Model with Application to Lipoprotein Particles. *J. Phys. Chem. B* **2006**, *110*, 3674–3684.
- (16) Saunders, M. G.; Voth, G. A. Coarse-Graining Methods for Computational Biology. *Annu. Rev. Biophys.* **2013**, *42*, 73–93.
- (17) Lu, L.; Dama, J. F.; Voth, G. A. Fitting coarse-grained distribution functions through an iterative force-matching method. *J. Chem. Phys.* **2013**, *139*, No. 121906.
- (18) Noid, W. G. Perspective: Coarse-grained models for biomolecular systems. *J. Chem. Phys.* **2013**, *139*, No. 090901.
- (19) Jin, J.; Pak, A. J.; Durumeric, A. E.; Loose, T. D.; Voth, G. A. Bottom-up coarse-graining: Principles and perspectives. *J. Chem. Theory Comput.* **2022**, *18*, 5759–5791.
- (20) Dama, J. F.; Sinitskiy, A. V.; McCullagh, M.; Weare, J.; Roux, B.; Dinner, A. R.; Voth, G. A. The Theory of Ultra-Coarse-Graining. 1. General Principles. *J. Chem. Theory Comput.* **2013**, *9*, 2466–2480.
- (21) Davtyan, A.; Dama, J. F.; Sinitskiy, A. V.; Voth, G. A. The Theory of Ultra-Coarse-Graining. 2. Numerical Implementation. *J. Chem. Theory Comput.* **2014**, *10*, 5265–5275.
- (22) Brennan, J. K.; Lisal, M.; Moore, J. D.; Izvekov, S.; Schweigert, I. V.; Larentzos, J. P. Coarse-Grain Model Simulations of Non-equilibrium Dynamics in Heterogeneous Materials. *J. Phys. Chem. Lett.* **2014**, *5*, 2144–2149.
- (23) Schilling, T. Coarse-grained modelling out of equilibrium. *Phys. Rep.* **2022**, *972*, 1–45.
- (24) Marrink, S. J.; Risselada, H. J.; Yefimov, S.; Tieleman, D. P.; de Vries, A. H. The MARTINI Force Field: Coarse Grained Model for Biomolecular Simulations. *J. Phys. Chem. B* **2007**, *111*, 7812–7824.
- (25) Noid, W. G. Perspective: Advances, Challenges, and Insight for Predictive Coarse-Grained Models. *J. Phys. Chem. B* **2023**, *127*, 4174–4207.
- (26) Palma Banos, M.; Popov, A. V.; Hernandez, R. Representability and Dynamical Consistency in Coarse-Grained Models. *J. Phys. Chem. B* **2024**, *128*, 1506–1514.
- (27) Craven, G. T.; Popov, A. V.; Hernandez, R. Structure of a tractable stochastic mimic of soft particles. *Soft Matter* **2014**, *10*, 5350–5361.
- (28) Moore, T. C.; Iacovella, C. R.; McCabe, C. Derivation of coarse-grained potentials via multistate iterative Boltzmann inversion. *J. Chem. Phys.* **2014**, *140*, No. 224104.
- (29) Dunn, N. J. H.; Noid, W. G. Bottom-up coarse-grained models that accurately describe the structure, pressure, and compressibility of molecular liquids. *J. Chem. Phys.* **2015**, *143*, No. 243148.
- (30) Moradzadeh, A.; Aluru, N. R. Transfer-Learning-Based Coarse-Graining Method for Simple Fluids: Toward Deep Inverse Liquid-State Theory. *J. Phys. Chem. Lett.* **2019**, *10*, 1242–1250.
- (31) Narros, A.; Moreno, A. J.; Likos, C. N. Influence of topology on effective potentials: coarse-graining ring polymers. *Soft Matter* **2010**, *6*, 2435–2441.
- (32) Craven, G. T.; Popov, A. V.; Hernandez, R. Effective Surface Coverage of Coarse-Grained Soft Matter. *J. Phys. Chem. B* **2014**, *118*, 14092–14102.
- (33) Zhang, P.; Wang, Q. Solvent entropy and coarse-graining of polymer lattice models. *Soft Matter* **2013**, *9*, 11183–11187.
- (34) Ricci, E.; Vergadou, N. Integrating Machine Learning in the Coarse-Grained Molecular Simulation of Polymers. *J. Phys. Chem. B* **2023**, *127*, 2302–2322.
- (35) Chennakesavalu, S.; Toomer, D. J.; Rotskoff, G. M. Ensuring thermodynamic consistency with invertible coarse-graining. *J. Chem. Phys.* **2023**, *158*, No. 124126.
- (36) Dunn, N. J. H.; Foley, T. T.; Noid, W. G. Van der Waals Perspective on Coarse-Graining: Progress toward Solving Representability and Transferability Problems. *Acc. Chem. Res.* **2016**, *49*, 2832–2840.
- (37) Rosenberger, D.; van der Vegt, N. F. A. Addressing the temperature transferability of structure based coarse graining models. *Phys. Chem. Chem. Phys.* **2018**, *20*, 6617–6628.
- (38) Mullinax, J. W.; Noid, W. G. Extended ensemble approach for deriving transferable coarse-grained potentials. *J. Chem. Phys.* **2009**, *131*, No. 104110.
- (39) Shen, K.; Sherck, N.; Nguyen, M.; Yoo, B.; Köhler, S.; Speros, J.; Delaney, K. T.; Fredrickson, G. H.; Shell, M. S. Learning composition-transferable coarse-grained models: Designing external potential ensembles to maximize thermodynamic information. *J. Chem. Phys.* **2020**, *153*, No. 154116.
- (40) Guenza, M. G.; Dinpajooh, M.; McCarty, J.; Lyubimov, I. Y. Accuracy, Transferability, and Efficiency of Coarse-Grained Models of Molecular Liquids. *J. Phys. Chem. B* **2018**, *122*, 10257–10278.
- (41) Craven, G. T.; Lubbers, N.; Barros, K.; Tretiak, S. Machine learning approaches for structural and thermodynamic properties of a Lennard-Jones fluid. *J. Chem. Phys.* **2020**, *153*, No. 104502.
- (42) Craven, G. T.; Lubbers, N.; Barros, K.; Tretiak, S. Ex Machina Determination of Structural Correlation Functions. *J. Phys. Chem. Lett.* **2020**, *11*, 4372–4378.
- (43) Rosenberger, D.; Barros, K.; Germann, T. C.; Lubbers, N. Machine learning of consistent thermodynamic models using automatic differentiation. *Phys. Rev. E* **2022**, *105*, No. 045301.

- (44) Jin, J.; Yu, A.; Voth, G. A. Temperature and Phase Transferable Bottom-up Coarse-Grained Models. *J. Chem. Theory Comput.* **2020**, *16*, 6823–6842.
- (45) Jin, J.; Pak, A. J.; Voth, G. A. Understanding Missing Entropy in Coarse-Grained Systems: Addressing Issues of Representability and Transferability. *J. Phys. Chem. Lett.* **2019**, *10*, 4549–4557.
- (46) Kidder, K. M.; Szukalo, R. J.; Noid, W. Energetic and entropic considerations for coarse-graining. *Eur. Phys. J. B* **2021**, *94*, No. 153.
- (47) Pretti, E.; Shell, M. S. A microcanonical approach to temperature-transferable coarse-grained models using the relative entropy. *J. Chem. Phys.* **2021**, *155*, No. 094102.
- (48) Harrison, J. A.; Schall, J. D.; Maskey, S.; Mikulski, P. T.; Knippenberg, M. T.; Morrow, B. H. Review of force fields and intermolecular potentials used in atomistic computational materials research. *Appl. Phys. Rev.* **2018**, *5*, No. 031104.
- (49) Kulichenko, M.; Smith, J. S.; Nebgen, B.; Li, Y. W.; Fedik, N.; Boldyrev, A. I.; Lubbers, N.; Barros, K.; Tretiak, S. The Rise of Neural Networks for Materials and Chemical Dynamics. *J. Phys. Chem. Lett.* **2021**, *12*, 6227–6243.
- (50) Smith, J. S.; Isayev, O.; Roitberg, A. E. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.* **2017**, *8*, 3192–3203.
- (51) Thaler, S.; Stupp, M.; Zavadlav, J. Deep coarse-grained potentials via relative entropy minimization. *J. Chem. Phys.* **2022**, *157*, No. 244103.
- (52) Wang, J.; Olsson, S.; Wehmeyer, C.; Pérez, A.; Charron, N. E.; De Fabritiis, G.; Noé, F.; Clementi, C. Machine learning of coarse-grained molecular dynamics force fields. *ACS Cent. Sci.* **2019**, *5*, 755–767.
- (53) Fu, X.; Xie, T.; Rebello, N. J.; Olsen, B. D.; Jaakkola, T. Simulate time-integrated coarse-grained molecular dynamics with geometric machine learning. Doctoral dissertation, Massachusetts Institute of Technology, 2022.
- (54) Husic, B. E.; Charron, N. E.; Lemm, D.; Wang, J.; Pérez, A.; Majewski, M.; Krämer, A.; Chen, Y.; Olsson, S.; de Fabritiis, G.; Noé, F.; Clementi, C. Coarse graining molecular dynamics with graph neural networks. *J. Chem. Phys.* **2020**, *153*, No. 194101.
- (55) Loose, T. D.; Sahrman, P. G.; Qu, T. S.; Voth, G. A. Coarse-Graining with Equivariant Neural Networks: A Path Toward Accurate and Data-Efficient Models. *J. Phys. Chem. B* **2023**, *127*, 10564–10572.
- (56) Zhang, L.; Han, J.; Wang, H.; Car, R.; E, W. DeePCG: Constructing coarse-grained models via deep neural networks. *J. Chem. Phys.* **2018**, *149*, No. 034101.
- (57) Ruza, J.; Wang, W.; Schwalbe-Koda, D.; Axelrod, S.; Harris, W. H.; Gómez-Bombarelli, R. Temperature-transferable coarse-graining of ionic liquids with dual graph convolutional neural networks. *J. Chem. Phys.* **2020**, *153*, No. 164501.
- (58) Duschatko, B. R.; Vandermause, J.; Molinari, N.; Kozinsky, B. Uncertainty driven active learning of coarse grained free energy models. *npj Comput. Mater.* **2024**, *10*, No. 9.
- (59) Airas, J.; Ding, X.; Zhang, B. Transferable Implicit Solvation via Contrastive Learning of Graph Neural Networks. *ACS Cent. Sci.* **2023**, *9*, 2286–2297.
- (60) Ge, P.; Zhang, L.; Lei, H. Machine learning assisted coarse-grained molecular dynamics modeling of meso-scale interfacial fluids. *J. Chem. Phys.* **2023**, *158*, No. 064104.
- (61) Larini, L.; Lu, L.; Voth, G. A. The multiscale coarse-graining method. VI. Implementation of three-body coarse-grained potentials. *J. Chem. Phys.* **2010**, *132*, No. 164107.
- (62) Das, A.; Andersen, H. C. The multiscale coarse-graining method. IX. A general method for construction of three body coarse-grained force fields. *J. Chem. Phys.* **2012**, *136*, No. 194114.
- (63) Wang, J.; Charron, N.; Husic, B.; Olsson, S.; Noé, F.; Clementi, C. Multi-body effects in a coarse-grained protein force field. *J. Chem. Phys.* **2021**, *154*, No. 164113.
- (64) Noid, W. G.; Chu, J.-W.; Ayton, G. S.; Krishna, V.; Izvekov, S.; Voth, G. A.; Das, A.; Andersen, H. C. The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models. *J. Chem. Phys.* **2008**, *128*, No. 244114.
- (65) Lubbers, N.; Smith, J. S.; Barros, K. Hierarchical modeling of molecular energies using a deep neural network. *J. Chem. Phys.* **2018**, *148*, No. 241715.
- (66) Chigaev, M.; Smith, J. S.; Anaya, S.; Nebgen, B.; Bettencourt, M.; Barros, K.; Lubbers, N. Lightweight and effective tensor sensitivity for atomistic neural networks. *J. Chem. Phys.* **2023**, *158*, No. 184108.
- (67) Peng, Y.; Pak, A. J.; Durumeric, A. E.; Sahrman, P. G.; Mani, S.; Jin, J.; Loose, T. D.; Beiter, J.; Voth, G. A. OpenMSCG: A Software Tool for Bottom-Up Coarse-Graining. *J. Phys. Chem. B* **2023**, *127*, 8537–8550.
- (68) Noid, W. G.; Chu, J.-W.; Ayton, G. S.; Voth, G. A. Multiscale coarse-graining and structural correlations: Connections to liquid-state theory. *J. Phys. Chem. B* **2007**, *111*, 4116–4127.
- (69) Ciccotti, G.; Kapral, R.; Vanden-Eijnden, E. Blue Moon Sampling, Vectorial Reaction Coordinates, and Unbiased Constrained Dynamics. *ChemPhysChem* **2005**, *6*, 1809–1814.
- (70) Lubbers, N. e. a. hippynn. 2024, <https://github.com/lanl/hippynn>.
- (71) Smith, J. S.; Lubbers, N.; Thompson, A. P.; Barros, K. Simple and efficient algorithms for training machine learning potentials to force data. arXiv preprint arXiv:2006.05475, 2020 <https://arxiv.org/abs/2006.05475>.
- (72) Fellman, A.; Byggmästar, J.; Granberg, F.; Nordlund, K.; Djurabekova, F. Fast and accurate machine-learned interatomic potentials for large-scale simulations of Cu, Al and Ni. 2024 <https://arxiv.org/abs/2408.15779>.
- (73) Byggmästar, J.; Hamedani, A.; Nordlund, K.; Djurabekova, F. Machine-learning interatomic potential for radiation damage and defects in tungsten. *Phys. Rev. B* **2019**, *100*, No. 144105.
- (74) MSCG Tutorials: Lesson 01. Force-Matching Single-Site Methanol. <https://software.rcc.uchicago.edu/mscg/tutorials/lesson-01/README.html>, Accessed: 2024–03–04.
- (75) Lindahl, E.; Hess, B.; Van Der Spoel, D. GROMACS 3.0: a package for molecular simulation and trajectory analysis. *J. Mol. Model.* **2001**, *7*, 306–317.
- (76) Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *J. Am. Chem. Soc.* **1996**, *118*, 11225–11236.
- (77) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; in 't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; Shan, R.; Stevens, M. J.; Tranchida, J.; Trott, C.; Plimpton, S. J. LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Comput. Phys. Commun.* **2022**, *271*, No. 108171.
- (78) Schmid, N.; Eichenberger, A. P.; Choutko, A.; Riniker, S.; Winger, M.; Mark, A. E.; van Gunsteren, W. F. Definition and testing of the GROMOS force-field versions 54A7 and 54B7. *Eur. Biophys. J.* **2011**, *40*, 843–856.
- (79) Malde, A. K.; Zuo, L.; Breeze, M.; Stroet, M.; Poger, D.; Nair, P. C.; Oostenbrink, C.; Mark, A. E. An Automated Force Field Topology Builder (ATB) and Repository: Version 1.0. *J. Chem. Theory Comput.* **2011**, *7*, 4026–4037.
- (80) Stroet, M.; Caron, B.; Visscher, K. M.; Geerke, D. P.; Malde, A. K.; Mark, A. E. Automated Topology Builder Version 3.0: Prediction of Solvation Free Enthalpies in Water and Hexane. *J. Chem. Theory Comput.* **2018**, *14*, 5834–5845.
- (81) Martínez, L.; Andrade, R.; Birgin, E. G.; Martínez, J. M. PACKMOL: A package for building initial configurations for molecular dynamics simulations. *J. Comput. Chem.* **2009**, *30*, 2157–2164.
- (82) Jewett, A. I.; Stelter, D.; Lambert, J.; Saladi, S. M.; Roscioni, O. M.; Ricci, M.; Autin, L.; Maritan, M.; Bashusqeh, S. M.; Keyes, T.; Dame, R. T.; Shea, J.-E.; Jensen, G. J.; Goodsell, D. S. Moltemplate: A Tool for Coarse-Grained Modeling of Complex Biological Matter and Soft Condensed Matter Physics. *J. Mol. Biol.* **2021**, *433*, No. 166841.

- (83) Rosenberger, D.; Lubbers, N.; Germann, T. C. Evaluating diffusion and the thermodynamic factor for binary ionic mixtures. *Phys. Plasmas* **2020**, *27*, No. 102705.
- (84) Trokhymchuk, A.; Nezbeda, I.; Jirsák, J.; Henderson, D. Hard-sphere radial distribution function again. *J. Chem. Phys.* **2005**, *123*, No. 024501.
- (85) Hansen, J. P.; McDonald, I. R. *Theory of Simple Liquids*; Academic Press, 2006.
- (86) Gray, C. G.; Gubbins, K. E. *Theory of Molecular Fluids*; Clarendon Press: Oxford, 1984; Vol. 1.
- (87) Foidl, C. Exact pair distribution function and structure factor for a one-dimensional hard rod mixture. *J. Chem. Phys.* **1986**, *85*, 410–417.
- (88) Krämer, A.; Durumeric, A. E. P.; Charron, N. E.; Chen, Y.; Clementi, C.; Noé, F. Statistically Optimal Force Aggregation for Coarse-Graining Molecular Dynamics. *J. Phys. Chem. Lett.* **2023**, *14*, 3970–3979.
- (89) Izvekov, S.; Voth, G. A. A multiscale coarse-graining method for biomolecular systems. *J. Phys. Chem. B* **2005**, *109*, 2469–2473.
- (90) Noid, W. G.; Liu, P.; Wang, Y.; Chu, J.-W.; Ayton, G. S.; Izvekov, S.; Andersen, H. C.; Voth, G. A. The multiscale coarse-graining method. II. Numerical implementation for coarse-grained molecular models. *J. Chem. Phys.* **2008**, *128*, No. 244115.
- (91) Durumeric, A. E. P.; Chen, Y.; Noé, F.; Clementi, C. Learning data efficient coarse-grained molecular dynamics from forces and noise. 2024 <https://arxiv.org/abs/2407.01286>.
- (92) Loose, T. D.; Sahrman, P. G.; Qu, T. S.; Voth, G. A. Coarse-Graining with Equivariant Neural Networks: A Path Toward Accurate and Data-Efficient Models. *J. Phys. Chem. B* **2023**, *127*, 10564–10572.
- (93) Jin, J.; Schweizer, K. S.; Voth, G. A. Understanding dynamics in coarse-grained models. I. Universal excess entropy scaling relationship. *J. Chem. Phys.* **2023**, *158*, No. 034103.
- (94) Izvekov, S.; Voth, G. A. Modeling real dynamics in the coarse-grained representation of condensed phase systems. *J. Chem. Phys.* **2006**, *125*, No. 151101.
- (95) Hijón, C.; Español, P.; Vanden-Eijnden, E.; Delgado-Buscalioni, R. Mori–Zwanzig formalism as a practical computational tool. *Faraday Discuss.* **2010**, *144*, 301–322.
- (96) Fritz, D.; Koschke, K.; Harmandaris, V. A.; van der Vegt, N. F. A.; Kremer, K. Multiscale modeling of soft matter: scaling of dynamics. *Phys. Chem. Chem. Phys.* **2011**, *13*, 10412–10420.
- (97) Kinjo, T.; Hyodo, S.-A. Equation of motion for coarse-grained simulation based on microscopic description. *Phys. Rev. E* **2007**, *75*, No. 051109.
- (98) Lei, H.; Caswell, B.; Karniadakis, G. E. Direct construction of mesoscopic models from microscopic simulations. *Phys. Rev. E* **2010**, *81*, No. 026704.
- (99) Gao, L.; Fang, W. Semi-bottom-up coarse graining of water based on microscopic simulations. *J. Chem. Phys.* **2011**, *135*, No. 184101.
- (100) Davtyan, A.; Dama, J. F.; Voth, G. A.; Andersen, H. C. Dynamic force matching: A method for constructing dynamical coarse-grained models with realistic time dependence. *J. Chem. Phys.* **2015**, *142*, No. 154104.
- (101) Davtyan, A.; Voth, G. A.; Andersen, H. C. Dynamic force matching: Construction of dynamic coarse-grained models with realistic short time dynamics and accurate long time dynamics. *J. Chem. Phys.* **2016**, *145*, No. 224107.
- (102) Izvekov, S. Microscopic derivation of particle-based coarse-grained dynamics: Exact expression for memory function. *J. Chem. Phys.* **2017**, *146*, No. 124109.
- (103) Jung, G.; Hanke, M.; Schmid, F. Iterative Reconstruction of Memory Kernels. *J. Chem. Theory Comput.* **2017**, *13*, 2481–2488.
- (104) Han, Y.; Jin, J.; Voth, G. A. Constructing many-body dissipative particle dynamics models of fluids from bottom-up coarse-graining. *J. Chem. Phys.* **2021**, *154*, No. 084122.
- (105) Izvekov, S. Mori–Zwanzig theory for dissipative forces in coarse-grained dynamics in the Markov limit. *Phys. Rev. E* **2017**, *95*, No. 013303.
- (106) Duschatko, B. R.; Fu, X.; Owen, C.; Xie, Y.; Musaelian, A.; Jaakkola, T.; Kozinsky, B. Thermodynamically Informed Multimodal Learning of High-Dimensional Free Energy Models in Molecular Coarse Graining. 2024 <https://arxiv.org/abs/2405.19386>.



CAS BIOFINDER DISCOVERY PLATFORM™

**PRECISION DATA
FOR FASTER
DRUG
DISCOVERY**

CAS BioFinder helps you identify
targets, biomarkers, and pathways

Unlock insights

CAS
A division of the
American Chemical Society