

ARTICLE OPEN



Uncertainty driven active learning of coarse grained free energy models

Blake R. Duschatko¹✉, Jonathan Vandermause^{1,2}, Nicola Molinari^{1,3} and Boris Kozinsky^{1,3}✉

Coarse graining techniques play an essential role in accelerating molecular simulations of systems with large length and time scales. Theoretically grounded bottom-up models are appealing due to their thermodynamic consistency with the underlying all-atom models. In this direction, machine learning approaches hold great promise to fitting complex many-body data. However, training models may require collection of large amounts of expensive data. Moreover, quantifying trained model accuracy is challenging, especially in cases of non-trivial free energy configurations, where training data may be sparse. We demonstrate a path towards uncertainty-aware models of coarse grained free energy surfaces. Specifically, we show that principled Bayesian model uncertainty allows for efficient data collection through an on-the-fly active learning framework and opens the possibility of adaptive transfer of models across different chemical systems. Uncertainties also characterize models' accuracy of free energy predictions, even when training is performed only on forces. This work helps pave the way towards efficient autonomous training of reliable and uncertainty aware many-body machine learned coarse grain models.

npj Computational Materials (2024)10:9; <https://doi.org/10.1038/s41524-023-01183-5>

INTRODUCTION

Molecular dynamics (MD) has long been used to efficiently study the statistical and kinetic properties of a wide variety of systems by integrating Newton's equations of motion^{1,2}. There are, however, substantial limitations on the applicability of molecular dynamics to situations where the behavior at long length and time scales is relevant, such as many biological processes. Fast degrees of freedom in the system prohibit the use of larger integration time steps while creating a more rugged energy landscape that typically slows important structural changes^{3,4}. Exploring a rough energy landscape with small time steps becomes unfavorable when coarser resolution is of interest. In addition, computing forces and updating the configuration state of every degree of freedom (DOF) at each time step in an all-atom (AA) system can be a computational burden. For a variety of problems, fast motions such as those of hydrogen vibrations do not play a significant role in long length- and time-scale properties, making it unnecessary to track each DOF.

In the context of biological systems, understanding protein folding pathways and the bulk properties of lipid bilayer membranes often require simulations of physical timescales on the order of microseconds or longer, while also typically modeling tens or hundreds of thousands of atoms^{5,6}. In the most cutting-edge applications of MD, on the other hand, it is not typically feasible to exceed these two limits simultaneously. Similarly, at the all-atom scale protein-protein interactions, or more generally polymer-polymer interactions, may not depend on all degrees of freedom. Accounting for the effects of solvents adds additional computational complexity. In such scenarios as described above, a variety of coarse grained (CG) techniques are often used to probe the system at longer time scales and lower spatial resolution^{3,4,7–26}.

Of central importance to modeling the thermodynamics of systems with CG approaches is identifying the reduced degrees of freedom, or CG beads, and determining the interactions between

them. To this end, two primary approaches are taken. Top-down methods are parametric models tuned to reproduce experimental observations, while bottom-up methods are built upon an underlying all-atom description of the system. Top-down coarse grained force fields are well established with widespread use across different applications^{3,4,12–18,20,21}, while the development of bottom-up coarse grained models has seen substantial activity within the past decade or so^{23,27–40}. Although top-down methods are appealing for many reasons including thermodynamic transferability and speed, they fail in many regards to be more than simply qualitative in nature. On the other hand, despite acquiring explicit dependence on thermodynamic state points, bottom-up approaches can rigorously reproduce the statistical behavior of the system of interest by targeting the many-body coordinate dependent free energy, often called the potential of mean force (PMF), represented by a reduced set of coordinates^{27,31}.

Due to the complexity of the PMF, machine learned (ML) forces fields have recently gained traction over empirically parameterized classical potentials. In particular, the success of all-atom machine learned force fields as surrogate models for ab initio molecular dynamics^{41–55} has resulted in increased interest in applying similar techniques to modeling the PMF^{23,33–40}. These approaches typically make use of regression over long all-atom trajectories via a multiscale coarse graining/force matching technique that reproduces the all-atom PMF in the limit of sufficient sampling of the canonical ensemble²⁷. In addition, machine learning has targeted related problems of choosing an optimal low-resolution representation of atomic systems³⁷ and more broadly the problem of reconstructing atomic details from CG models^{56–58}. So far, most ML approaches to CG models were based on neural networks (NN), which possess a number of benefits. For example, they serve as highly flexible representations of complex functions that can be trained on large training sets. Kernel-based methods have also been proposed^{38–40}. However, all

¹John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. ²Department of Physics, Harvard University, Cambridge, MA 02138, USA. ³Robert Bosch LLC, Watertown, MA 02472, USA. ✉email: bduchatko@g.harvard.edu; bkoz@seas.harvard.edu

ML-based CG models of the PMF to date lack a measure of uncertainty through Bayesian techniques or even through neural network ensembles.

The wide array of problems that are faced in CG modeling motivate a new perspective. For instance, in exploratory applications, such as the discovery of structure of large proteins or rapid screening of soft materials, training CG models with a set of relevant configurations may become infeasible. Some configurations may never be sampled due to their rarity or the time it takes for a system to evolve into these states. For this reason, current approaches that do not use uncertainty metrics rely heavily on substantial a priori knowledge of the target behavior of the all-atom system, posing a limitation on their potential for wide spread applicability.

Moreover, it becomes increasingly difficult to assess the quality of models by estimating the true error on the test set, since this requires long constrained dynamics trajectories to obtain well-converged PMF data. In addition, traditional aggregate metrics such as mean absolute errors of forces may not capture subtle deficiencies in a model. Also, bottom-up CG models are highly dependent on the configuration and chemical make up of the all-atom system used for training, making the transferability of these CG models difficult to anticipate.

Compared to NNs, kernel-based Bayesian regression methods, such as Gaussian process (GP) regression, provide access to predictive uncertainty and have been demonstrated to efficiently select sufficiently representative training sets via active learning in all-atom settings^{47,49,50}. One limitation is the increase of the computational cost of the training and inference tasks with the size of the training set. However, in many cases, the full GP can be mapped onto an exact model for predicting both the mean and uncertainty with a cost that is independent of the training set size. Application of these recent methods have been demonstrated to simulations of complex heterogeneous systems at record speed and size, reaching 500 billion atoms^{51,59,60}. In the context of coarse graining, active learning frameworks have seen less activity, and the existing literature on these methods is limited to non-bottom-up approaches⁶¹.

In this work, we present an active learning regression framework based on principled Bayesian uncertainty inherent to sparse Gaussian processes (SGP's) for autonomous development of coarse grained force field models. First, we demonstrate the ability of Gaussian process regression to learn the coarse grained PMF on-the-fly, thereby reducing the guess work typically needed to construct the training set. In practice, this allows the model to discover unknown configurations that may appear at long timescales by bypassing the more predictable motions of faster degrees of freedom. We emphasize that this differs from current approaches that rely on all-atom simulations to sufficiently explore the full configuration space on their own before removing any degrees of freedom. Second, we show that uncertainty-aware active learning enables the development of more transferable coarse grained models. In particular, we demonstrate how uncertainty, along with the locality, allows models of one molecular system to be transferred to a new system by updating the training database. In addition, we explore the implications of the design choice of how to label the species of CG beads on model accuracy and transferability, which has an important impact on the final speed of the model. Third, we find that uncertainty allows for a rapid and direct assessment of model robustness and limitations where traditional metrics such as force errors are insufficient, thereby accelerating deployment and facilitating automatic refinement.

RESULTS

Validity of uncertainty aware on-the-fly learning of coarse grained models

Bottom-up training of coarse grained models is typically done by regressing the PMF derivatives to time averages of forces, the precise form of which is given by Eqs. (7) and (8). A common approach is to utilize instantaneous forces in a long unconstrained molecular dynamics trajectory to minimize appropriate functionals that reproduce the PMF^{27,33}. This approach requires care in dealing with two implicit timescales in the problem. For one, the simulation times of the fast degrees of freedom must be long enough to ensure that the sampled all-atom configurations are not highly correlated. Also, collecting sufficient training data requires simulating long enough time scales in order to visit a full range of coarse grained configurations. Even so, depending on the shape of the PMF, some regions in CG configuration space separated by barriers may not be visited during the training simulation. Principled quantitative uncertainty provides a rigorous way of identifying configurations that lie within and outside the training set.

Here, we introduce an on-the-fly CG workflow implemented in the FLARE framework^{49,50}, depicted in Fig. 1, that utilizes the predictive Bayesian uncertainty of SGPs. The goal of this FLARE-CG approach is to automate the collection of the PMF labels by deriving a decision threshold of data acquisition from the uncertainty associated with every prediction during a CG MD simulation. This is implemented by directly comparing each local CG environment for which predictions are made with those in the training set using a pre-defined kernel function operating on geometric descriptors of local configuration environments. Specifically, we consider the local environment of CG sites within a defined cutoff radius from the central site. Many-body symmetry preserving descriptors based on the atomic cluster expansion (ACE) are constructed for this purpose, with different hyperparameters allowing for different levels of expressiveness^{50,62}.

The workflow begins from an initial coarse grained structural configuration, for which a constrained dynamics trajectory is performed to acquire force labels for training. Note that the acquired forces are the gradients of, and therefore directly related to, the PMF that we are trying to learn (see Methods A). Subsequently, for each configuration, the predicted models local PMF uncertainties are used to decide whether to evolve the system forward in time. If the uncertainty of the model on the new configuration is above a user-defined threshold, more constrained dynamics data is collected to augment the Gaussian process model database; otherwise, the CG MD step is accepted, and the system evolves forward by a time step. This process is repeated at every step of the CG MD trajectory, forming an autonomous active learning loop. Eventually, after the model no longer makes frequent calls to the all-atom baseline, the trained SGP model's explicit dependence on the training set size can be eliminated by mapping it onto an exactly equivalent and much faster parametric model⁵⁰.

During the active learning process, the model hyperparameters are adaptively optimized by maximizing the log-likelihood in response to new training data. By maintaining well calibrated uncertainties, models are then capable of discovering new configurations on their own, circumventing a major problem of potentially missing unknown structures in a model's training set. Further, this methodology allows us to rely on the all-atom models only to remove fast degrees of freedom and obtain CG force training labels, so that the slow degrees of freedom are evolved directly using the learned low-resolution CG surrogate model. We illustrate the performance of SGPs trained on-the-fly for a pentane liquid structure consisting of 70 molecules, with the hydrogen degrees of freedom integrated out (Fig. 2).

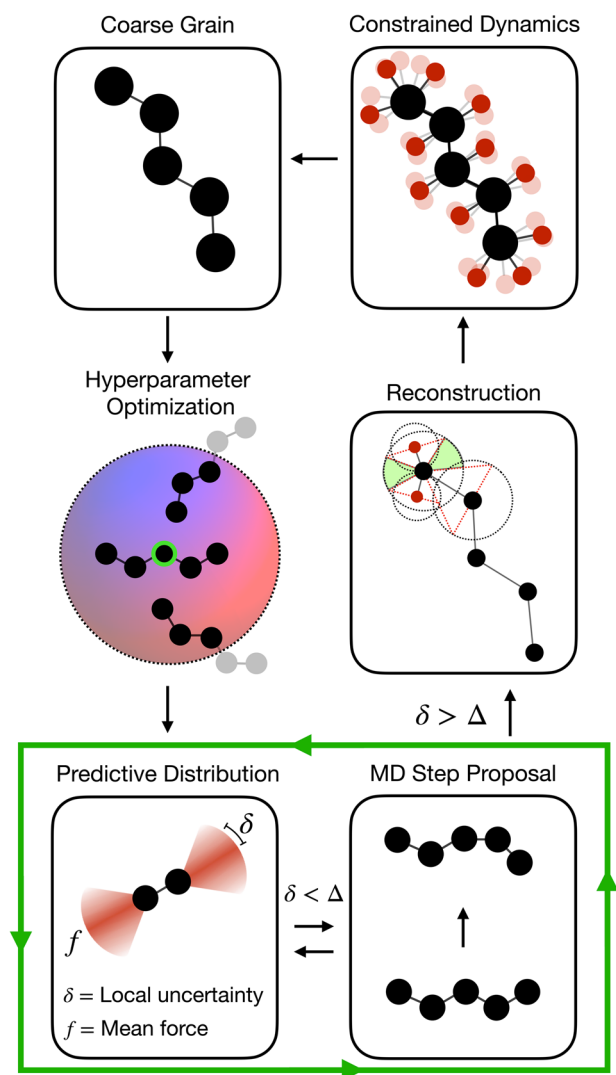


Fig. 1 A Schematic representation of the active learning workflow. FLARE-CG is an extension of the Fast Learning of Atomic Rare Events software^{49–51}. Here, we schematically demonstrate the workflow of the on-the-fly active learning training loop. An initial all-atom frame is run under constrained dynamics and coarse grained to obtain initial force labels. A select number of sparse environments are randomly added to the training set of the Gaussian process. The construction of model descriptors is performed, and the hyperparameters of the SGP updated by maximizing the log marginal likelihood. The model proposes a molecular dynamics step, along with force and local free energy uncertainties. If all local free energy uncertainties are below a tolerance threshold, the step is accepted. Otherwise, reconstruction is performed in order to collect more constrained dynamics training data. To perform reconstruction, we construct excluded volumes around already placed atomic centers in the system. New atom placements are proposed by randomly drawing an azimuthal and polar angle pair, which are subsequently accepted if they do not lie within regions of overlap between existing excluded volumes.

The all-atom to CG mapping sequence is shown schematically in Fig. 2a. Our approach requires a prior definition of CG sites. Although recent works have considered automating the CG site selection process^{37,56–58,63,64}, more work is needed to integrate these approaches with our proposed framework. We instead rely on chemical intuition to choose the CG mapping, where for the pentane system, we choose to map the all-atom system to carbon sites. In this example, the hydrogen atoms are integrated out,

effectively setting to zero their weight in the mapping function, defined in Eq. (3) of Methods.

Existing empirical and ML CG models often label coarse grained species by their underlying all-atom composition and bond topology, suggesting the end carbon atoms of pentane should be treated differently than those on the interior of the chain^{3,33}. In particular, CG carbons can be treated as different species types in the descriptor, defined in Eq. (15), depending on whether they are at the end of the molecular chain. Alternatively, all carbons can be treated as the same species type, and we can rely on the ML model to correctly learn CG forces from the geometry of the local environment structure.

We investigate the impact of this choice by exploring the performance of models of pentane with and without explicit labeling of end carbons as a different species in the descriptor. Hydrocarbon systems such as pentane liquid are a convenient test case, as many top-down empirical models are parameterized for them. In the following, we compare the accuracy of our ML CG models, whose all atom baseline is the OPLS force field, to an empirical CG force field with the same CG site mapping, OPLS-UA^{17,65}. Both the OPLS parameter set and its united atom variant were designed to capture the densities and heats of vaporization of organic liquids. While of the same molecular form, the OPLS-UA force field is further refined such that groups consisting of hydrogen atoms bound to a carbon atom are treated as a single interaction site.

For the ML CG approach, we find in Fig. 2c that the single-species and the two-species models provide similar force accuracy as well as the reproduction of structural properties compared to the all-atom baseline. Such comparisons are valuable because the less complex single-species models have shorter inference times. This arises from the scaling of the dimensions of ACE descriptors and kernel matrices with increasing number of species, associated with more computationally heavy linear algebra computations at inference. We are not claiming, however, that this finding for pentane is a general result expected to hold for other molecular systems.

Similarly, comparing to each other the learning of on-the-fly models for both single- and two-species realizations in Fig. 2c, we find both types of models achieve comparable force accuracy. We note that the predictive uncertainties of each model differ between single- and two-species as a result of having inherently different model complexity. In practice, this means that single-species models will tend to make fewer calls to the reference method (constrained all-atom dynamics). In this example, we use a predictive uncertainty value threshold, defined in Eq. (19) of Methods, of 0.02 for the two species case and 0.01 for the single species case to further highlight this contrast in behavior with respect to uncertainty. The final mean absolute force error, as a percentage of the mean absolute force component of the model predictions, on a test set lies around 9% for both single and two species models. For comparison, we also compute the error in forces predicted by the OPLS-UA force field on the same test set and find an error of 30%.

To emphasize that uncertainties are indeed predictive, in Fig. 2c, we show the (unitless) mean local PMF uncertainty given by Eq. (19) of Methods in each test set frame. A frame refers to a structure snapshot together with atomic force information. As a function of training set size, the uncertainty correlates directly with the trends followed by the mean absolute force errors on the test set. Even though the quantitative values of uncertainty and force test-set error do not agree, the crucial point is that the same trend behavior persists.

Figure 2b, d demonstrate that the inter- and intra-molecular properties, respectively, of the CG carbon sites match the behavior of the all-atom carbon atoms with high fidelity. The full carbon-carbon radial distribution functions are well reproduced, as are the end-to-end molecule chain distance. This distance is defined as the linear separation between the two carbons on the ends of each molecule.

In addition to pair distributions, more complex structural correlations have been suggested for assessing the fidelity of

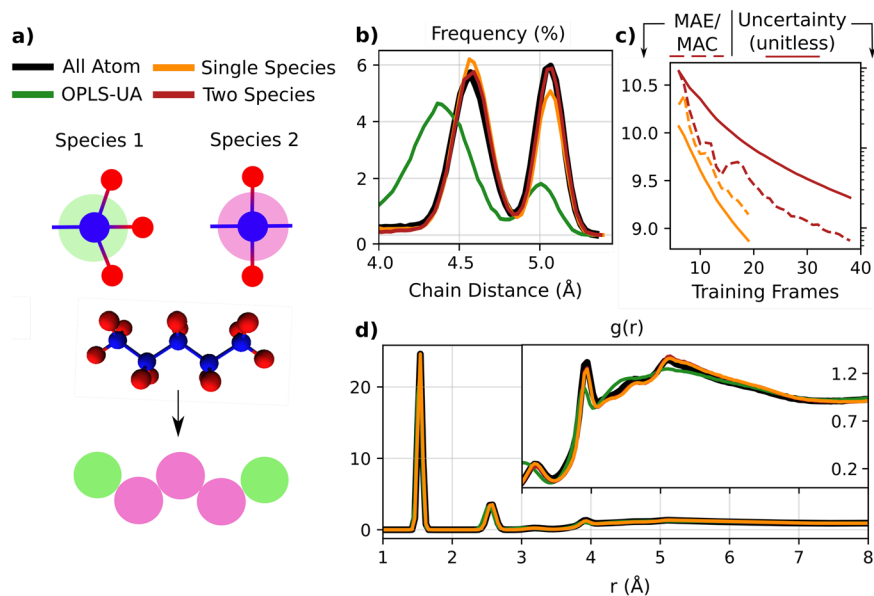


Fig. 2 Demonstration of active learning to coarse grained pentane. **a** Atoms in a pentane liquid are mapped directly to their carbon sites, integrating out hydrogen degrees of freedom. In two experiments, the beads are treated as either one or two different species, based on the underlying bond topology. **b** The end-to-end chain distance distribution of single and two species models are compared to the all-atom training baseline, as well as a common coarse grained force field for hydrocarbon liquids, OPLS-UA. **c** The learning rate for two and single species models is reported by showing the mean absolute force error on a test set as a relative percentage to the mean absolute force component in each test frame. We also report the mean free energy uncertainty of the model on the set, defined in Eq. (19), which is a unitless quantity. **d** The full carbon-carbon radial distribution function for single species, two species, all atom, and OPLS-UA simulations. The inset shows a zoomed-in view of the long-range structure.

coarse grained models^{66–70}. For example, it is possible for CG models to reproduce a set of distribution functions such as those shown, but the relative sampling between a set of states may still be inconsistent⁶⁷. To further establish the strength of our method, we analyze the possible correlated states for pentane.

As in Fig. 3, pentane bond and angle distributions are unimodal, and therefore do not meaningfully contribute to understanding correlations between structural states. Instead, we consider the trans (T) and gauche (G) dihedral angle regimes, akin to ref. 67. With two dihedral angles, we label the structural state pairings of pentane as TT, GG, and TG. We show in Fig. 3 that both the single and two species CG models capture the relative sampling of these states quite well compared to the all-atom baseline.

For completeness, we compare the accuracy of structural property predictions of our coarse grained models to an existing empirical (non-ML) model, OPLS-UA¹⁷, in both Figs. 2 and 3. We find far greater fidelity in the reproduction of structural properties with our SGPs compared to this empirical model, as well as the relative sampling of dihedral angle pairs. In particular, we emphasize that the OPLS-UA model describes two, three, and four body interactions that are typical in classical force fields. Despite this, the simple functional form of such models is insufficient for capturing structural correlations. This further motivates the use of ML approaches to PMF modeling.

Transferability across molecular systems

Bottom-up CG models are typically developed to preserve the partition function of the Boltzmann distribution maintained by thermostats in the all-atom MD simulation, thereby ensuring consistency with the AA thermodynamics. However, this implicitly assumes that the model will be specific to both the chemical make up of the all-atom system as well as the thermodynamic state points. Particularly in this setting, where models are specific to a given system, predictive model uncertainty is helpful in quantifying the distance of local CG configurations in the test set from

those in the training set. This enables us to systematically adjust and improve CG models.

To examine the degree to which SGP CG models trained on one system can be transferred to another system, we consider several ways in which a model trained on the pentane liquid can be applied and adapted to an octane liquid, as summarized graphically in Fig. 4a. Direct CG models of octane are trained on all-atom octane data on-the-fly using active learning. Unadapted CG models are trained on-the-fly on the pentane system for 100,000 time steps and then deployed directly on octane. Here, the SGP evolves the system in CG space over time for 100,000 steps, checking the uncertainty against the user defined threshold at each step, before being deployed on octane. Note that 100,000 does not refer to the number of AA constrained dynamics steps used when collecting more data. The models labeled as adapted (50) and adapted (100) have identical pentane training data, but their training sets are subsequently augmented on-the-fly for 50k and 100k time steps with octane training data, respectively, via the same procedure defined for the pentane on-the-fly loop. We also compare the performance of trained CG pentane models to the underlying AA results for octane. The goal of this experiment is to determine whether or not uncertainty can enable models previously trained on other systems to be extended to a new system more efficiently than starting over and without loss of accuracy. The results of each model type shown in Fig. 4 and Table 1 is the average of 40 independently trained models.

We find in Fig. 4c for the two species models that the inter- and intra-molecular structure of the carbon atoms is reproduced substantially better with the adapted models than their unadapted counterparts. This is further reflected in the end-to-end chain distance RMSE values given in Table 1. Note, however, that the unadapted single-species models shown in Fig. 4b give higher fidelity radial distribution functions compared to the unadapted two species models while using far less training data on pentane (see Table 1). This result can be attributed to the higher dimensionality of the two-species descriptors, which makes the

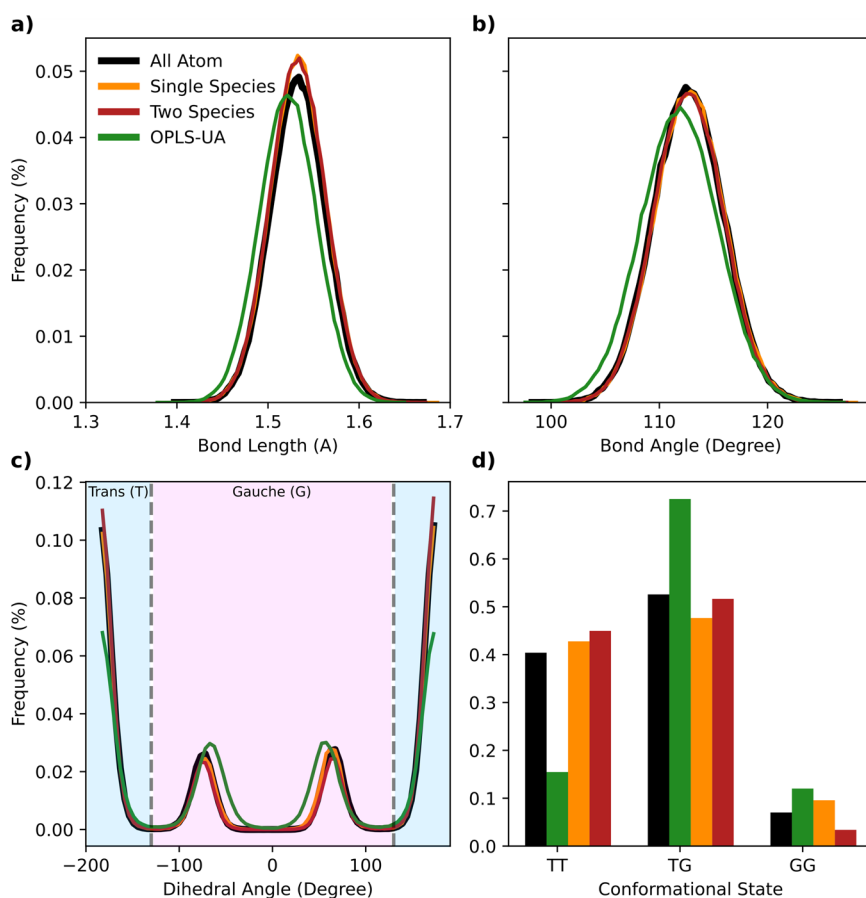


Fig. 3 Structural correlations captured via active learning. **a** The distribution of bond lengths in pentane molecules for single species, two species, and OPLS-UA models relative to the all-atom baseline. **b** Bond angle distributions of the same set of models. **c** Dihedral angle distributions of the same set of models. Vertical lines indicate the separation between trans and gauche dihedral conformations. **d** The relative sampling of the three dihedral pair states for pentane, trans-trans (TT), trans-gauche (TG) and gauche-gauche (GG).

two-species model more sensitive to changes in local environments. In the single-species case, for example, local environments near the end of the hydrocarbon chains look highly similar to those of octane. For the two-species case, this is not true as the local environments, within the 4.5 Å cutoff, of octane carbon sites will almost never contain both ends of the chain, which is likely to happen in pentane. As a result, to a single species kernel, octane looks more similar to pentane and it is able to extrapolate with a limited amount of data.

To emphasize the utility of active learning, in particular in the context of designing transferable models, we note the stark contrast in accuracy between adapted models trained with and without using uncertainty. In Fig. 5, we examine difficulties that arise in modeling the end-to-end chain lengths of single and two species approaches. Here, the non-active learning based models are trained from randomly selected octane data, with a single frame being used for the single species models and 15 frames for the two species, in line with the amount of data collected by on-the-fly models reported in Table 1. Each curve is the average over an ensemble of 40 models.

A relatively small amount of octane data is added while training these adapted models, and it is therefore crucial that the added data be maximally informative. In this regard, two-species models are more susceptible to poor performance when the available data is not representative of the differences between the original (pentane) system and the new (octane) system. Being a less descriptive model, the single species examples do not suffer from this to the same degree. It is quite evident that uncertainty based active learning enables accurate transferability of coarse grained

models where non-active learning could not. For completeness, we examine the learned distributions of active and non-active learning based models in the Supplementary Discussion, along with a detailed discussion on the structural correlations for this system.

To highlight the benefits of using an adapted model on systems outside of the training set as opposed to starting from scratch, we compare the computational cost of AA constrained dynamics reference computations as well as the resulting force errors and structural properties of adapted models compared to training an octane model from scratch (Table 1). Assuming the pentane data is already available, we find that by using an SGP containing data from a chemically similar system, far fewer AA reference calls to the new system are needed compared to starting from scratch. In both the single- and two-species cases, the force accuracy of adapted models with fewer active learning calls to the octane reference AA constrained dynamics method is higher than in direct models trained on octane from scratch which also requests more data in active learning. Additionally, over the course of the 50,000 and 100,000 on-the-fly training steps, we report in Table 1 the average number of frames in which the models request more octane data. We see explicitly that in all cases, models starting from pentane request fewer frames of data for the new system. This is a crucial point, as the reduction in required constrained dynamics calls of adapted models significantly reduces the computational cost of model training. We observe that the single- and two-species models display comparable force errors despite their disagreement in the reproduction of chain distance distributions, which we explore further in the next section.

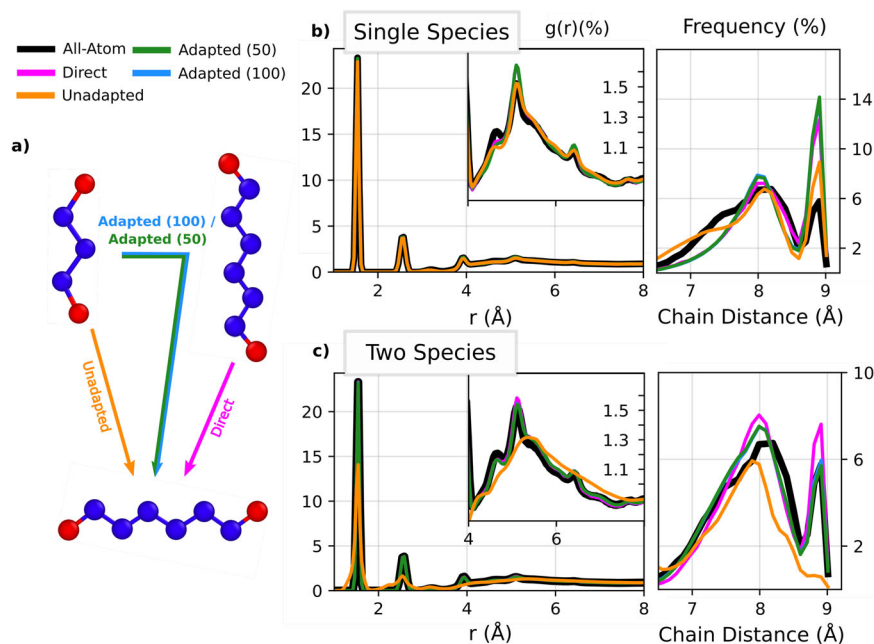


Fig. 4 Transferable coarse graining enabled by active learning. **a** A schematic representation of the computational experiment considered. Adapted (50 and 100) pentane models are those that have seen extra octane data, generated on the fly for 50,000 and 100,000 steps, respectively. Unadapted models are pentane models deployed directly on octane with no additional training. **b** The carbon-carbon radial distribution function and the end-to-end chain distance frequency are shown for single species models. The inset shows the long-range structure of the liquid **(c)** the carbon-carbon radial distribution function and the end-to-end chain distance frequency are shown for two species models, with the inset showing the long-range structure.

Table 1. Accuracy and efficiency of transferable coarse grained models.

	GP Model	Pentane Frames	Octane Frames	Force MAE/MAC (%)	Population Error (%)
Single Species	Unadapted Pentane	11.45	0	15.29	34.98
	Adapted (50) Pentane	11.45	1.77	10.10	53.71
	Adapted (100) Pentane	11.45	2.02	10.10	52.09
	Direct	0	12.1	10.18	49.72
Two Species	Unadapted Pentane	42.42	0	78.18	597.05
	Adapted (50) Pentane	42.42	15.26	10.14	9.84
	Adapted (100) Pentane	42.42	17.52	10.04	5.85
	Direct	0	54.73	10.20	13.86

The average number of all-atom LAMMPS calls over a set of 40 models, as well as the average total number of training frames in the GP, is reported for both single- and two-species models. In addition, we report the mean absolute force error on a test set, averaged over 10 models, as a percentage of the mean absolute force component (MAC) in each test frame, as well as the population error of the averaged chain distance distribution from the all-atom ground truth, defined in Eq. (1).

Correlating model robustness and PMF uncertainty

In this section, we aim to examine the correlation between the force error and error in structural properties determined by the PMF. Specifically, we focus our attention on the end-to-end chain distance distribution. Physically, the multi-modal distributions that appear in pentane and octane systems arise from the PMF minima at two values of dihedral angles along the chain. For shorter chains, fewer pairs of dihedral angles exist that are energetically favorable, whereas longer chains effectively become more flexible to bending.

To quantify the variation between a model and the AA system in the end-to-end chain distance distribution, we define the population error as

$$100 \times \left| \frac{p_1^{cg}/p_2^{cg}}{p_1^{aa}/p_2^{aa}} - 1 \right| \quad (1)$$

where $p_{1,2}$ are the populations (i.e., frequencies of sampling end-to-end chain distance values) of the first and second PMF basins of the AA or CG system.

With this definition, we can motivate this discussion by considering the relationship between mean force errors and population errors for the octane models discussed in the previous section. The mean force errors of the adapted single-species models are quite similar to the two-species models, while the population error between the single- and two-species approaches differ substantially (see Table 1). The population error of the adapted (100) single- and two-species models differ by nearly a factor of 10. Similarly, adapted versions of the single-species models achieve comparable force accuracy to their direct model counterparts, despite varied performance in capturing the relative sampling of states. This result points to a clear disconnect between force and PMF errors.

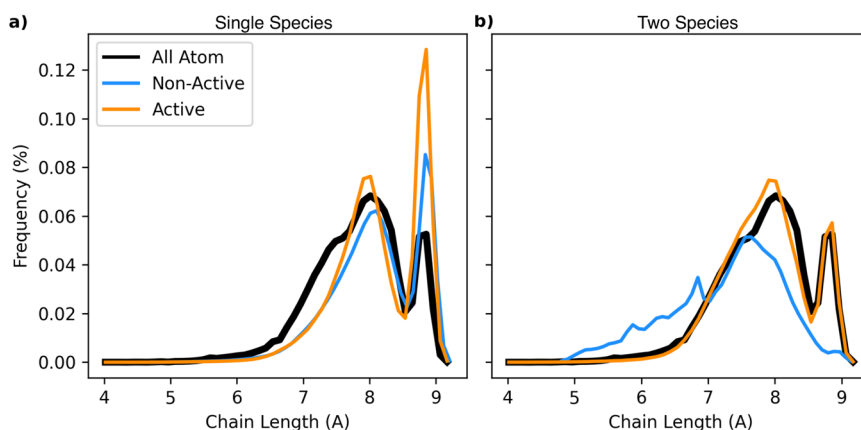


Fig. 5 Comparison of transferability with and without active learning. We show the end-to-end chain length distribution for single and two species CG models trained with and without active learning. **a** Compared to the all-atom baseline, single species models are shown. Each CG curve is an average over the results of 40 models. The non-active learned models has seen one additional frame of octane data, while the active learning models see on average 1.7 frames of data. **b** The same data is shown but for two species models. In this case, the non-active learned models each see 15 frames of octane data, while the active learning models see on average 15.2 additional frames.

For the remainder of this section, we focus our attention on the case of pentane specifically as it has a much more bimodal distribution in chain distances compared to octane, corresponding to a more challenging energy landscape in which to study population ratio errors. Even longer chains do not display such bimodal distributions at all (see Supplementary Fig. 4). To this end, in Fig. 6a, we show the end-to-end chain distance distributions for a set of 40 pentane models, each trained with subsets of 50 training frames and 50 sparse points per frame, drawn randomly from the same set of AA constrained dynamics data. It is clear that models vary widely in predicting end-to-end pentane chain distance distributions, despite all models exhibiting low errors on forces. Quantitatively, we show in Fig. 6b, the learning curves for the set of 40 models along with their corresponding population errors defined by equation (1). Black and red lines correspond to monotonically and non-monotonically decreasing quantities, respectively. Each model in the ensemble exhibits monotonic decrease in the force error with more force training data. At the same time, many models show non-monotonic evolution in the population error as a function of the training set size.

These trends can be understood by noting that pentane chain distance distributions are characterized by the difference in PMF values between the two basins. However, in our training, only PMF derivatives (forces) are used, which only implicitly constrain the model's estimate of PMF. As a result, in a test set with minimal representation of transition structures and high sampling of equilibrium ones, we expect that mean force errors will decrease with added data more rapidly than the PMF errors. Thus, improving the fidelity with which these models reproduce distributions of structural properties requires a large amount of force data, especially in the rarely sampled transition region. We note that despite the variability across models in the ensemble, the mean population error as a function of training set size indeed decreases, as does the mean force error. The insets of Fig. 6b show the mean force error and population error of the ensemble of models, along with the standard deviation represented with error bars.

The overlap in the distribution of force errors is far smaller than that of the population errors. In each step of the learning process, a model's PMF predictions are less constrained than the force predictions. As a result, there are effectively more learning pathways that quantities arising from the PMF, such as population errors considered here, can take towards a converged value. Because of the substantial overlap in distributions at different stages of training, many of these pathways are not necessarily monotonically decreasing functions.

One possible solution to minimize this variability would be to include PMF labels in the training set. However, free energies are difficult to compute. In the absence of such PMF labels, however, we argue that by training with force labels alone, the local PMF uncertainties can still be meaningfully interpreted, and that the uncertainties are able to capture useful information regarding the impact force data will have on the model's PMF predictions.

To connect uncertainty more directly to the performance of observable properties, we explore the local PMF uncertainties of CG environments and their relationship to Eq. (1). First, we define the molecular uncertainty as the average local PMF uncertainty on CG sites, i , belonging to molecules with an end-to-end chain distance L . Formally, this is given by

$$\sigma(L) = \frac{1}{N} \sqrt{\sum_t \sum_{m \in m_t(L)} \sum_{i \in m_t(L)} \delta \varepsilon_i^2} \quad (2)$$

where $\delta \varepsilon_i^2$ is the local PMF uncertainty of CG site i belonging to a molecule m of length L , evaluated on a frame in a test set t . The average is performed over all molecules whose end-to-end chain distance lies within a range $L \pm \Delta L$ in the test set.

In addition to the mean and standard deviation of force errors and population errors shown in Fig. 6b, we plot the molecular PMF uncertainty in equation (2) as a function of the training set size for a single model averaged over molecules within the transition state region. The molecular uncertainty follows the monotonic trend of the true population error averaged over the ensemble of models, while individual models do not necessarily follow such monotonic trends in population errors, as mentioned above. We suggest that, due to the wide variation across models, the local PMF uncertainty for a single model does not necessarily predict the true error of populations on a single model, but rather correlates with the expected average error of an ensemble of models.

We also find that force training labels alone constrain PMF predictions in a meaningful way. In particular, the usefulness of the force data that is added is also reflected in the local PMF uncertainties, supporting its use as a metric in on-the-fly training. To demonstrate this, we examine the molecular uncertainty defined by Eq. (2) as a function of chain distance for two different scenarios. Starting from a baseline model trained with 50 frames of pentane force data, we compute $\sigma(L)$ on an independent test set. Subsequently, we consider the addition of force labels in a molecule whose chain distance lies within the transition region (~ 4.8 Å, model A) compared to the addition of sites in a stretched molecule of chain distance ~ 6.3 Å, model B. The stretched configuration in model B is highly energetically unfavorable, and

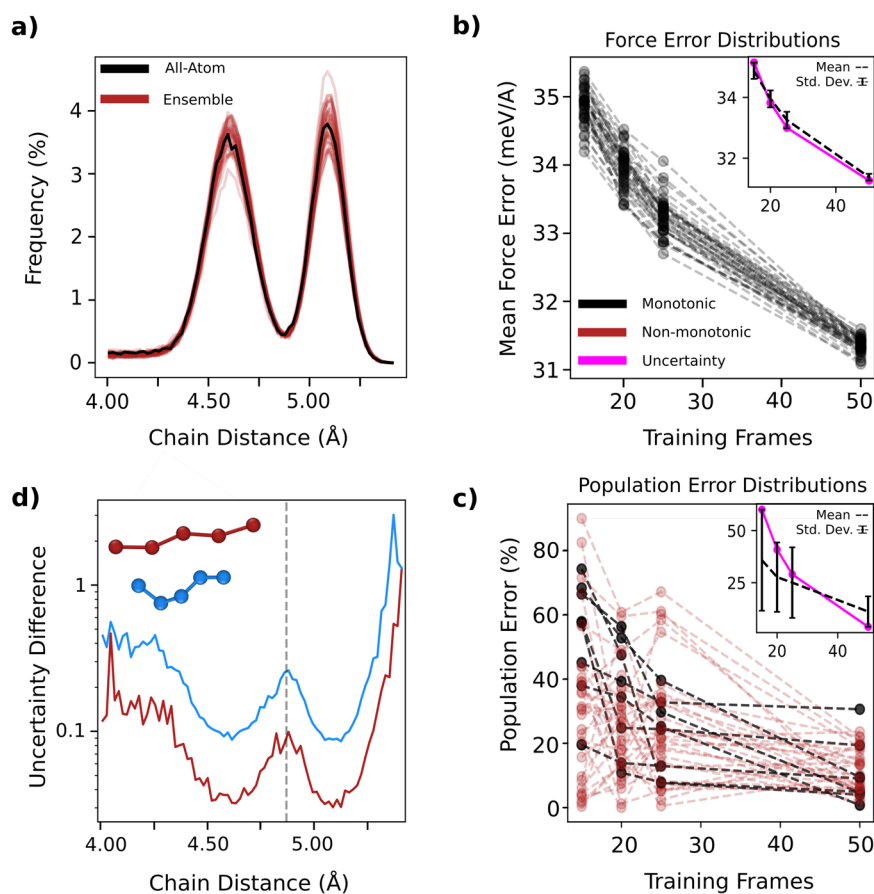


Fig. 6 Insufficiency in relative energies captured by uncertainty. **a** The end-to-end chain distance distribution of an ensemble of 40 pentane models, each trained with 50 frames of data **(b)** the distribution of mean absolute force errors and **(c)** the population errors defined in Eq. (2) for the pentane ensemble as a function of training set size. Red lines correspond to non-monotonic trends, while black are monotonic. The insets show the mean and standard deviation of the computed property (black), as well as the uncertainty defined in Eq. (2) averaged over models in the ensemble and over molecules within the transition region (magenta). The uncertainty does not share an axis with the force or population errors, but is simply a unitless quantity as described in the Methods **(d)** the effect of adding different types of data is shown. The carbon sites of a transition state molecule (blue) and unphysically stretched molecule (red) are added as new data to a model having 50 frames of data. The resulting change in molecular uncertainty from the baseline model is shown as a function of chain distance.

we expect that this data would have a less meaningful impact on the local PMF uncertainties of the model within the more well-explored regions of phase space.

In Fig. 6c, we plot differences of molecular uncertainties, $\sigma(L) - \sigma_{A,B}(L)$, between the baseline model and models A and B as a function of chain distance. Indeed, the transition state model shows a sharper decrease in uncertainty overall compared to the stretched model, directly indicating that the local PMF uncertainties are capable of identifying how force data will constrain the corresponding PMF predictions.

DISCUSSION

In this work, we have explored the utility of uncertainty aware machine learning models for coarse graining. The principled Bayesian uncertainty measure of Gaussian processes enables an on-the-fly active learning scheme for CG models that allows for automating the creation of training sets. A key aspect of the on-the-fly ML CG method is that it overcomes the time- and length-scale limitations of AA models by integrating the AA system in time only over fast degrees of freedom. In addition, by collecting training data only when necessary during MD, we eliminate the need for specifying a priori which configurations our model will need to be trained on. Instead, as the algorithm explores configurations, it automatically decides if they are new enough

to be added to the training set. We have made a direct comparison of all-atom and CG simulations using the same time step, although the CG models generalize well and are stable for larger time steps. This is expanded on in the Supplementary Discussion.

In addition, computational speed is a key metric in CG force fields. While fast classical all-atom force fields are used as a reference in this work, this need not be the case. For example, one could target ML models based on ab initio data that would be far more accurate. To this end, we quantify in Supplementary Fig. 1 the efficiency gained by using ML models at a coarse grained resolution as opposed to the full all-atom alternative. We find that these coarse grained models have the potential to be up to 50 times more efficient than an equivalent all-atom model when taking into account the largest stable timesteps.

Further to the point of designing more efficient CG models, the analysis we have done on the performance of single- versus two-species models is valuable. In this Gaussian process framework, there is unfavorable scaling with the number of species. As such, a single-species representation equates to a much faster model. Systems where this choice does not lead to significant accuracy loss in pair-wise and correlated distributions, such as the pentane system considered, will benefit greatly from improved model efficiency.

Moreover, the on-the-fly framework enables models to be transferred across molecular systems in a principled way. By the

locality assumption inherent in the structural descriptors, for local environments that are similar across systems forces will be predicted to be similar, at a given thermodynamic state. Similarly, environments that differ will result in small kernel values and contribute proportionally less to the prediction of forces on new environments, and the framework will request more data.

In practice, this approach could be used to develop CG models for common polymer backbones that could then be efficiently adapted to systems with differing functionalization. In this case, only data around new functional groups would be required, as opposed to starting from scratch. This would allow for more rapid development of coarse grained models in materials screening settings. We emphasize that there is a tradeoff in single- and two-species representations for transferable models. In the octane example considered in this work, there is an ease of direct transferability of single-species models without adaption. On the other hand, the achievable accuracy upon the addition of more data of the more descriptive two-species model may be more desirable. In general, it is difficult to anticipate the extent of this complex tradeoff.

The issue of model variance is further addressed herein as arising from limited direct PMF information. By training on time-averaged forces only, obtained from unbiased MD configurations, the PMF in transition regions is difficult to capture. What's more, examining force errors alone can seem to suggest that models should always improve with more data. In fact, we show that over an ensemble of models, average force errors indeed decrease monotonically with more data, while PMF errors do not necessarily decrease for each model. As a result, properties arising from PMF values, such as population ratios of stable configurations, can vary significantly and converge slowly in the training process. We find only that the average over a set of models will improve such property predictions. We demonstrate that despite the lack of direct PMF labels in the training set, providing forces still imparts meaningful PMF information to the model that is reflected in the monotonic behavior of uncertainties. This allows a better understanding of model robustness where traditional metrics such as force errors are insufficient.

We note that the molecular local PMF uncertainty plotted in Fig. 6b, c does not correspond quantitatively to the mean absolute force error or population error. While some work has been done on understanding how to more concretely link uncertainty and performance⁷¹, it is not well understood how we can, for example, translate specific values of uncertainty directly to the variance of observed structural properties over an ensemble of models. This remains an open question and demands systematic investigation even for force-fields in all-atom simulations. Moreover, the particular form in which we analyze molecular uncertainty in Eq. (2) is a choice. Other functional forms of molecular uncertainty can be conceptualized that may or may not provide deeper physical insight. This will require careful considerations in future works.

Finally, we note that a particular challenge in making the proposed method more generally applicable is the reconstruction of all-atom configurations to enable constrained dynamics during the active learning loop. In order to seamlessly go between the all-atom and coarse grained representations, a scheme for recovering lost degrees of freedom must be designed. In many cases, this is done with techniques designed for specific systems (as we have utilized in this work), or brute force methods such as a multi-stage compression and expansion of the system box to equilibrate replaced degrees of freedom. Recent works have begun to examine this problem from the perspective of machine learning^{56–58,63,64}, but such approaches require pre-selected training sets to learning the CG mapping that would require additional work to reconcile with our active learning scheme. Doing so will enable a much broader study of materials at a variety of resolutions.

METHODS

Potential of mean force

While many approaches exist for designing coarse grained models, top-down approaches lack thermodynamic consistency. A thermodynamically consistent coarse grained model is one for which the Boltzmann distribution of the coarse grained sites is the same as that implied by the underlying all-atom model²⁷. In principle, all of the thermodynamics of the all-atom system may be recovered from the resulting potential of mean force.

Suppose we have an all-atom model consisting of n atoms at a set of positions $\mathbf{r}^n = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n\}$ and governed by a potential $u(\mathbf{r}^n)$. We consider a mapping of the AA system to N coarse grained sites, given by $\mathbf{R}^N = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$, of the form

$$\mathbf{R}_j = \mathbf{M}_j(\mathbf{r}^n) = \sum_{i=1}^n c_{ij} \mathbf{r}_i \quad (3)$$

where j denotes the index of the coarse grained site and i the index of each atom. The mapping coefficients c_{ij} must satisfy $\sum_j c_{ij} = 1$ to ensure translation invariance of the resulting model²⁷.

The above information is enough to define the Boltzmann probability distribution of the AA system within the canonical ensemble, so that

$$p(\mathbf{r}^n) = \frac{\exp(-u(\mathbf{r}^n)/k_B T)}{Z_r} \quad (4)$$

with Boltzmann constant k_B , temperature T , and partition function Z_r . Let the coarse grained system be governed by its own potential, $U(\mathbf{R}^N)$, so that we can define a Boltzmann probability for the coarse grained sites at the same state points as

$$P(\mathbf{R}^N) = \frac{\exp(-U(\mathbf{R}^N)/k_B T)}{Z_R} \quad (5)$$

For a consistent coarse grained model, we would like the probability of sampling a given coarse grained configuration to be the same as if we had sampled the sites from the AA system. In other words,

$$P(\mathbf{R}^N) = \int p(\mathbf{r}^n) \prod_{j=1}^N \delta(\mathbf{M}_j(\mathbf{r}^n) - \mathbf{R}_j) d\mathbf{r}^n \quad (6)$$

This requirement defines the potential of mean force, $U(\mathbf{R}^N)$, as a coordinate-dependent free energy surface in terms of an underlying AA model. In particular,

$$\exp(-U(\mathbf{R}^N)/k_B T) \propto \int p(\mathbf{r}^n) \prod_{j=1}^N \delta(\mathbf{M}_j(\mathbf{r}^n) - \mathbf{R}_j) d\mathbf{r}^n \quad (7)$$

The PMF allows us to define a force on each coarse grained site as a gradient of the free energy. This force captures not only energetic effects, but also the entropic contributions to the free energy arising from the degrees of freedom being integrated out.

Formally, for mappings that take a group of atoms to their center of mass, and for which no atom belongs to more than one coarse grained site, the mean force on each site is Eq.²⁷

$$\mathbf{F}_j(\mathbf{R}^N) = \left\langle \sum_{i \in S_j} \mathbf{f}_i \right\rangle_{\mathbf{R}^N} \quad (8)$$

where S_j is the set of atoms involved in the mapping to site j , \mathbf{f}_i is the atomistic force on atom i , and the average is a weighted ensemble average defined as

$$\langle g(\mathbf{r}^n) \rangle_{\mathbf{R}^N} = \frac{\int g(\mathbf{r}^n) e^{-u(\mathbf{r}^n)/k_B T} \prod_{j=1}^N \delta(\mathbf{M}_j(\mathbf{r}^n) - \mathbf{R}_j) d\mathbf{r}^n}{\int e^{-u(\mathbf{r}^n)/k_B T} \prod_{j=1}^N \delta(\mathbf{M}_j(\mathbf{r}^n) - \mathbf{R}_j) d\mathbf{r}^n} \quad (9)$$

for an arbitrary function g of the atomistic coordinates. Note that it is not a requirement that every atom have a non-zero weight in at least one mapping to a coarse grained site. In the hydrocarbon systems we consider, the hydrogen atoms have corresponding coefficients $c_{ij} = 0$ for all CG sites, j .

To model the PMF independent of atomic details, we must define a functional approximation to $U(\mathbf{R}^N)$ as well as how the parameters of this function will be fit. In the following section, we describe the technique used to reconstruct all-atom information for data collection, as well as the sparse Gaussian process model used to approximate the PMF. The mean force labels in Eq. (8) can be estimated using constrained molecular dynamics.

All-atom reconstruction

In order to go from the coarse grained resolution back to an AA representation, we must define a method for recovering AA degrees of freedom. For the n-alkane systems studied in this work, we use the idea of excluded volumes. In all cases, we have chosen to map AA structures onto a carbon site based structure representation, leaving only hydrogen atoms to be reconstructed.

A schematic of this reconstruction scheme is provided in Fig. 1. Initially, excluded volumes of radius R_C are placed around each CG carbon site. For each site, we define a procedure for reconstructing all hydrogen atoms that are bonded to the central site. We wish to place a hydrogen atom a distance R_{C-H} away from its bonded carbon, which requires a polar and azimuthal angle to define a complete coordinate. A random pair of angles is drawn, and it is determined whether or not the resulting point lies within a restricted region.

Such restricted regions are defined in terms of overlapping excluded volumes. In particular, a conical region is constructed whose central axis is defined by the vector from the central carbon to the neighboring atom creating an excluded volume overlap. The angular extent of the conical region is determined by the angle between the central axis and the lines connecting the central carbon to the circular intersection of the edges of the excluded environments.

Should the angle between the vector from the carbon to the new hydrogen site and the central axis vector of any restricted region be less than that regions angular extent, a new pair of polar and azimuthal angles is drawn for the hydrogen until a position that does not lie in an overlapping region is found. This process is iterated for each hydrogen atom bonded to the carbon site before proceeding to the next carbon atom. Note that in general, R_{C-H} and R_C need not be equal, and in our implementation, R_{C-H} is taken to be the experimental value of 1.118 Å.

In addition, for each hydrogen atom placed, a new excluded volume is introduced into the system centered at this location, with a radius R_H . In this way, the presence of reconstructed degrees of freedom can influence subsequent placements, allowing for easy avoidance of non-physical AA reconstructions. We find stable reconstruction results when setting R_C and R_H to 1.6 Å and 1.1 Å, respectively. These quantities are comparable with the experimental C-H and C-C bond lengths of 1.118 Å and 1.531 Å.

In order to reduce the time spent on sampling, we place a restriction on our reconstruction algorithm such that if an acceptable sample is not drawn in a predefined amount of time, the excluded volume radii are reduced by 10%. However, we rarely find this to be necessary.

Once all hydrogen atoms have been reconstructed, energy minimization (see Supplementary Methods) is performed, after which hydrogen velocities are drawn at random from a Boltzmann distribution at 250K and equilibrated with a 0.5 fs timestep for 20,000 steps.

Sparse Gaussian processes

Bottom-up coarse graining models often require an accurate description of many-body interactions. Here, we use sparse Gaussian processes as implemented in the FLARE software to do

so. We outline the methodology here, with more detail available in ref. 49,50. The methods below closely mirror the procedures presented by FLARE, but more specifically with application to coarse grained potentials.

We assume that the PMF can be modeled with a purely local function. In particular, given a set of coarse grained coordinates $\mathbf{R}^N = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$ with chemical-type identities $\{s_1, s_2, \dots, s_N\}$, the PMF is given by a sum of local free energy contributions from the CG sites in the form

$$U(\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N; s_1, s_2, \dots, s_N) = \sum_i^N \epsilon(s_i, \rho_i) \quad (10)$$

where ρ_i is some description of the local environment of CG site i . In particular, the environment is the set of distance vectors between neighboring sites, such that $j \neq i$, within a cutoff radius of $r_{\text{cut}}^{(s_i, s_j)}$ that can, in general, be species dependent. Formally,

$$\rho_i = \{(s_j, \mathbf{r}_{ij}) \mid r_{ij} < r_{\text{cut}}^{(s_i, s_j)}\} \quad (11)$$

Following the Atomic Cluster Expansion approach introduced by Drautz⁶², the local environment of a site can be projected onto single-particle basis functions

$$\langle \phi_{n\ell m} | \rho_i \rangle = \phi_{n\ell m}(\mathbf{r}_{ij}) \quad (12)$$

where we choose for ϕ a decomposition into radial and spherical harmonic components with cutoff function, c ,

$$\phi_{n\ell m}(\mathbf{r}_{ij}) = R_n \left(\frac{r_{ij}}{r_{\text{cut}}^{(s_i, s_j)}} \right) \cdot Y_{\ell m}(\hat{\mathbf{r}}_{ij}) \cdot c(r_{ij}, r_{\text{cut}}^{(s_i, s_j)}) \quad (13)$$

In this work, we use Chebyshev polynomials for the radial basis and spherical harmonics for the angular basis. A covariant tensor can be constructed by summing the basis functions over sites within the local environment as

$$C_{isn\ell m} = \sum_{j \in \rho_i} \delta_{s, s_j} \phi_{n\ell m}(\mathbf{r}_{ij}) \quad (14)$$

Utilizing the sum rule of the spherical harmonic functions, a rotationally invariant descriptor, \mathbf{d}_i is

$$d_{is_1 s_2 n_1 n_2 \ell} = \sum_{m=-\ell}^{\ell} C_{is_1 n_1 \ell m} C_{is_2 n_2 \ell m} \quad (15)$$

The parameters n , ℓ , and r_{cut} are hyperparameters that we manually specify. To do so, we maximize the marginal log-likelihood (see Supplementary Fig. 3), and use the values $n = 12$, $\ell = 5$, and $r_{\text{cut}} = 4.5$ Å, independent of chemical type.

To complete the description of the sparse Gaussian process, we must define a kernel that compares the local environment descriptors. As in Ref. 50, we choose a kernel that resembles the smooth overlap of atomic potentials (SOAP) kernel^{43,54,55} and takes the form

$$k(\mathbf{d}_1, \mathbf{d}_2) = \sigma^2 \left(\frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{d_1 d_2} \right)^\xi \quad (16)$$

Here, the hyperparameter, σ , is optimized by maximizing the marginal log-likelihood during on-the-fly training, and ξ is a chosen parameter that can be used to increase the body-order of the model.

For a set of sparse coarse grained environments, S , that is a subset of a larger training set, F , a prediction of the local free energy on a new environment ρ_i can be cast as a sum over the sparse points:

$$\epsilon(\rho_i) = \sum_{s \in S}^{N_s} k(\mathbf{d}_i, \mathbf{d}_s) \alpha_s \quad (17)$$

where

$$\boldsymbol{\alpha} = (\sigma_n^{-1} \mathbf{K}_{SF} \mathbf{K}_{FS} + \mathbf{K}_{SS})^{-1} \mathbf{K}_{SF} \mathbf{y} \quad (18)$$

where \mathbf{y} is the vector of training force labels, \mathbf{K}_{SF} the matrix of kernel values between the sparse set and training set, and \mathbf{K}_{SS} the matrix of kernel values between points in the sparse set alone. The noise hyperparameter, σ_n , quantifies the inherent uncertainty present in the training labels and is another hyperparameter that is tuned via maximizing the marginal log-likelihood.

Such mean predictions with the sparse Gaussian process also allow for posterior predictive distribution variances. For SGP's, computing the variance requires approximate methods, where we choose in this work to use the Deterministic Training Conditional (DTC) approximation⁷². As in ref. ⁵⁰, we use a further simplified form that is the predictive variance of a fictitious Gaussian process trained on local free energies of the sparse environments alone. The resulting uncertainty on local free energies is

$$\tilde{V}(\epsilon) = \frac{k_{\epsilon\epsilon} - \mathbf{k}_{\epsilon S} \mathbf{K}_{SS}^{-1} \mathbf{k}_{S\epsilon}}{\sigma^2} \quad (19)$$

Lying between 0 and 1, this form gives us a unitless measure in defining uncertainty thresholds during on-the-fly training. We find a relative tolerance of 0.02 to perform well. This is found to be stable from our empirical observations and is consistent with similar values used in ref. ⁵⁰.

The resulting models, following the on-the-fly training trajectory, can be simplified such that the summation over sparse points in the predictive distribution can be computed once and used for all future predictions⁵⁰. This simplification allows for efficient inference upon deployment of the SGPs.

Computational details

The LAMMPS package⁷³ was used for all production simulations. The OPLS-AA⁶⁵ and OPLS-UA¹⁷ force fields were used in addition to the SGP models. The on-the-fly MD active learning loop was performed using the Atomic Simulation Environment (ASE)⁷⁴, and mapped SGP potentials were used in a custom implementation of a LAMMPS pair-style, available as an executable in the FLARE repository^{49,50}. We have adapted the on-the-fly framework to coarse grained applications, and the software package, FLARE CG, is available upon request. The specific version of the FLARE code used in the production of these results is also available upon request. Parsing of LAMMPS output files for RDFs was done primarily through Ovito⁷⁵. Parameters for all simulations are provided in the Supplementary Methods.

DATA AVAILABILITY

All input and output files from the FLARE training and LAMMPS simulations are available upon request.

CODE AVAILABILITY

The code used to perform coarse graining and the integration with FLARE will be made public on github. The version used for the data collected herein is otherwise available upon request.

Received: 28 October 2022; Accepted: 5 December 2023;

Published online: 08 January 2024

REFERENCES

- Allen, M. P. & Tildesley, D. J. *Computer Simulation of Liquids* (Oxford University Press, 1987).
- Frenkel, D. & Smit, B. *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, 2002).
- 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.* **111**, 7812–7824 (2007).

- Kmieciak, S. et al. Coarse-grained protein models and their applications. *Chem. Rev.* **116**, 7898–7936 (2016).
- Skjerve, A. A. et al. All-atom lipid bilayer self-assembly with the amber and charmm lipid force fields. *Chem. Commun.* **51**, 4402 (2015).
- Woo, S. Y. & Lee, H. All-atom simulations and free energy calculations of coiled-coil peptides with lipid bilayers: binding strength, structural transition, and effect on lipid dynamics. *Sci. Rep.* **6**, 22299 (2016).
- Bonomi, M., Barducci, A. & Parrinello, M. Reconstructing the equilibrium Boltzmann distribution from well-tempered metadynamics. *J. Comput. Chem.* **30**, 1615–1621 (2009).
- Torrie, G. M. & Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *J. Comput. Phys.* **23**, 187–199 (1977).
- Barducci, A., Bonomi, M. & Parrinello, M. Metadynamics. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **1**, 826–843 (2011).
- Souaille, M. & Roux, B. Extension to the weighted histogram analysis method: combining umbrella sampling with free energy calculations. *Comput. Phys. Commun.* **135**, 40–57 (2001).
- Baran, I & Ržysko, W. Application of a coarse-grained model for the design of complex supramolecular networks. *Mol. Syst. Des. Eng.* **5**, 484–492 (2020).
- Ponder, J. W. & Case, D. A. Force fields for protein simulations. *Adv. Prot. Chem.* **66**, 27–85 (2003).
- Monticelli, L. et al. The martini coarse-grained force field: extension to proteins. *J. Chem. Theory Comput.* **4**, 819–834 (2008).
- López, C. A. et al. Martini coarse-grained force field: extension to carbohydrates. *J. Chem. Theory Comput.* **5**, 3195–3210 (2009).
- de Jong, D. H. et al. Improved parameters for the martini coarse-grained protein force field. *J. Chem. Theory Comput.* **9**, 687–697 (2013).
- Yesylevskyy, S. O., Schäfer, L. V., Sengupta, D. & Marrink, S. J. Polarizable water model for the coarse-grained martini force field. *PLoS Comput. Biol.* **6**, 1000810 (2010).
- Beveridge, D. L. & Jorgensen, W. L. The opls potential functions for proteins. energy minimizations for crystals of cyclic peptides and crambin. *Annu. Rev. Biophys. Bioeng.* **110**, 18 (1988).
- Koliński, A. Protein modeling and structure prediction with a reduced representation. *Acta Biochim. Pol.* **51**, 349–71 (2004).
- Martin, M. G. & Siepmann, J. I. Transferable potentials for phase equilibria. 1. united-atom description of n-alkanes. *J. Phys. Chem. B* **102**, 2569–2577 (1998).
- Rohl, C. A., Strauss, C. E. M., Misura, K. M. S. & Baker, D. Protein structure prediction using Rosetta. *Methods Enzymol.* **383**, 66–93 (2004).
- Liwo, A. et al. A unified coarse-grained model of biological macromolecules based on mean-field multipole-multipole interactions. *J. Mol. Model.* **20**, 2306 (2014).
- Noid, W. G. Perspective: coarse-grained models for biomolecular systems. *J. Chem. Phys.* **139**, 90901 (2013).
- Chen, Y. et al. Machine learning implicit solvation for molecular dynamics. *J. Chem. Phys.* **155**, 084101 (2021).
- Mukherjee, B., Site, L. D., Kremer, K. & Peter, C. Derivation of coarse grained models for multiscale simulation of liquid crystalline phase transitions. *J. Phys. Chem. B* **116**, 8474–8484 (2012).
- Sinitskiy, A. V. & Voth, G. A. Quantum mechanics/coarse-grained molecular mechanics (QM/CG-MM). *J. Chem. Phys.* **148**, 014102 (2018).
- Mironenko, A. V. & Voth, G. A. Density functional theory-based quantum mechanics/coarse-grained molecular mechanics: theory and implementation. *J. Chem. Theory Comput.* **16**, 6329–6342 (2020).
- Noid, W. G. et al. The multiscale coarse-graining method. i. a rigorous bridge between atomistic and coarse-grained models. *J. Chem. Phys.* **128**, 244114 (2008).
- Hills, Jr., R. D., Lu, L. & Voth, G. A. Multiscale coarse-graining of the protein energy landscape. *PLoS Comput. Biol.* **6** (2010).
- 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* **111**, 4116–4127 (2007).
- Larini, L., Lu, L. & Voth, G. A. The multiscale coarse-graining method. vi. implementation of three-body coarse-grained potentials. *J. Chem. Phys.* **132**, 164107 (2010).
- Shell, M. S. The relative entropy is fundamental to multiscale and inverse thermodynamic problems. *J. Chem. Phys.* **129**, 144108 (2008).
- Izvekov, S. & Voth, G. A. Multiscale coarse graining of liquid-state systems. *J. Chem. Phys.* **123**, 134105 (2005).
- Wang, J. et al. Machine learning of coarse-grained molecular dynamics force fields. *ACS Cent. Sci.* **5**, 755–767 (2019).
- Zhang, L., Han, J., Wang, H., Car, R. & E. W. Deepcg: constructing coarse-grained models via deep neural networks. *J. Chem. Phys.* **149**, 034101 (2018).
- Ruza, J. et al. Temperature-transferable coarse-graining of ionic liquids with dual graph convolutional neural networks. *J. Chem. Phys.* **153**, 164501 (2020).
- Husic, B. E. et al. Coarse graining molecular dynamics with graph neural networks. *J. Chem. Phys.* **153**, 194101 (2020).

37. Wang, W. & Gómez-Bombarelli, R. Coarse-graining auto-encoders for molecular dynamics. *Npj Comput. Mater.* **5** (2019).
38. John, S. T. & Csányi, G. Many-body coarse-grained interactions using Gaussian approximation potentials. *J. Phys. Chem. B* **121**, 10934–10949 (2017).
39. Wang, J., Chmiela, S., Müller, K.-R., Noé, F. & Clementi, C. Ensemble learning of coarse-grained molecular dynamics force fields with a kernel approach. *J. Chem. Phys.* **152**, 194106 (2020).
40. Scherer, C., Scheid, R., Andrienko, D. & Bereau, T. Kernel-based machine learning for efficient simulations of molecular liquids. *J. Chem. Theory Comput.* **16**, 3194–3204 (2020).
41. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
42. Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **134**, 074106 (2011).
43. Bartók, A. P., Payne, M. C., Kondor, R. & Csányi, G. Gaussian approximation potentials: the accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.* **104**, 136403 (2010).
44. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 019902 (2013).
45. Li, Z., Kermode, J. R. & De Vita, A. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.* **114**, 096405 (2015).
46. Glielmo, A., Sollich, P. & De Vita, A. Accurate interatomic force fields via machine learning with covariant kernels. *Phys. Rev. B* **95**, 214302 (2017).
47. Glielmo, A., Zeni, C. & De Vita, A. Efficient nonparametric n-body force fields from machine learning. *Phys. Rev. B* **97**, 184307 (2018).
48. Schütt, K. T., Sauceda, H. E., Kindermans, P. J., Tkatchenko, A. & Müller, K. R. SchNet—a deep learning architecture for molecules and materials. *J. Chem. Phys.* **148**, 241722 (2018).
49. Vandermause, J. et al. On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events. *Npj Comput. Mater.* **6**, 20 (2020).
50. Vandermause, J., Xie, Y., Lim, J. S., Owen, C. J. & Kozinsky, B. Active learning of reactive Bayesian force fields applied to heterogeneous catalysis dynamics of h/pt. *Nat. Commun.* **13**, 5183 (2022).
51. Xie, Y., Vandermause, J., Sun, L., Cepellotti, A. & Kozinsky, B. Bayesian force fields from active learning: study of inter-dimensional transformation of stanene. *Npj Comput. Mater.* **7**, 40 (2021).
52. Bätzner, S. et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453 (2022).
53. Musaelian, A. et al. Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **14**, 579 (2023).
54. Bartók, A. P., Kondor, R. & Csányi, G. On representing chemical environments. *Phys. Rev. B* **87**, 184115 (2013).
55. Bartók, A. P. & Csányi, G. Gaussian approximation potentials: a brief tutorial introduction. *Int. J. Quantum Chem.* **115**, 1051–1057 (2015).
56. Peng, J., Yuan, C., Ma, R. & Zhang, Z. Backmapping from multiresolution coarse-grained models to atomic structures of large biomolecules by restrained molecular dynamics simulations using Bayesian inference. *J. Chem. Theory Comput.* **15**, 3344–3353 (2019).
57. Li, W., Burkhart, C., Políńska, P., Harmandaris, V. & Doxastakis, M. Backmapping coarse-grained macromolecules: an efficient and versatile machine learning approach. *J. Chem. Phys.* **153** (2020).
58. Wang, W. et al. Generative coarse-graining of molecular conformations. Preprint at <http://arXiv.org/abs/2201.12176> (2022).
59. Johansson, A. et al. Micron-scale heterogeneous catalysis with Bayesian force fields from first principles and active learning. Preprint at <https://arxiv.org/abs/2204.12573> (2022).
60. Xie, Y. et al. Uncertainty-aware molecular dynamics from Bayesian active learning: phase transformations and thermal transport in sic. *Npj Comput. Mater.* **9**, 36 (2023).
61. Loeffler, T. D., Patra, T. K., Chan, H. & Sankaranarayanan, S. K. R. S. Active learning a coarse-grained neural network model for bulk water from sparse training data. *Mol. Syst. Des. Eng.* **5**, 902–910 (2020).
62. Drautz, R. Atomic cluster expansion for accurate and transferable interatomic potentials. *Phys. Rev. B* **99**, 014104 (2019).
63. Chen, L.-J., Qian, H.-J., Lu, Z.-Y., Li, Z.-S. & Sun, C.-C. An automatic coarse-graining and fine-graining simulation method: application on polyethylene. *J. Phys. Chem. B* **110**, 24093–24100 (2006).
64. Hunkler, S., Lemke, T., Peter, C. & Kukharensko, O. Back-mapping based sampling: coarse grained free energy landscapes as a guideline for atomistic exploration. *J. Chem. Phys.* **151** (2019).
65. 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.* **118**, 11225–11236 (1996).
66. Mullinax, J. W. & Noid, W. G. A generalized-lyon-born-green theory for determining coarse-grained interaction potentials. *J. Phys. Chem. C* **114**, 5661–5674 (2010).
67. Rudzinski, J. F. & Noid, W. G. Investigation of coarse-grained mappings via an iterative generalized yvon-born-green method. *J. Phys. Chem. B* **118**, 8295–8312 (2014).
68. Scherer, C. & Andrienko, D. Understanding three-body contributions to coarse-grained force fields. *Phys. Chem. Chem. Phys.* **20**, 22387–22394 (2018).
69. Lu, L., Dama, J. F. & Voth, G. A. Fitting coarse-grained distribution functions through an iterative force-matching method. *J. Chem. Phys.* **139**, 121906 (2013).
70. Rudzinski, J. F. & Noid, W. G. The role of many-body correlations in determining potentials for coarse-grained models of equilibrium structure. *J. Phys. Chem. B* **116**, 8621–8635 (2018).
71. Imbalzano, G. et al. Uncertainty estimation for molecular dynamics and sampling. *J. Chem. Phys.* **154**, 074102 (2021).
72. Candela-Quinonero, J. & Rasmussen, C. E. A unifying view of sparse approximate Gaussian process regression. *J. Mach. Learn. Res.* **6**, 1939–1959 (2005).
73. Plimpton, S. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.* **117**, 1–19 (1995).
74. Larsen, A. H. et al. The atomic simulation environment - a Python library for working with atoms. *J. Condens. Matter Phys.* **29**, 273002 (2017).
75. Stukowski, A. Visualization and analysis of atomistic simulation data with ovito - the open visualization tool. *Model. Simul. Mater. Sci. Eng.* **18**, 015012 (2010).

ACKNOWLEDGEMENTS

The authors thank Yu Xie, Lixin Sun, Cameron Owen, Anders Johansson, Noah Bice, and Benjamin Jensen for helpful discussions. This work was supported by a NASA Space Technology Graduate Research Opportunity, by the NSF through the Harvard University Materials Research Science and Engineering Center Grant No. DMR-2011754, and by a Multidisciplinary University Research Initiative sponsored by the Office of Naval Research, under Grant N00014-20-1-2418. All computational experiments utilized the resources provided and maintained by Harvard FAS Research Computing.

AUTHOR CONTRIBUTIONS

B.R.D. developed the coarse graining extension code to FLARE and performed all computational experiments. J.V. led the development of the all-atom FLARE code and helped with its adaptation and with interpretation of trained models. N.M. assisted in assessing molecular dynamics results and the physical interpretations of the data, as well as contributed to figure design. B.K. supervised the work and contributed to the theoretical developments. B.R.D. wrote the manuscript, and all authors contributed to manuscript preparation.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-023-01183-5>.

Correspondence and requests for materials should be addressed to Blake R. Duschatko or Boris Kozinsky.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024