

# AMARO: All Heavy-Atom Transferable Neural Network Potentials of Protein Thermodynamics

Antonio Mirarchi, Raúl P. Peláez, Guillem Simeon, and Gianni De Fabritiis\*



Cite This: *J. Chem. Theory Comput.* 2024, 20, 9871–9878



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

**ABSTRACT:** All-atom molecular simulations offer detailed insights into macromolecular phenomena, but their substantial computational cost hinders the exploration of complex biological processes. We introduce Advanced Machine-learning Atomic Representation Omni-force-field (AMARO), a new neural network potential (NNP) that combines an O(3)-equivariant message-passing neural network architecture, TensorNet, with a coarse-graining map that excludes hydrogen atoms. AMARO demonstrates the feasibility of training coarser NNP, without prior energy terms, to run stable protein dynamics with scalability and generalization capabilities.



## 1. INTRODUCTION

Molecular events at the individual macromolecule level reveal emergent and collective macroscopic behaviors, emphasizing the need for a comprehensive and hierarchical approach to deeply investigate the complexities of biophysical complexes critical to all cellular functions.<sup>1</sup> Over the past decades, the integration of advanced computational methods, with the latest hardware and software, and the increasing availability of experimental molecular structure data, has deeply transformed molecular biology and drug discovery.<sup>2,3</sup> These enhancements, supported by significant advancements in theoretical frameworks, have evolved molecular simulations from simple proofs of concept to detailed *in silico* studies of protein folding<sup>4,5</sup> and dynamics.<sup>6,7</sup>

Molecular dynamics (MD) simulations, in particular, serve as a powerful tool for capturing the behaviors of proteins and biomolecules at the atomic level, providing fine temporal resolution. However, classical all-atom MD simulations present multiple limitations due to the substantial increase in computational resources required for larger systems and extended time scales; furthermore, postprocess and data analysis demands considerable effort, particularly in terms of human expertise and time.<sup>10,11</sup> To investigate larger systems over an extended time scale, a leading approach involves reducing computational demands via coarse-grained (CG) simulations, where molecular systems are simulated using fewer degrees of freedom than those associated with the atomic positions.<sup>12,13</sup>

Several CG models have been developed, each tailored to optimize molecular simulations and capture critical biophysical features. The MARTINI model,<sup>14,15</sup> for instance, excels due to its adaptability across a range of biomolecular systems, including membrane structure formation and protein interactions. Models such as AWSEM<sup>16</sup> and UNRES<sup>17</sup> have been

successful in simulating intramolecular protein dynamics, though they occasionally struggle to capture alternative metastable states. Similarly, Primo<sup>18,19</sup> and Rosetta<sup>20</sup> focus on specific molecular interactions and the design of protein structures and complexes, enhancing their accuracy in targeted applications. Currently, there is no universally accepted theory that precisely determines the most effective coarse-graining mapping for any given system, which in general is related to the intended application and computational constraints.<sup>21,22</sup> Once a coarse-graining mapping has been established, different strategies exist to define the model energy function, either to reproduce the reference fine-grained statistic (bottom-up)<sup>23</sup> or to match experimental observables (top-down).<sup>24</sup> Neural network potentials (NNPs) demonstrate remarkable efficiency in rapidly learning accurate potential energy<sup>25</sup> and effectively model many-body atomic interactions.<sup>26,27</sup> These features are extremely beneficial for developing coarse-grained (CG) force-fields that need multibody functional forms to accurately represent protein thermodynamics and incorporate implicit solvation effects. Many CG-NNP models<sup>28–32</sup> have been presented in the past years, but most rely on prior energies for stability and accuracy. In the context of CG-NNPs, the prior energy terms are defined as contributions to the final energy prediction of a model that are predefined (either constant or via some function of the atomic labels) or based on physical principles, independent of the machine-learning model. Here


**Received:** September 20, 2024

**Revised:** October 29, 2024

**Accepted:** October 30, 2024

**Published:** November 8, 2024





**Figure 1.** Pipeline for developing all-heavy atom NNPs is reported here. A CG map is applied to the mdCATH data set,<sup>8</sup> and an embedding  $z$  for each domain is created. TensorNet<sup>9</sup> is then trained using the CG data. Generalization and scale-up properties are evaluated on a set of four fast-folding proteins and larger domains in the final stage.

we present the first version of our Advanced Machine-learning Atomic Representation Omni-force-field (AMARO), a bottom-up CG-NNP without the need of prior terms to achieve stability and transferability (Figure 1).

## 2. MATERIAL AND METHODS

A CG model typically comprises two main components: selecting the CG resolution (or mapping) and designing an effective energy function once the mapping has been established. Mapping schemes often draw from physical or chemical intuition. CG sites may represent functional groups, residues, or monomers, or be tailored to a specific resolution. In this study, we adopt a no-hydrogen coarse-graining map, and a machine learning model is employed to learn the potential energy function.<sup>33</sup> Traditional force fields have historically treated hydrogens primarily as charge carriers due to their negligible Lennard-Jones interactions with heavier atoms. However, significant advances in implicit atom models, including the united atoms representation and implicit water machine learning potentials, illustrate the feasibility of accounting for these interactions using a many-body approach.<sup>30,34</sup>

**2.1. Neural Network Model.** In bottom-up CG approaches the interactions between CG beads are determined based on a more detailed model, and the many-body potential of mean force (PMF) is used to describe the free energy landscape of a system as a function of a collective coordinate or reaction coordinate, which is in other terms a configurational free energy in a reduced space. The term "many-body" in the context of a CG model refers to the fact that the potential energy landscape is considered in terms of interactions between groups of atoms rather than individual atoms, and network models offer a straightforward approximation in this context.<sup>13</sup> Our machine learning model of choice in this work is TensorNet,<sup>9</sup> a new neural network architecture that integrates  $O(3)$ -equivariance in message-passing and utilizes rank-2 Cartesian tensor representations.  $O(3)$ -equivariant NNPs<sup>35–37</sup> ensuring that tensor outputs transform correctly under rotations and reflections. In practice, in this work, we only predict the scalar energy, therefore invariance would have been enough. TensorNet, using a Cartesian representation, does not add a significant extra cost in incorporating these features, while it might provide higher expressive power and accuracy in energy and force prediction.<sup>38</sup>

**2.2. No Hydrogen CG Map.** The selection of an optimal mapping approach for transitioning from a fine- to coarse-grained representation is a critical aspect of a CG model definition. An effective CG map should significantly reduce the computational burden of the all-atom model while preserving sufficient information to prevent an excessively flat energy

surface. In this study, a no-hydrogen (*noh*) and no-water coarse-grained map has been chosen to reduce the degrees of freedom by almost half compared to the all-atom counterpart and to align with the basic force aggregation method.<sup>39</sup>

Consider a data set  $\mathcal{D}$  consisting of coordinate-force pairs obtained using an all-atom MD force field. The data set has  $M$  systems, each with a potentially different number of atoms. For each system, the coordinates are denoted by  $\mathbf{r} \in \mathbb{R}^{3N}$  and the forces by  $\mathbf{f} \in \mathbb{R}^{3N}$ , where  $N$  is the number of atoms in that particular system. A linear map operator  $\Xi: \mathbb{R}^{3N} \rightarrow \mathbb{R}^{3n}$  is defined to map an all-atom conformation  $\mathbf{r}$  to a coarse-grained conformation  $\mathbf{R}$ , where  $n$  denotes the number of non-hydrogen atoms. This mapping is applied separately to each system, accommodating their atomic compositions. A paired coarse-grained force  $\mathbf{F}$  is then considered for any conformation  $\mathbf{R}$ , and the force of the  $i$ -th *noh-bead*  $\mathbf{F}_i$  is defined as

$$\mathbf{F}_i = \mathbf{f}_{ih} + \sum_{j \in \mathcal{B}} \mathbf{f}_j \quad (1)$$

where  $\mathbf{f}_{ih}$  represents the force of the  $i$ -th heavy atom, and  $\mathcal{B}$  denotes the set of hydrogen atoms bonded to the  $i$ -th heavy atom in the fine-grained representation. Accounting for the *noh* coarse-graining map, this approach uniquely embeds each *noh-bead* by considering both its corresponding heavy atom and the number of bonded hydrogen atoms. A final set of 12 embedding values, as outlined in Table S1, was obtained enhancing the model's ability to differentiate between various electronic hybridizations. This selection empowers the model to more effectively learn the atoms' properties and molecular geometry.

**2.3. Neural Network Training.** As demonstrated in prior research, the acquisition of the many-body PMF involves minimizing the mean-squared deviation between a CG candidate force field and atomistic forces appropriately mapped. This method, known as variational force matching,<sup>40</sup> establishes a robust approach to effectively learning the PMF, laying the foundation for the exploration of intricate biomolecular interactions.<sup>33,41</sup> At CG resolution the force matching method becomes more complicated since the training data contain less information than their atomistic counterparts: energies are not available, and forces are noisy.<sup>31</sup> To enhance the quality of input data, we employed the *basic force aggregation* method in data set preparation. Coupled with the *noh*-CG map, this approach significantly improves PMF learning.<sup>39</sup> In essence, the model now learns the force of a *noh-bead* as the sum of the forces acting on its heavy atoms and the constrained (i.e., 'bonded') hydrogen atoms. Notably, this model introduces a groundbreaking feature in the CG-NNP field moving from delta-learning to directly learning the forces

acting on particles, no prior terms are considered. So, the energy function is parametrized by the network parameters  $\theta$  which are optimized to minimize the mean-squared deviation between predicted and labeled forces via the loss function

$$L_{FM}(\theta) = \frac{1}{\sum_{k=1}^K N_k} \sum_{k=1}^K \sum_{i=1}^{N_k} \|\nabla_{\mathbf{R}_i^k} \tilde{U}(\mathbf{R}^k; \theta) - \mathbf{F}_i^k\|^2 \quad (2)$$

where  $N_k$  is the number of beads in conformation  $k$ , and  $K$  is the total number of conformations in a batch. Predicted forces are obtained as the negative gradient of the potential energy  $\tilde{U}$  with respect to the *noh*-bead coordinates  $\mathbf{R}$ , and  $\mathbf{F}$  represents the labeled coarse-grained forces.

**2.4. Data Set.** The mdCATH data set<sup>8</sup> was the basis for applying the coarse-grained mapping approach. Specifically, the initial data were processed to retain only the heavy atoms' coordinates and forces, with a basic force aggregation map applied to the latter. Additionally, the *z* data set within each HDF5 file was modified to serve as the embedding for each system. A series of filters were implemented to exclude certain domains based on the following criteria: 1) domains containing more than 150 residues, 2) domains comprising more than 1000 *noh*-atoms, or 3) domains with less than 50% combined helix and sheet fractions. Various temperatures (320 K, 348 K, 379 K, 413 K, 450 K) were considered to facilitate the model's learning of atomic proximity or separation dynamics. Since we have trained the model using multiple temperatures, we are not expected to reproduce the energetics at 350 K exactly. Nonetheless, we find the agreement reasonable. The reason for training at multiple temperatures is that it increases the variability of the training data. In total, 2,834 domains and more than 26 million conformations were selected from the mdCATH data set after considering the filters and temperatures as described above.

The final data set used to train the model was obtained by applying a stride of 25 to the 26 M residual conformations, resulting in an approximate split of 900,000 conformations for training, 50,000 for validation, and 100,000 for testing.

The model's scalability was then tested on larger mdCATH domains, specifically those with more than 150 residues and a combined helix and sheet content greater than 50%.

At the same time, transferability was assessed using four fast-folding proteins: Chignolin, Trp-cage, Villin, and  $\alpha$ 3D; sequence similarity details are reported in Table S2.

**2.5. AMARO Molecular Simulations.** All MD simulations reported in this work, involving AMARO, were conducted with the mass of each CG-bead calculated as the combined mass of the heavy atoms and hydrogen atoms that constitute the bead.


**2.6. Markov State Models.** In this study, we analyze the dynamics of CG simulations using Markov State Models (MSMs)<sup>42–44</sup> available in HTMD<sup>45</sup> and compare them with those from corresponding all-atom simulations. MSMs partition the entire dynamics of a system into  $n$  discrete states and are particularly suited for systems that exhibit Markovian behavior—where future states depend only on the current state without memory of the past. We construct a transition probability matrix for these Markovian systems, characterized by  $n$  states and a lag time  $\tau$ , which records the system's state. This matrix enables us to determine state populations and conditional pairwise transition probabilities, from which free energies are derived. We employ time-lagged independent component analysis (TICA)<sup>46</sup> to enhance our analysis, reducing the high-dimensional conformational space into a

lower-dimensional, optimally reduced space. We then discretize this space using K-means clustering to construct the MSM. For each fast-folding protein, all-atom simulation data, characterized by pairwise  $C\alpha$  distances, are projected onto the first four components using TICA. A reference free energy surface was then constructed for each system by binning the first two TICA dimensions into an  $80 \times 80$  grid and averaging the weights of the equilibrium probability in each bin, as computed by the Markov state model. For the CG simulations, we adopt the approach outlined in,<sup>47</sup> utilizing covariance matrices from all-atom molecular dynamics (MD) to project the first three components. This method aligns with established methodologies, ensuring consistency and facilitating further analysis. HTMD provided the necessary computational tools and framework to perform these analyses.<sup>45</sup> Finally, to avoid biasing the model with starting conformations, 10% of the initial frames of each trajectory were removed from the analysis except for  $\alpha$ 3D, where 5% of the initial frames were discarded due to the longer simulation times considered.

### 3. RESULTS

TensorNet was trained on the filtered mdCATH data set, as described in Section 2.4, using TorchMD-Net<sup>48</sup> for 100 epochs. Detailed information on model architecture and training hyper-parameters can be found in the Supporting Information (see Tables S3 and S4).

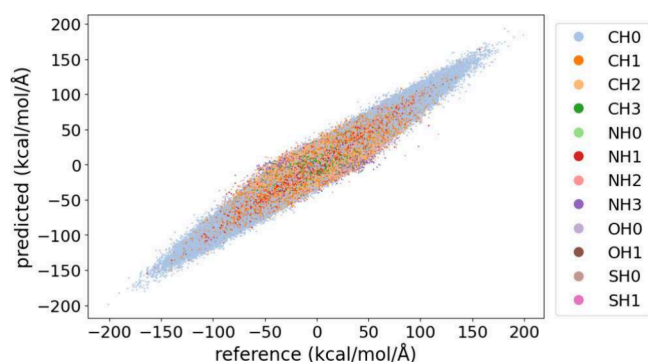
The final L1 test loss for the model is reported as 5.07 kcal/mol/Å, while MSE loss for training and validation are reported in Figure 2.



**Figure 2.** Training- and validation- MSE loss, in blue and orange respectively, for AMARO as a function of training epoch.

**3.1. Generalization to Larger Domains.** We explore the ability of the AMARO to scale up when larger protein domains than those considered in the training set. To achieve this, we selected a subset of 5,000 conformations from the mdCATH data set, specifically targeting domains with between 150 and 250 residues and a combined helix and sheet fraction  $>50\%$ . This selection criteria ensured that the domains were representative of complex protein structures while still maintaining a manageable size for computational analysis. The forces acting on the *noh*-bead within these larger domains were evaluated using the CG-NNP. AMARO exhibited a mean absolute error (MAE) of 4.98 kcal/mol/Å, assessed for each force component ( $x$ ,  $y$ ,  $z$ ). The error value here recorded, compared also to the one obtained at the end of the training,

proves that the learned potential can scale up without loss in accuracy. Moreover, Figure 3 presents a direct comparison



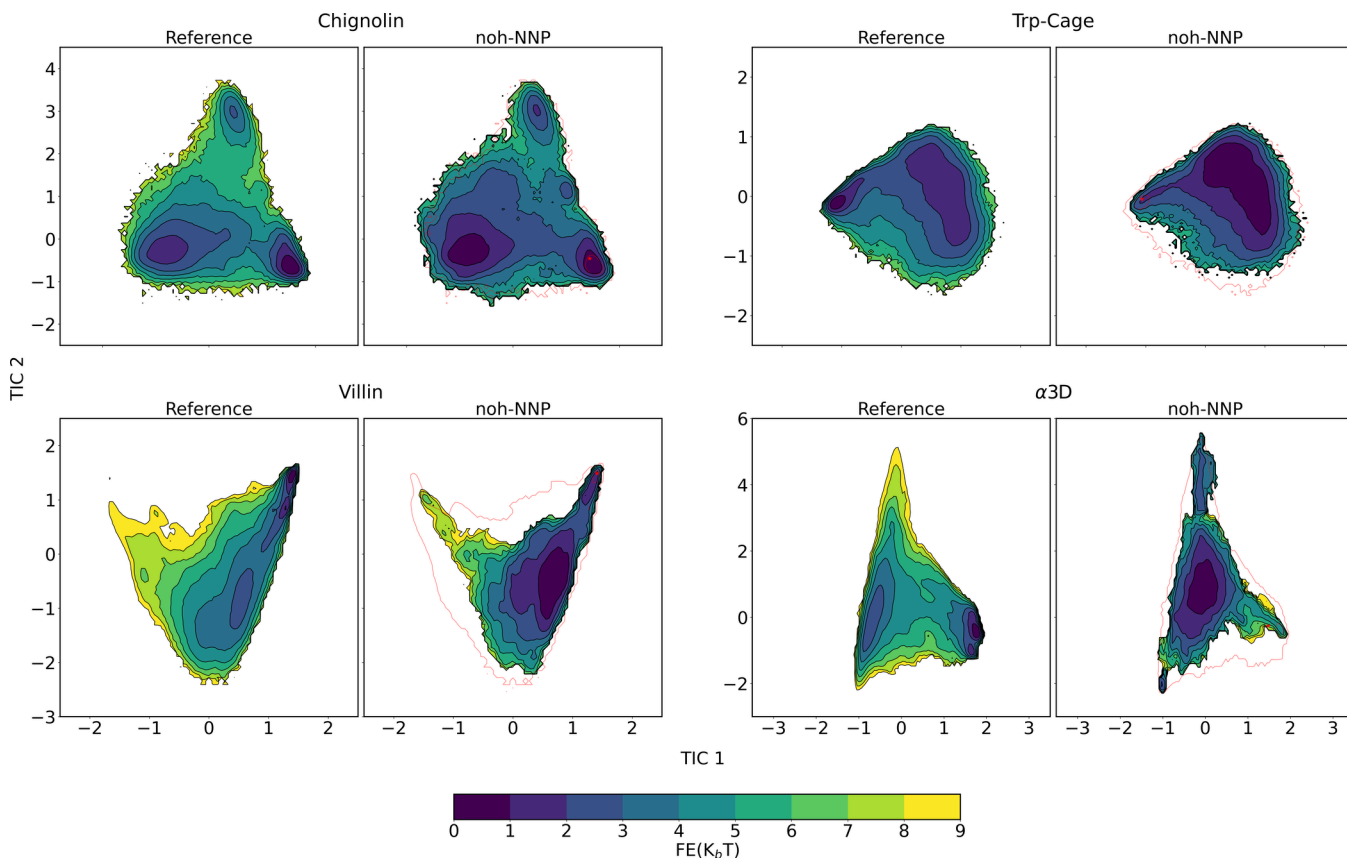
**Figure 3.** Scale-up validation of AMARO on larger domains from mdCATH. Comparison between labeled (i.e., reference) and predicted force component values ( $x$ ,  $y$ ,  $z$ ). Each data point in the scatter plot is color-coded according to the CG atom type.

between the expected and observed force values, with each dot color-coded by CG atom type, illustrating the model's precision. The results underscore the robustness of the CG-NNP model and its potential applicability to larger biological systems, confirming its feature to maintain performance across an expanded range of system sizes and complexities.

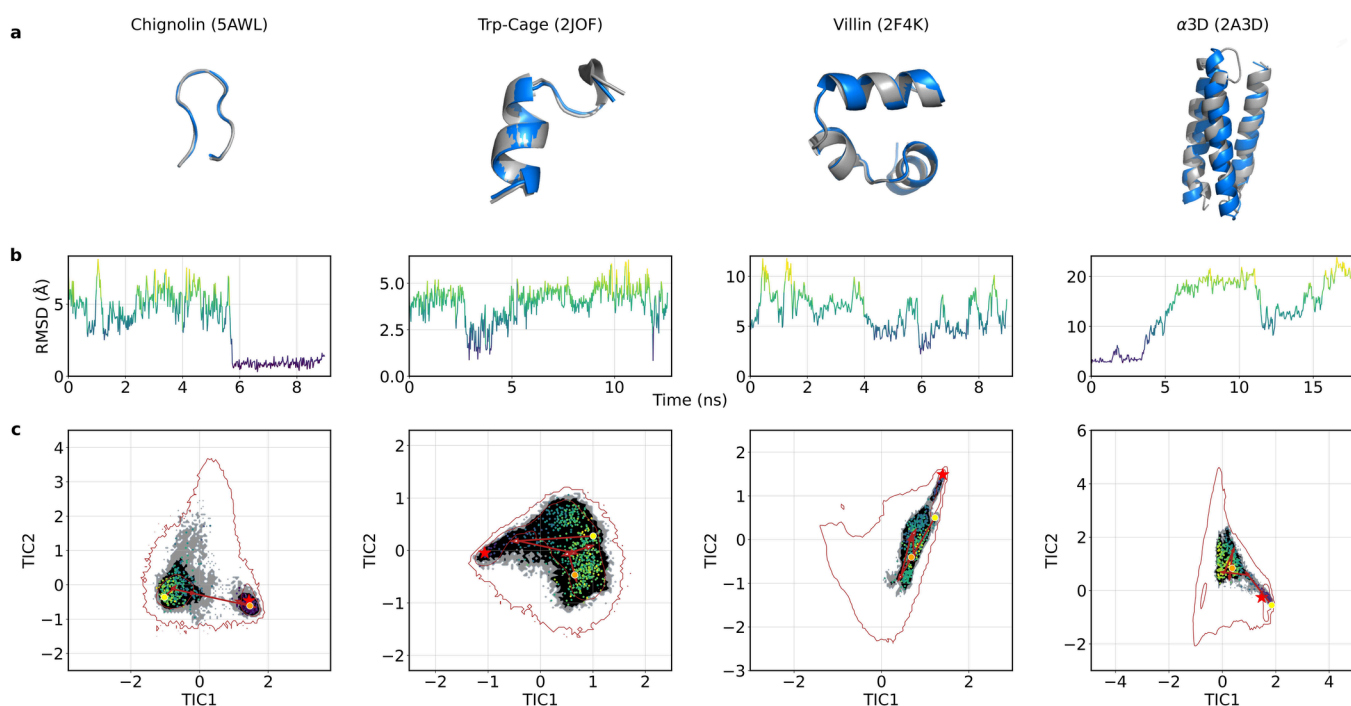
Detailed errors per atom type are reported in Table S5, with  $\text{NH}_3^+$  exhibiting the highest error, a mean absolute error (MAE) of 8.10 kcal/mol/Å. However, this is not due to  $\text{NH}_3$

bead-type being underrepresented in the training data sets, as shown in Figure S1, but rather due to the physical and chemical properties of this group. Five domains with the highest number of  $\text{NH}_3^+$  groups (i.e.,  $\geq 20$ ) were selected for further investigation: 1kvnA00, 1nu7D01, 1w9rA00, 2c5zA00, 2jzvA00, 2nc9A00 and 3qneA01. In all of these cases, the protonated amino group corresponds to the terminal amino group along the lysine side chain. These groups are oriented outward, suggesting solvent interactions, which are not accounted for in the current modeling. Moreover, in more compact structures, these groups might be involved in salt bridges with carboxyl groups. Such complex interactions contribute to the challenge of generalizing the model's parameters for this specific atom type.

**3.2. Validation on Fast-Folding Proteins.** To assess the model's transferability, we selected four fast-folding proteins not included in the training set: Chignolin (175 atoms), Trp-Cage (210 atoms), Villin (573 atoms), and  $\alpha$ 3D (1149 atoms); the number of atoms in their all-atom representations are indicated in parentheses. After applying the *noh* mapping, the sizes were reduced to 97, 112, 286, and 576 atoms, respectively. TorchMD<sup>49</sup> was used to run 32 replica MD simulations for each protein and considering at least 320 ns of aggregated simulation time, see Table S6. The initial coordinates for these simulations were uniformly sampled from the respective TICA surfaces of each fast folder, as shown in Figure S2. For Trp-Cage, only residues from 2 to 16 have been considered to focus on capturing the overall folding and not the cis–trans proline (residue 20) isomerization.



**Figure 4.** Comparative analysis of the free energy landscape obtained from all-atom simulations (left) and NNP coarse-grained simulations (right) across the first two TICA dimensions for four fast-folding proteins: Chignolin, Trp-cage, Villin and  $\alpha$ 3D.



**Figure 5.** CG trajectories of Chignolin, Trp-Cage, Villin, and  $\alpha$ 3D, selected based on the inclusion of microstates from the lowest RMSD macrostate. (a) Minimum RMSD conformation (blue) aligned with the experimental structure (gray) for each protein, labeled with the protein name and PDB ID. (b)  $C\alpha$  RMSD of each trajectory compared to the crystal structure. (c) CG free energy surface, projected over the first two TICs with the folded state (red star) and sampled states indicated by RMSD color-coded dots. The trajectory's progression is illustrated with arrows connecting the starting (yellow point) and ending (orange point) conformations. The all-atom equilibrium density is shown by a red contour.

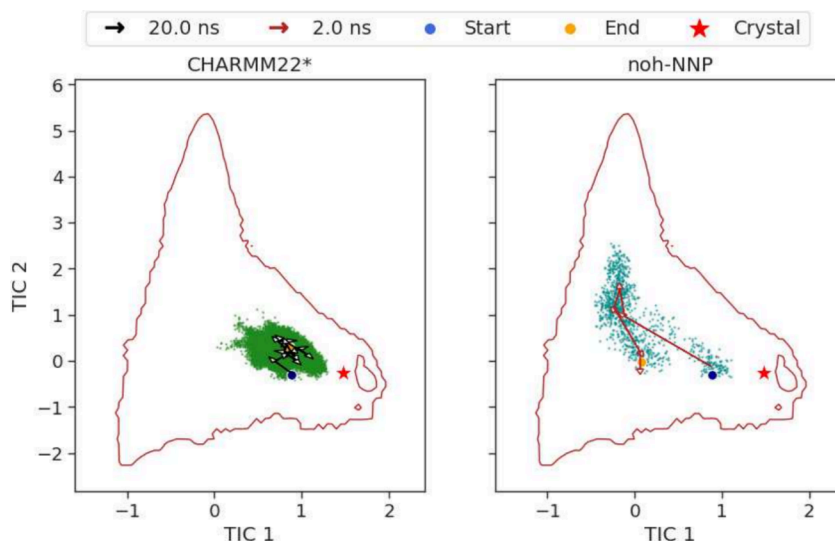
**3.3. Recovering the Energetic Landscape.** For three of the four fast-folding proteins the TICA landscape has been successfully recovered, see Figure 4. In contrast, for  $\alpha$ 3D, the recovery is only partial, with most microstates populating a middle region between the global minimum on the right and a local minimum on the left. This behavior could be attributed to the particular secondary structure of  $\alpha$ 3D, which is characterized by a high proportion of  $\alpha$ -helices. Additionally, the relative shape anisotropy (RSA) of the system, 0.3, may not be well-represented in the training data set, which has an average RSA of  $0.17 \pm 0.16$ . This discrepancy, coupled with the larger system size and the presence of  $\text{NH}_3^+$  groups in the unstructured region between the first and second helices of the folded state, could contribute to the observed differences.

While Figure 4 displays the two TICA dimensions as the principal axes and uses free energy as the third dimension, Figure S3 presents the free energy on the  $y$ -axis, providing a quantitative analysis. For the simplest target, Chignolin, the absolute minimum is perfectly captured by AMARO. For more complex structures like Trp-Cage and Villin, the overall shape of the profile is well approximated. Conversely, for  $\alpha$ 3D, extensive sampling in the central region results in a shift, leading to the identification of an absolute minimum between the two minima.

**3.4. Sampling the Native Structures of Unseen Training Proteins.** Analysis of the CG simulations using MSMs, detailed in Section 2.6 and Table S7, revealed that the model successfully reproduced the experimental structure of the corresponding fast-folding proteins, as illustrated in Figure 5a. Sampling originated from the native macrostate, defined as the macrostate containing the conformation with the minimum root-mean-square deviation (RMSD) with respect to the

experimental crystal structure. The models accurately predicted secondary and tertiary structural elements, with loops and unstructured terminal regions showing minimal variation, except for  $\alpha$ 3D. Table S8 provides detailed information on the equilibrium probability, mean, and minimum RMSDs of the native conformations, as reflected by the high equilibrium probabilities for these macrostates. For Chignolin, Trp-Cage, and Villin, an average RMSD below 1 Å was observed when comparing the native macrostate to the folded structure, more in detail: 0.15 Å, 0.30 Å, and 0.6 Å, respectively. In contrast, the larger and more complex  $\alpha$ 3D system reported an RMSD of 2.30 Å. These results underscore the model's accuracy in structural prediction and its adaptability across different molecular systems, though further investigation is required to fully understand the particular case of  $\alpha$ 3D.

**3.5. Computational Efficiency.** To assess the computational efficiency of AMARO, we conducted a comparative analysis against traditional all-atom simulations, focusing on sampling free energy (FE) landscapes within the TICA space. Both simulations were constrained to a 12-h time frame to ensure a fair comparison, with the all-atom simulations using the CHARMM22\* force field,<sup>50,51</sup> and explicit solvent, while the AMARO simulation focused on heavy atoms using a *noh* mapping approach. Both models were run on openMM<sup>52</sup> using an NVIDIA RTX 4090, and employed a Langevin thermostat at 350 K, with differing friction coefficients:  $0.1 \text{ ps}^{-1}$  for CHARMM22\* and  $1 \text{ ps}^{-1}$  for AMARO (standard setup, see Section 2.5). Additionally, the hydrogen mass repartitioning (HMR) scheme<sup>53</sup> was set at 4 a.m.u. for the classical force field, and *increased mass* was applied for the NNP system as



**Figure 6.** Sampling efficiency comparison in the TICA landscape for the  $\alpha$ 3D system under two force fields: CHARMM22\* (left) and AMARO (right). Both simulations, starting from a randomly selected initial conformation (blue dot) and ending at a final conformation (orange dot), were run within a fixed 12-h time window for direct comparison.

detailed in the methods section. All-atom simulations employed long-range electrostatics using the particle-mesh Ewald (PME) summation method<sup>54,55</sup> with a cutoff of 9 Å, while van der Waals interactions used a cutoff of 9 Å and a switching distance of 7.5 Å. Hydrogen atoms were constrained using the SHAKE<sup>56</sup> algorithm. In contrast, AMARO does not use explicit treatment for long-range interactions, as they are modeled by the NNP with a total receptive field of 10 Å.

The analysis revealed that the all-atom simulations achieved 265.5 ns of molecular simulation while AMARO completed 8.4 ns within the same operational time window. However, to account for differences in temporal resolution between the coarse-grained and all-atom models, we compared the areas of the TICA landscapes recovered by both models, see Figure 6. The starting point (blue dot), the same for both cases, was selected randomly and resulted in a conformation near the global minimum, while the red star represents the crystal structure (PDB ID: 2A3D).

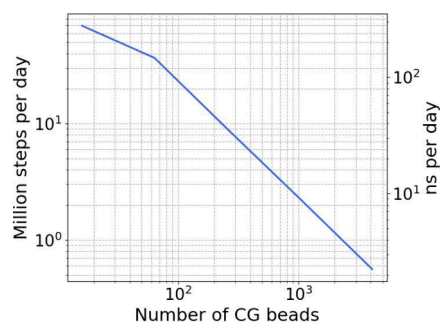
The all-atom simulation (left panel) explores a more localized region of the TICA space, with the trajectory (black arrows) predominantly staying within a single basin. This indicates limited sampling within the time frame due to the high dimensionality and energy barriers typical of fully atomistic representations. The final conformation is marked by a yellow dot, indicating that the simulation does not move significantly away from the initial state.

In contrast, the CG-NNP simulation (right panel), despite the shorter trajectory length, explores a broader region of the conformational space (red arrows), as seen by the wider spread of sampled conformations. The trajectory covers a larger portion of the landscape, crossing energy barriers more easily and reaching areas of higher free energy. This behavior, characteristic of coarse-grained models, is due to the reduced degrees of freedom and the smoothing of energy landscapes which produces an effective faster kinetics. Despite this broader sampling, the path still shows consistency with the overall free energy landscape, see Figure 4 for reference, suggesting thermodynamic consistency.

Notably, while the all-atom simulation reaches a minimum RMSD of 5.03 Å with respect to the crystal structure, the CG-

NNP model achieves a lower minimum RMSD of 4.38 Å. Moreover, the average RMSD values of  $13.31 \pm 4.82$  Å for CG-NNP and  $6.41 \pm 0.5$  Å for CHARMM22\* suggest that the coarse-grained model explores a larger conformational space, as expected.

The CG-NNP's efficiency was further quantified relative to system size, measured in CG beads (Figure 7, providing a



**Figure 7.** Log–log plot of AMARO's computational speed relative to system size (number of CG Beads). The left y-axis shows the number of simulation steps (in millions) that can be executed within a single day. The right y-axis displays the performance in nanoseconds per day (ns/day), assuming a time-step of 4 fs.

quantitative benchmark of AMARO's performance. Both millions of simulation steps (left y-axis) and ns/day (right y-axis) are considered variables. The term “simulation step” refers to a forward and backward step of the model. The efficiency demonstrated by the CG-NNP in small-to-medium systems (up to 1000 CG beads) makes it a valuable tool for exploring larger conformational landscapes or long-time scale events.

#### 4. CONCLUSIONS

This paper introduces the first version of AMARO a new fully machine-learning coarse-grained force field offering a new framework for molecular dynamics simulations. AMARO uses an all-heavy-atoms coarse-graining strategy paired with variational force matching, which simplifies protein representation

while retaining essential dynamical information. This approach addresses previous challenges in balancing model simplicity with the retention of critical dynamics, which often require energy priors for stabilization. Notably, our model demonstrates remarkable transferability and scaling-up ability. However, further enhancements in computational efficiency and memory usage must be achieved in the future. The current study serves as a proof of concept, challenging the reliance on prior energies in developing stable NNPs for studying protein thermodynamics.

## ■ ASSOCIATED CONTENT

### Data Availability Statement

Code and relevant data to reproduce this work are available at <https://github.com/compsciencelab/amaro>.

### SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.4c01239>.

Data and hyperparameters used in our computational experiments to ensure transparency and reproducibility, including specific parameters for training, model, and fast-folding protein MSMs, details of the values used to achieve the main text results, and additional analyses and statistics (PDF)

## ■ AUTHOR INFORMATION

### Corresponding Author

Gianni De Fabritiis – Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona 08003, Spain; Acellera Labs, Barcelona 08005, Spain; Institutio Catalana de Recerca i Estudis Avançats (ICREA), Barcelona 08010, Spain; [orcid.org/0000-0003-3913-4877](https://orcid.org/0000-0003-3913-4877); Email: [g.defabritiis@gmail.com](mailto:g.defabritiis@gmail.com)

### Authors

Antonio Mirarchi – Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona 08003, Spain; [orcid.org/0000-0001-7798-0519](https://orcid.org/0000-0001-7798-0519)  
Raúl P. Peláez – Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona 08003, Spain; [orcid.org/0000-0003-3393-7329](https://orcid.org/0000-0003-3393-7329)  
Guillem Simeon – Computational Science Laboratory, Universitat Pompeu Fabra, Barcelona 08003, Spain; [orcid.org/0000-0003-3225-1632](https://orcid.org/0000-0003-3225-1632)

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.4c01239>

### Notes

The authors declare no competing financial interest.

## ■ ACKNOWLEDGMENTS

AM is financially supported by Generalitat de Catalunya's Agency for Management of University and Research Grants (AGAUR) PhD grant 2024 FI-1-00278; the project PID2023-151620OB-I00 has been funded by MCIN/AEI/10.13039/501100011033. Research reported in this publication was partially supported by the National Institute of General Medical Sciences (NIGMS) of the National Institutes of Health under award number R01GM140090. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

## ■ REFERENCES

- (1) Williamson, J. R. Cooperativity in macromolecular assembly. *Nat. Chem. Biol.* **2008**, *4*, 458–465.
- (2) McCammon, J. A.; Gelin, B. R.; Karplus, M. Dynamics of folded proteins. *nature* **1977**, *267*, 585–590.
- (3) Hollingsworth, S. A.; Dror, R. O. Molecular dynamics simulation for all. *Neuron* **2018**, *99*, 1129–1143.
- (4) Baxa, M. C.; Haddadian, E. J.; Jumper, J. M.; Freed, K. F.; Sosnick, T. R. Loss of conformational entropy in protein folding calculated using realistic ensembles and its implications for NMR-based calculations. *Proc. Natl. Acad. Sci. U. S. A.* **2014**, *111*, 15396–15401.
- (5) Noé, F.; Schütte, C.; Vanden-Eijnden, E.; Reich, L.; Weikl, T. R. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. U. S. A.* **2009**, *106*, 19011–19016.
- (6) Muller, M. P.; Jiang, T.; Sun, C.; Lihan, M.; Pant, S.; Mahinthichaichan, P.; Trifan, A.; Tajkhorshid, E. Characterization of lipid-protein interactions and lipid-mediated modulation of membrane protein function through molecular simulation. *Chem. Rev.* **2019**, *119*, 6086–6161.
- (7) McGeagh, J. D.; Ranaghan, K. E.; Mulholland, A. J. Protein dynamics and enzyme catalysis: insights from simulations. *Biochimica et Biophysica Acta (BBA)-Proteins and Proteomics* **2011**, *1814*, 1077–1092.
- (8) Mirarchi, A.; Giorgino, T.; Fabritiis, G. D. mdCATH: A Large-Scale MD Dataset for Data-Driven Computational Biophysics. *arXiv*, 2407.14794, 2024.
- (9) Simeon, G.; de Fabritiis, G. TensorNet Cartesian Tensor Representations for Efficient Learning of Molecular Potentials. *NIPS '23: Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023, Article No.: 1623, Pages 37334–37353.
- (10) Chodera, J. D.; Noé, F. Markov state models of biomolecular conformational dynamics. *Curr. Opin. Struct. Biol.* **2014**, *25*, 135–144.
- (11) Plattner, N.; Doerr, S.; De Fabritiis, G.; Noé, F. Complete protein-protein association kinetics in atomic detail revealed by molecular dynamics simulations and Markov modelling. *Nature Chem.* **2017**, *9*, 1005–1011.
- (12) Pak, A. J.; Voth, G. A. Advances in coarse-grained modeling of macromolecular complexes. *Curr. Opin. Struct. Biol.* **2018**, *52*, 119–126.
- (13) Noid, W. G. Perspective Coarse-grained models for biomolecular systems. *J. Chem. Phys.* **2013**, *139*, 090901.
- (14) 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.
- (15) Monticelli, L.; Kandasamy, S. K.; Periole, X.; Larson, R. G.; Tieleman, D. P.; Marrink, S.-J. The MARTINI coarse-grained force field: extension to proteins. *J. Chem. Theory Comput.* **2008**, *4*, 819–834.
- (16) Davtyan, A.; Schafer, N. P.; Zheng, W.; Clementi, C.; Wolynes, P. G.; Papoian, G. A. AWSEM-MD: Protein structure prediction using coarse-grained physical potentials and bioinformatically based local structure biasing. *J. Phys. Chem. B* **2012**, *116*, 8494–8503.
- (17) Liwo, A.; Sieradzan, A. K.; Lipska, A. G.; Czaplowski, C.; Joung, I.; Zmudzńska, W.; Halabis, A.; Oldziej, S. A general method for the derivation of the functional forms of the effective energy terms in coarse-grained energy functions of polymers. III. Determination of scale-consistent backbone-local and correlation potentials in the UNRES force field and force-field calibration and validation. *J. Chem. Phys.* **2019**, *150*, 155104.
- (18) Gopal, S. M.; Mukherjee, S.; Cheng, Y.-M.; Feig, M. PRIMO/PRIMONA: a coarse-grained model for proteins and nucleic acids that preserves near-atomistic accuracy. *Proteins: Struct., Funct., Bioinf.* **2010**, *78*, 1266–1281.
- (19) Kar, P.; Gopal, S. M.; Cheng, Y.-M.; Predeus, A.; Feig, M. PRIMO: a transferable coarse-grained force field for proteins. *J. Chem. Theory Comput.* **2013**, *9*, 3769–3788.

- (20) Das, R.; Baker, D. Macromolecular modeling with rosetta. *Annu. Rev. Biochem.* **2008**, *77*, 363–382.
- (21) Foley, T. T.; Kidder, K. M.; Shell, M. S.; Noid, W. Exploring the landscape of model representations. *Proc. Natl. Acad. Sci. U. S. A.* **2020**, *117*, 24061–24068.
- (22) Boninsegna, L.; Banisch, R.; Clementi, C. A data-driven perspective on the hierarchical assembly of molecular structures. *J. Chem. Theory Comput.* **2018**, *14*, 453–460.
- (23) 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.
- (24) Navarro, C.; Majewski, M.; De Fabritiis, G. Top-down machine learning of coarse-grained protein force-fields. *arXiv*, 2306.11375, 2023.
- (25) Noé, F.; Tkatchenko, A.; Müller, K.-R.; Clementi, C. Machine learning for molecular simulation. *Annu. Rev. Phys. Chem.* **2020**, *71*, 361–390.
- (26) Behler, J.; Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical review letters* **2007**, *98*, No. 146401.
- (27) 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*, 164113.
- (28) 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 central science* **2019**, *5*, 755–767.
- (29) Husic, B. E.; Charron, N. E.; Lemm, D.; Wang, J.; Pérez, A.; Majewski, M.; Krämer, A.; Chen, Y.; Olsson, S.; de Fabritiis, G. Coarse graining molecular dynamics with graph neural networks. *J. Chem. Phys.* **2020**, *153*, 194101.
- (30) Chen, Y.; Krämer, A.; Charron, N. E.; Husic, B. E.; Clementi, C.; Noé, F. Machine learning implicit solvation for molecular dynamics. *J. Chem. Phys.* **2021**, *155*, 084101.
- (31) Durumeric, A. E.; Charron, N. E.; Templeton, C.; Musil, F.; Bonneau, K.; Pasos-Trejo, A. S.; Chen, Y.; Kelkar, A.; Noé, F.; Clementi, C. Machine learned coarse-grained protein force-fields: Are we there yet? *Curr. Opin. Struct. Biol.* **2023**, *79*, No. 102533.
- (32) Charron, N. E.; Musil, F.; Guljas, A.; Chen, Y.; Bonneau, K.; Pasos-Trejo, A. S.; Venturin, J.; Gusew, D.; Zaporozhets, I.; Krämer, A. et al. Navigating protein landscapes with a machine-learned transferable coarse-grained model. *arXiv*, 2310.18278, 2023.
- (33) Kmiecik, S.; Gront, D.; Kolinski, M.; Wieteska, L.; Dawid, A. E.; Kolinski, A. Coarse-grained protein models and their applications. *Chem. Rev.* **2016**, *116*, 7898–7936.
- (34) Yang, L.; Tan, C.-h.; Hsieh, M.-J.; Wang, J.; Duan, Y.; Cieplak, P.; Caldwell, J.; Kollman, P. A.; Luo, R. New-generation amber united-atom force field. *J. Phys. Chem. B* **2006**, *110*, 13166–13176.
- (35) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **2022**, *13*, 2453.
- (36) Batatia, I.; Kovacs, D. P.; Simm, G.; Ortner, C.; Csányi, G. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems* **2022**, *35*, 11423–11436.
- (37) Musaelian, A.; Batzner, S.; Johansson, A.; Sun, L.; Owen, C. J.; Kornbluth, M.; Kozinsky, B. Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **2023**, *14*, 579.
- (38) Duval, A.; Mathis, S. V.; Joshi, C. K.; Schmidt, V.; Miret, S.; Malliaros, F. D.; Cohen, T.; Lio, P.; Bengio, Y.; Bronstein, M. A Hitchhiker's Guide to Geometric GNNs for 3D Atomic Systems. *arXiv*, 2312.07511, 2023.
- (39) Krämer, A.; Durumeric, A. E.; 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.
- (40) Izvekov, S.; Voth, G. A. A multiscale coarse-graining method for biomolecular systems. *J. Phys. Chem. B* **2005**, *109*, 2469–2473.
- (41) 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*, 244114.
- (42) Noé, F.; Fischer, S. Transition networks for modeling the kinetics of conformational change in macromolecules. *Curr. Opin. Struct. Biol.* **2008**, *18*, 154–162.
- (43) Husic, B. E.; Pande, V. S. Markov state models: From an art to a science. *J. Am. Chem. Soc.* **2018**, *140*, 2386–2396.
- (44) Pan, A. C.; Roux, B. Building Markov state models along pathways to determine free energies and rates of transitions. *J. Chem. Phys.* **2008**, *129*, 064107.
- (45) Doerr, S.; Harvey, M.; Noé, F.; De Fabritiis, G. HTMD: high-throughput molecular dynamics for molecular discovery. *J. Chem. Theory Comput.* **2016**, *12*, 1845–1852.
- (46) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. *J. Chem. Phys.* **2013**, *139*, 015102.
- (47) Majewski, M.; Pérez, A.; Thölke, P.; Doerr, S.; Charron, N. E.; Giorgino, T.; Husic, B. E.; Clementi, C.; Noé, F.; De Fabritiis, G. Machine learning coarse-grained potentials of protein thermodynamics. *Nat. Commun.* **2023**, *14*, 5739.
- (48) Pelaez, R. P.; Simeon, G.; Galvelis, R.; Mirarchi, A.; Eastman, P.; Doerr, S.; Thölke, P.; Markland, T. E.; De Fabritiis, G. TorchMD-Net 2.0: Fast Neural Network Potentials for Molecular Simulations. *J. Chem. Theory Comput.* **2024**, *20*, 4076–4087.
- (49) Doerr, S.; Majewski, M.; Pérez, A.; Krämer, A.; Clementi, C.; Noe, F.; Giorgino, T.; Fabritiis, G. D. TorchMD: A deep learning framework for molecular simulations. *J. Chem. Theory Comput.* **2021**, *17*, 2355.
- (50) MacKerell, A. D., Jr; Bashford, D.; Bellott, M.; Dunbrack, R. L., Jr; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; et al. All-atom empirical potential for molecular modeling and dynamics studies of proteins. *J. Phys. Chem. B* **1998**, *102*, 3586–3616.
- (51) Mackerell, A. D., Jr; Feig, M.; Brooks, C. L., III Extending the treatment of backbone energetics in protein force fields: Limitations of gas-phase quantum mechanics in reproducing protein conformational distributions in molecular dynamics simulations. *Journal of computational chemistry* **2004**, *25*, 1400–1415.
- (52) Eastman, P.; Galvelis, R.; Peláez, R. P.; Abreu, C. R.; Farr, S. E.; Gallicchio, E.; Gorenko, A.; Henry, M. M.; Hu, F.; Huang, J. OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials. *J. Phys. Chem. B* **2024**, *128*, 109.
- (53) Feenstra, K. A.; Hess, B.; Berendsen, H. J. C. Improving efficiency of large time-scale molecular dynamics simulations of hydrogen-rich systems. *J. Comput. Chem.* **1999**, *20*, 786–798.
- (54) Darden, T.; York, D.; Pedersen, L. Particle mesh Ewald: An N log(N) method for Ewald sums in large systems. *J. Chem. Phys.* **1993**, *98*, 10089–10092.
- (55) Harvey, M.; De Fabritiis, G. An implementation of the smooth particle mesh Ewald method on GPU hardware. *J. Chem. Theory Comput.* **2009**, *5*, 2371–2377.
- (56) Ryckaert, J.-P.; Ciccotti, G.; Berendsen, H. J. Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes. *J. Comput. Phys.* **1977**, *23*, 327–341.