

Challenges for machine learning force fields in reproducing potential energy surfaces of flexible molecules

Cite as: J. Chem. Phys. **154**, 094119 (2021); <https://doi.org/10.1063/5.0038516>

Submitted: 23 November 2020 • Accepted: 11 February 2021 • Published Online: 03 March 2021

 Valentin Vassilev-Galindo, Gregory Fonseca,  Igor Poltavsky, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[SchNet – A deep learning architecture for molecules and materials](#)

The Journal of Chemical Physics **148**, 241722 (2018); <https://doi.org/10.1063/1.5019779>

[Machine learning for interatomic potential models](#)

The Journal of Chemical Physics **152**, 050902 (2020); <https://doi.org/10.1063/1.5126336>

[When do short-range atomistic machine-learning models fall short?](#)

The Journal of Chemical Physics **154**, 034111 (2021); <https://doi.org/10.1063/5.0031215>

The Journal
of Chemical Physics

SPECIAL TOPIC: Low-Dimensional
Materials for Quantum Information Science

Submit Today!





Challenges for machine learning force fields in reproducing potential energy surfaces of flexible molecules

Cite as: J. Chem. Phys. 154, 094119 (2021); doi: 10.1063/5.0038516

Submitted: 23 November 2020 • Accepted: 11 February 2021 •

Published Online: 3 March 2021



Valentin Vassilev-Galindo,  Gregory Fonseca, Igor Poltavsky,  and Alexandre Tkatchenko^{a)}

AFFILIATIONS

Department of Physics and Materials Science, University of Luxembourg, L-1511 Luxembourg City, Luxembourg

^{a)} Author to whom correspondence should be addressed: alexandre.tkatchenko@uni.lu

ABSTRACT

Dynamics of flexible molecules are often determined by an interplay between local chemical bond fluctuations and conformational changes driven by long-range electrostatics and van der Waals interactions. This interplay between interactions yields complex potential-energy surfaces (PESs) with multiple minima and transition paths between them. In this work, we assess the performance of the state-of-the-art Machine Learning (ML) models, namely, sGDML, SchNet, Gaussian Approximation Potentials/Smooth Overlap of Atomic Positions (GAPs/SOAPs), and Behler–Parrinello neural networks, for reproducing such PESs, while using limited amounts of reference data. As a benchmark, we use the *cis* to *trans* thermal relaxation in an azobenzene molecule, where at least three different transition mechanisms should be considered. Although GAP/SOAP, SchNet, and sGDML models can globally achieve a chemical accuracy of 1 kcal mol^{−1} with fewer than 1000 training points, predictions greatly depend on the ML method used and on the local region of the PES being sampled. Within a given ML method, large differences can be found between predictions of close-to-equilibrium and transition regions, as well as for different transition mechanisms. We identify key challenges that the ML models face mainly due to the intrinsic limitations of commonly used atom-based descriptors. All in all, our results suggest switching from learning the entire PES within a single model to using multiple local models with optimized descriptors, training sets, and architectures for different parts of the complex PES.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0038516>

I. INTRODUCTION

Thermodynamic and dynamical properties of molecules can be computed if an accurate model for a potential-energy surface (PES) is provided. Among these properties, transition paths connecting pairs of minima on the PES are crucial for understanding the dynamics of complex systems,¹ such as conformational changes in molecules,^{2,3} nucleation events during phase transitions,^{4,5} and folding and unfolding of proteins.^{6–8} The state-of-the-art methods for finding transition pathways range from the optimization of a single direction on the PES^{9,10} or a chain of states connecting both minima, e.g., the string¹¹ and nudged elastic band (NEB)^{12,13} methods, to the more sophisticated transition path sampling techniques.^{14,15} Most of them often provide only a single “optimal” transition path. The rate of success in finding the path highly depends on the dimensionality and complexity of the PES: flexible

molecules containing a few tens of atoms, such as organic photoswitches and peptides, are already challenging to deal with. Moreover, due to the non-trivial interplay between covalent and non-covalent interactions, the transitions in such molecules may happen following several different pathways. In this case, one needs to consider the contribution of every path to the transition process, and the knowledge about just one optimal pathway is insufficient. Practical studies of such transitions require reliable force fields (FFs) able to accurately reconstruct broad regions of the PES, including multiple local minima and all the relevant pathways connecting them.

In recent years, the use of Machine Learning (ML) in chemistry and materials science has been a subject of intensive research.^{16–63} Namely, the advent of ML potentials has offered new tools for meeting the constantly increasing demand for accurate simulations of realistic systems since such potentials aim to keep the

accuracy of *ab initio* calculations with an efficiency closer to that of classical force fields. Among the available methodologies, neural networks (NNs)^{33–41} and kernel-based methods^{16–30} are the most used to learn the PES of molecules. However, this learning task is not easy, and it has encouraged the improvement of data sampling,^{29–33} molecular representations,^{23–27,44–55} and NN architectures.^{37–41,56,57} In addition, some efforts have been directed toward improving ML-aided search and sampling of transition states and pathways.^{32,64,65} For instance, Noé *et al.*³² showed a promising method to sample rare events between equilibrium states using Boltzmann generators. The method is by many orders of magnitude more efficient than “brute force” molecular dynamics (MD) simulations. Other approaches^{64,65} are built on the state-of-the-art methods for calculating transition states enhanced with ML techniques. ML-enhanced transition state search methods are more efficient than their precursors but present the same limitations. ML methods are often data demanding, making their application infeasible when computationally expensive *ab initio* methods are required. Hence, constructing robust ML models for flexible molecules is the necessary next step for practical applications of ML potentials in chemistry and biology.

There are two main challenges in building accurate ML models for flexible molecules: First, generating enough data around the transition regions of the PES. Second, building a highly accurate and data-efficient ML model that describes the resulting complex PES. In this work, we address both of these challenges in an azobenzene ($C_{12}H_{10}N_2$) molecule, for example. While being small in size, azobenzene is flexible enough to feature a *cis* to *trans* thermal relaxation following at least three possible channels: a rotation, an inversion, and a rotation assisted by inversion mechanisms.^{66–70} We start by discussing the problem of building reliable reference datasets for these transitions. Then, we assess the performance of the state-of-the-art ML methods in the prediction of forces and energies along the obtained transition paths. The methods include NNs, such as Behler–Parrinello neural networks (BPNNs)^{36,37} and SchNet^{41–43} architectures, and kernel-based methods, such as sGDML^{16–19} and Gaussian Approximation Potentials (GAPs)^{26,27} using the Smooth Overlap of Atomic Position (SOAP) representation.²⁸ To highlight how the complexity in learning the PES increases with the flexibility of a molecule, we compare the results of azobenzene with those of a simpler glycine molecule. We limit the training datasets to 1000 geometries. The ML models unable to predict the PES of the considered small molecules correctly within this limit would face considerable problems for large flexible molecules where the cost of reference calculations increases very steeply.

The structure of this article is given as follows: in Sec. II, we present the isomers of glycine and azobenzene, build possible transition paths between them, and construct the datasets for training ML models. In Sec. III, we discuss the frequent pitfalls of the state-of-the-art ML methods that describe different configurations of flexible molecules. Then, in Sec. IV, we assess the performance of ML models (such as sGDML,^{16–19} GAP/SOAP,^{26–28} BPNN,^{36,37} and SchNet^{41–43}) for azobenzene and glycine molecules trained for both equilibrium states, as well as the transition paths between them. In Sec. V, we describe the challenges of ML force fields (MLFFs) when applied to flexible molecules. Section VI contains the conclusions and an outlook.

II. CONSTRUCTING REFERENCE DATASETS FOR ISOMERIZATION

The starting point for building any MLFF is collecting reference data covering the relevant parts of the PES of interest. When modeling transition pathways, the reference data can be split into two parts: (i) data covering the vicinity of the equilibrium states between which the transition process happens and (ii) data of “far-from-equilibrium” parts of the PES defining the transition path(s). While the equilibrium states are normally readily available, configurations describing the transition paths connecting them are, in most cases, not trivial to find. Moreover, the complexity of this task rapidly grows with the increase in flexibility and size of the molecule. In view of this, we had to employ two different strategies for generating the datasets for glycine and azobenzene isomerization. Below, we discuss in detail the process followed for each molecule separately. All calculations (unless specified otherwise) were performed in FHI-aims software⁷¹ using the Perdew–Burke–Ernzerhof (PBE) exchange–correlation functional⁷² with tight settings and the Tkatchenko–Scheffler (TS) method⁷³ to account for van der Waals (vdW) interactions. For all MD simulations, the i-PI package⁷⁴ was wrapped with FHI-aims. Detailed information of all methods and MD simulations, as well as all relevant configurations and datasets, are available in the [supplementary material](#).

A. Glycine

Glycine, being a rather small molecule, possesses numerous planar and non-planar conformers in the gas phase whose relative energies have been extensively studied.^{75,76} Here, we consider the isomerization from the global minimum geometry, called Ip, to the IIIp conformer because it is the closest “directly connected” minimum. Transitions to any other metastable state from Ip go through this conformer. The Ip–IIIp transformation can be characterized by a change in torsional angles τ_1 and τ_2 (see Fig. 1), both around the C bond. They go from 180.0° and 0.0° in isomer Ip to 0.0° and 180.0° in isomer IIIp, respectively.

1. Transition path

To construct the transition path between the equilibrium states of glycine, we used the string¹¹ and NEB^{12,13} methods. Both methods converge without any issues, providing similar pathways. The transition state obtained by employing the PBE+TS method lies only $2.4 \text{ kcal mol}^{-1}$ above the Ip isomer (see Fig. S1 of the [supplementary material](#)). The mechanism is defined by almost equal rotations of both τ_1 and τ_2 torsional angles (see Table S1 of the [supplementary material](#) for important geometric details).

2. Dataset

Since the relative energy between the Ip isomer and the highest-energy structure found on the minimum energy path (MEP) is less than 3 kcal mol^{-1} , the Ip–IIIp transformation is accessible via standard constant-temperature MD simulations. Therefore, to construct the dataset for glycine isomerization, we ran two dynamics starting from both equilibrium geometries. A total of 5000 configurations at 500 K with a time step of 1 fs were obtained from each simulation.




FIG. 1. Optimized configurations of the minima considered in this work and labels of the main degrees of freedom: the bonds a (C_1-N_2 and N_3-C_4 in azobenzene) and b ($N_2=N_3$ in azobenzene), the bending angles θ and θ' ($C_1-N_2=N_3$ and $N_2=N_3-C_4$ in azobenzene), and the torsional angles ϕ ($C_1-N_2=N_3-C_4$ in azobenzene) and τ_1 and τ_2 ($N_1-C_2-C_3-O_4$ and $N_1-C_2-C_3-O_5$ in glycine).

A transition was observed in the simulation starting from the IIp isomer.

B. Azobenzene

Azobenzene is a photochemical compound; however, it also exhibits a *cis* to *trans* thermal relaxation, on which we focus in this work. Hence, we avoid issues with electronic multi-reference states in azobenzene and use the generalized gradient approximation density functional theory (DFT)-PBE functional and include vdW interactions with the TS-vdW method for generating the reference data. The *cis* and *trans* configurations of azobenzene (Fig. 1) differ mostly by a change in the torsional angle ϕ around the $N=N$ double bond from close to 10.0° to 180.0° during the isomerization. Although the existence of the two forms has been known since the works of Hartley in the 1930s,^{77,78} there is still an open debate regarding whether azobenzene primarily follows a rotation (changes around the dihedral angle ϕ), an inversion (changes in the angles θ and θ'), or a rotation assisted by an inversion (changes in both ϕ and θ and θ') mechanism. DFT, multi-reference methods, and ML approaches have been used in an attempt to unveil the actual mechanism of isomerization,^{66–70,79–81} but conclusive evidence favoring a particular mechanism is yet to be found.

1. Transition paths

Although azobenzene is not much larger than glycine, all the transition pathways that can be found in the literature for this molecule are constructed manually. One can easily check that neither the string nor NEB methods converge to a reasonable path for *cis* to *trans* transition. Following the previous works,^{66–70} we also constructed the transition pathways manually as follows:

- The rotation path, which is defined by a change in the torsional angle ϕ around the central double bond (see Fig. 1).

- The inversion path whose main feature is the bending of either θ or θ' (see Fig. 1).
- The rotation assisted by an inversion path, which is the combination of the first two.

Each path is comprised of 15 intermediate geometries linking the minima. In all cases, the molecule was forced to follow the desired mechanism by linearly interpolating the main degree(s) of freedom between both minima (Tables S2–S4 show important geometric data, and Fig. S1 shows the energy profiles). The obtained highest-energy geometries are in good agreement with those found elsewhere.^{66–69}

Table I shows the relative energies (with respect to cisazobenzene) of the highest-energy structures found for each transition path. The rotation mechanism is the most favorable path at the initial and final steps of the isomerization, but it has the highest-energy barrier among the three transition paths considered here. The inversion mechanism is the one with the lowest-lying highest-energy structure within PBE+TS calculations. The rotation assisted by an inversion path is the least favored at the zones close to the minima and presents a plateau region at the top of the curve with two “peaks” with relative energies close to that of the highest-energy configuration in the inversion mechanism.

The pathways introduced in the previous paragraphs are just linear interpolations between the *cis* and *trans* geometries. This introduces constraints on how the different degrees of freedom can

TABLE I. Relative energies (ΔE in kcal mol^{−1}) of the highest-energy structures on each mechanism computed with the PBE+TS method.

Mechanism	Rotation	Inversion	Rot+Inv
ΔE	30.2	27.4	27.5, 27.6

evolve through the transition. To obtain a path affected by the contributions of all the important degrees of freedom, we “optimized” the rotation path by choosing the values of θ , θ' , and a and b (Fig. 1) that minimize the energy at each step. The optimized rotation path is the most favorable with the PBE+TS method with an energy barrier of 26.1 kcal mol⁻¹. The geometric details of this path can be found in Table S5.

All the paths described here can be considered as good insights into the real isomerization process. It has been found that the activation barrier of the *cis* to *trans* thermal relaxation in *n*-heptane solution is between 22.7 kcal mol⁻¹ and 25.1 kcal mol⁻¹.⁸⁰ Hence, in the gas phase, we would expect greater values like those presented here. In what follows, we will focus on the optimized rotation (named simply rotation from now on) and the inversion mechanisms.

2. Datasets

Constructing a dataset for a molecule such as azobenzene requires a more elaborate procedure compared to the simpler glycine molecule. First of all, the transition process is a rare event at ambient conditions and cannot be easily accessed; second, there are more than one possible transition pathways.

Here, we build separate datasets for the rotation and inversion mechanisms. We first combined two types of MD simulations: (a) long constant-temperature MD runs with a time step of 1 fs at 300 K at the PBE+TS/light level of theory starting from the equilibrium geometries, from which we selected a configuration every 25 steps and carried out single-point calculations with the PBE+TS/tight method (around 3500 configurations were collected for each minimum); (b) constant-temperature MD runs of 300 steps with a 0.5 fs time step at 100 K starting from each of the intermediate steps of the rotation and inversion paths. From the energy distributions shown in Fig. 2(a), one can conclude that the configurations visited during our MD simulations are bounded by the temperature to certain energy ranges, as indicated by the well-defined peaks representing *trans*- and *cis*-like configurations. The addition of configurations obtained from the rotation and inversion mechanisms had little impact on the energy distribution, which leads us to conclude that the transition parts are still poorly sampled. Thus, we generated additional configurations by performing (a) four constant-temperature MD simulations (of 2500 steps each) at 750 K with a time step of 1 fs starting from structures close to equilibrium and (b) constant-temperature MD simulations of 2500 steps at 50 K with a very small time step (0.025 fs) starting from steps 7 and 9 of the inversion path and the steps 8 and 10 of the rotation path (see Tables S3 and S5 for details). The former provide the data required to model the cooling down process from transition states to minima, which involves high kinetic energies. The latter allow us to include slow changes in the degrees of freedom during the transition process. Figure 2(b) shows that the new reference geometries sample different energy distributions for close-to-equilibrium, rotation, and inversion datasets. The final datasets combine the results of all four types of simulations containing 26 455 data points for the rotation and 25 528 data points for the inversion mechanisms.

An alternative to the above-described procedure would be to use an enhanced sampling technique (e.g., metadynamics or

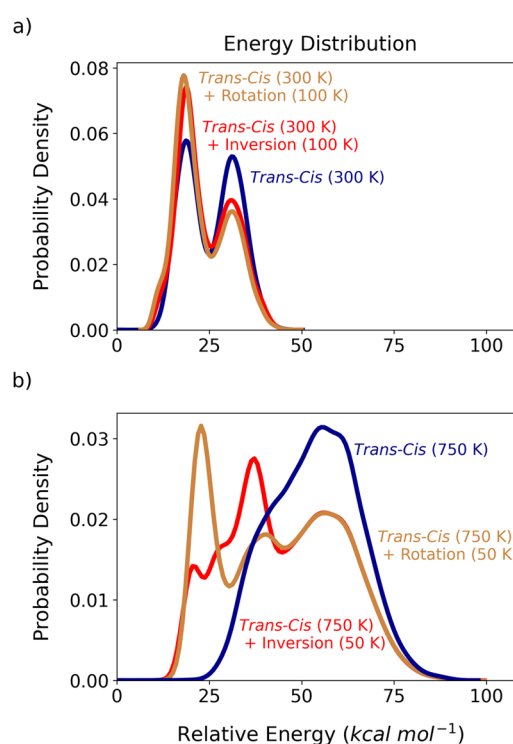


FIG. 2. Distribution of relative energies (in kcal mol⁻¹, with respect to the optimized configuration of transazobenzene) of (a) “close-to-equilibrium” configurations at 300 K and geometries close to the transition paths at 100 K and (b) high-energy “close-to-equilibrium” configurations and properly sampled geometries close to the transition paths.

umbrella sampling). The key component in such methods is the introduction of a bias potential that penalizes already visited system configurations. To do so, one defines a low dimensional projection of the configurational space of the system, the so-called collective variables (CVs), where energy minima and transition paths are well distinguished. During the simulation, a repulsive (bias) potential is added every n steps in the CV space to prevent revisiting already sampled parts of the configuration space. As a result of this procedure, all configurations have an equal probability of being sampled, regardless of their potential energy. Thereby, broad regions of the PES containing multiple minima and transition pathways can be sampled. In our work, we are interested in specific parts of the PES: the regions surrounding the *cis* and *trans* isomers of the azobenzene molecule and the transition pathways between them. While the two angles of azobenzene could be used as CVs, it is not clear that these two variables are sufficient to properly resolve different transition mechanisms. In addition, the transition state energies are high compared to the thermal energy of the molecule at ambient conditions. As a result, choosing the parameters of the bias potential for an efficient metadynamics simulation yields a dataset where the transition mechanisms are undersampled. For a more detailed analysis, see Sec. V. On the contrary, setting the width and the amplitude of the bias potential that allows for a good

resolution in the transition region makes the metadynamics simulation inefficient. Hence, the construction of datasets combining manual and MD-generated configurations is more efficient for sampling high-energy transitions, giving us the opportunity to assess the performance of ML models on the transition paths of flexible molecules.

III. ADVANTAGES AND LIMITATIONS OF DIFFERENT ML METHODS

Before applying any ML method, the molecular configurations must be encoded into a rotationally, translationally, and permutationally invariant representation or “descriptor”. There are many descriptors for MLFF available in the literature,^{22–27,44–55} and efforts to find suitable representations are still ongoing. We remark that a descriptor must balance efficiency with accuracy; hence, different descriptors are applicable to different scenarios. One can divide them into local and global descriptors. For the former, neural networks (e.g., SchNet^{41–43}) and kernel-based potentials (e.g., when using descriptors such as FCHL^{22,55} or SOAP²⁸) assume locality through the introduction of a cutoff radius, and the interactions between atoms are modeled as a sum of individual atomic contributions. Conversely, global descriptors (such as inverse pair-wise distances^{16,44,62}) can serve to build models where the prediction is obtained for the whole structure. Both approaches have their own advantages: for instance, while local descriptors can identify similar neighborhoods in small molecules that can later be transferred to larger systems, global descriptors can capture all interactions of a given system whenever the reference calculations contain the relevant information. However, descriptors also have their pitfalls, some of which arise with large flexible molecules and might become problematic when dealing with complex processes, such as those happening along transition paths.

The first immediate issue that one can foresee is the limited reach of local descriptors, as imposed by the selection of the cutoff radius of the atomic environments. Figure 3 shows the interatomic distance distribution in glycine and azobenzene rotation datasets. While for glycine, the largest distances remain below 6 Å and values

lower than 4 Å are the most populated, in azobenzene, the distances present values of up to 12 Å, and distances between 5 Å and 8 Å are rather common. Thus, local descriptors might already face problems with molecules as large as azobenzene when relevant interactions fall outside their scope (see the results with GAP/SOAP in Sec. IV). An example of such interactions is the long-range ones, which play an important role in azobenzene isomerization as suggested from the paths constructed in Sec. II. Specifically, vdW interactions decrease going from the *cis* to the *trans* configuration, and these interactions lead to an increment in the energy barriers of all paths of more than 1.0 kcal mol^{−1} (the details are shown in Tables S2–S5). Increasing the cutoff radius appears to be a straightforward solution, but the potential gain in accuracy might lead to a significant loss in efficiency.

The second problem affects both local and global descriptors and is related to the scope of the descriptor itself. If some important features are not included or not properly represented in the selected descriptor, the method will fail to achieve appropriate performance (see the results with sGDML in Sec. V). This could well apply to transition paths, where small changes in specific degrees of freedom of the molecule result in considerable energy variations. The addition of the relevant features to the descriptor might alleviate this issue but requires *a priori* knowledge of the studied system.

IV. ACCURACY OF ML MODELS FOR TRANSITION PATHS

Although ML potentials have evolved successfully, there are many open challenges. Among others, the problem of building accurate and data-efficient ML models for flexible molecules describing equilibrium states and the transition pathways between them deserves special attention. Below, we assess the performance of the state-of-art ML models (BPNN,^{36,37} SchNet,^{41–43} GAP/SOAP,^{26–28} and sGDML^{16–19}) on the PES of glycine and azobenzene molecules. The cutoff radius in BPNN, SchNet, and GAP/SOAP was set to a typical value of 5 Å, although additional tests with GAP/SOAP were done with larger cutoffs. All other important settings for each ML method are given in the [supplementary material](#). Namely, we used (a) the glycine dataset and (b) the inversion and rotation datasets of azobenzene (see Sec. II).

The training and test sets were created as follows: first, we used the training set selection process of sGDML (which draws a sample from the dataset that preserves its original energy distribution) to create subsets from each of the datasets used in this work. The subsets have a size equal to five times the number of training points (e.g., for 1000 training points, a subset of 5000 configurations was constructed). We then performed fivefold cross-validation on each subset, using a single fold for training and the rest for testing. For instance, when using 1000 training points, we tested our model with the remaining 4000 configurations (which represents around 15% of the datasets of azobenzene and 40% of the glycine dataset). We randomly created the cross-validation tasks while ensuring that in each fold the energy distribution of the whole dataset was preserved. In this way, each of the folds is representative of the original dataset such that testing our models on these folds will give the same results as testing on the complete dataset. The proposed training/test set selection procedure is comparable to a default random scheme, as




FIG. 3. Distribution of interatomic distances (in Å) in glycine and azobenzene rotation datasets.

implemented in SchNet, while providing more reliable and accurate ML models.

Figure 4 shows the energy and force prediction accuracy for the best models out of all cross-validation tasks. For SchNet, BPNN, and GAP/SOAP, the best model for a given training set size is the one

showing the lowest energy root mean squared error (RMSE) in the test set because these methods are trained on energies and forces. In the case of sGDML, the best model usually is the one with the lowest force RMSE in the test set because sGDML is trained only on forces. However, if two or more sGDML models with similar force RMSE

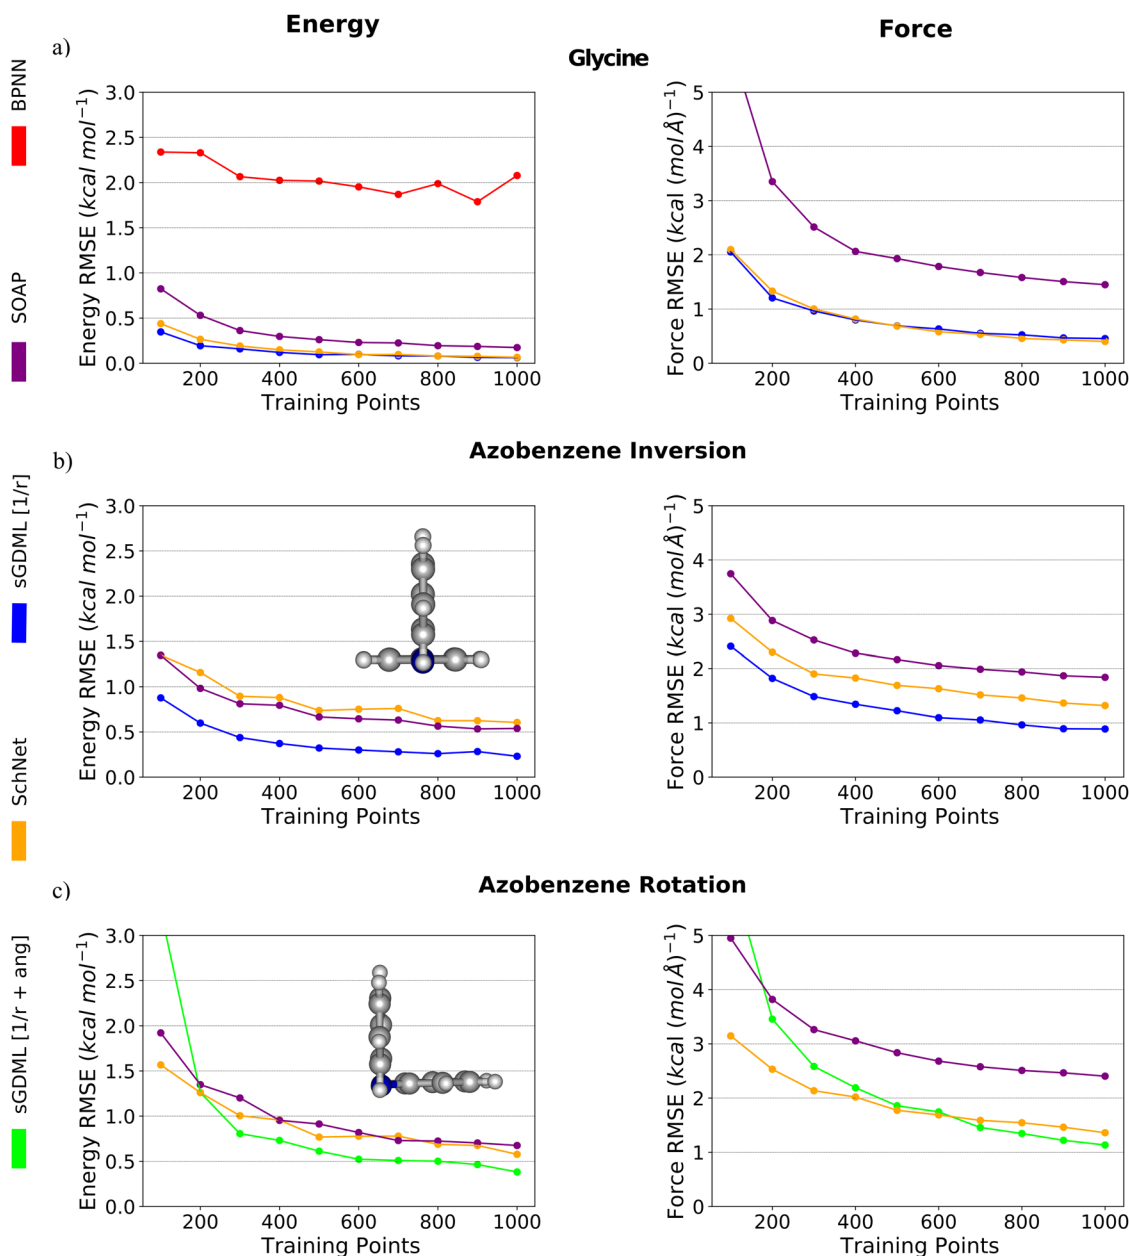


FIG. 4. Energy (in kcal mol⁻¹) and force [in kcal (mol Å)⁻¹] prediction accuracy in terms of root mean squared error (RMSE) as a function of training set size for (a) glycine, (b) inversion, and (c) rotation datasets of azobenzene using the best models of BPNN, GAP/SOAP, SchNet, and sGDML with the default descriptor (sGDML [1/r]) or with the extended descriptor (sGDML [1/r + ang]) out of all cross-validation tasks. Only models with errors below 5.0 kcal (mol Å)⁻¹ and 3.0 kcal mol⁻¹ are shown. (b) and (c) show a side view of the highest-energy structure on the inversion and rotation paths of azobenzene, respectively.

[$\sim 0.1 \text{ kcal (mol } \text{\AA})^{-1}$] present substantially different energy RMSEs ($\sim 1.0 \text{ kcal mol}^{-1}$), we favored the one with the lowest energy RMSE. In the following paragraphs, the errors we discuss correspond to those of the best models.

For the small glycine molecule [Fig. 4(a)], BPNN presents high errors, with a RMSE above $6.0 \text{ kcal (mol } \text{\AA})^{-1}$ for forces and around $2.0 \text{ kcal mol}^{-1}$ for energies, even after using 1000 training points. sGDML and SchNet perform much better with errors below $1.0 \text{ kcal mol}^{-1}$ and $1.0 \text{ kcal (mol } \text{\AA})^{-1}$ with 300 and 400 training points, respectively. GAP/SOAP also shows a good performance in energy prediction with errors below $1.0 \text{ kcal mol}^{-1}$ after using 100 training points, although it is less accurate when predicting forces [errors remain around $1.5 \text{ kcal (mol } \text{\AA})^{-1}$ with 1000 training points]. Based on this analysis, we are henceforth not considering BPNN as a valid candidate to reproduce a more complex PES of flexible molecules using a limited amount of training points, instead focusing on GAP/SOAP, sGDML, and SchNet.

For the azobenzene datasets, the results for different models show high variability. GAP/SOAP obtains an error in energies below $1.0 \text{ kcal mol}^{-1}$ for the inversion mechanism with only 200 training points [Fig. 4(b)]. However, for the rotation mechanism, this performance is achieved with 400 training points [Fig. 4(c)]. In addition, force prediction accuracy is worse for the rotation mechanism [remains above $2.4 \text{ kcal (mol } \text{\AA})^{-1}$ with 1000 training points] than for the inversion mechanism [remains above $1.8 \text{ kcal (mol } \text{\AA})^{-1}$ with 1000 training points] along the whole learning curves. This means that the parts of the PES that are covered by each transition process cannot be evenly described with the same mapping. Indeed, SOAP learns the local information (within the selected cutoff radius), which quickly saturates with the increase in the training set, but cannot equally capture the relevant interactions of both mechanisms.

The sGDML model achieves an outstanding performance for both transition mechanisms. For the inversion dataset, errors in energy go below $1.0 \text{ kcal mol}^{-1}$ with 100 training points and errors in forces go below $1.0 \text{ kcal (mol } \text{\AA})^{-1}$ with 800 training points [Fig. 4(b)]. For the rotation dataset, 300 training points are needed to obtain an energy RMSE lower than $1.0 \text{ kcal mol}^{-1}$ and the force RMSE gets close to $1.1 \text{ kcal (mol } \text{\AA})^{-1}$ after using 1000 training points. To achieve this performance, however, one requires different descriptors for different mechanisms: the default descriptor (inverse pair-wise distances) for the inversion mechanism and an extended descriptor [inverse pair-wise distances, and information on bonded angles and dihedrals in the form $D_\Theta = (1 - e^{-\Theta})^2 - 1$ and $D_\Phi = 1 + \cos \Phi$, respectively]. The features in a given descriptor are not evenly important for all parts of the PES. Thus, a single descriptor might be unable to properly resolve all relations between the relevant degrees of freedom for the PES and each of its transition mechanisms.

SchNet also achieves a “chemical accuracy” of $1.0 \text{ kcal mol}^{-1}$ for both inversion and rotation datasets, but, contrary to sGDML, it can deal with both mechanisms using the same settings [Figs. 4(b) and 4(c)]. However, SchNet is less reliable than sGDML when predicting forces, with an overall RMSE of around $1.4 \text{ kcal (mol } \text{\AA})^{-1}$ for both the inversion and the rotation datasets with 1000 training points. Better performance can be expected with larger training sets. However, this is a trivial solution limited, in practice, by the increased computational costs of the reference data for larger

molecular sizes. The reason that SchNet outperforms GAP/SOAP is also clear: even though SchNet primarily learns local features, it can learn other interactions by embedding such features into the local environments for different parts of the molecule.

V. CHALLENGES FOR ML MODELS IN FLEXIBLE MOLECULES

Even though GAP/SOAP, sGDML, and SchNet methods are able to learn the PES of the azobenzene molecule with chemical accuracy, there is a considerable difference between the predictions for different methods, as well as for different transition mechanisms within the same method. Below, we demonstrate that these contrasting results are caused by imperfections of the implemented training set selection schemes (suboptimal for a complex PES with multiple local minima), as well as intrinsic limitations of the employed descriptors (unable to equally capture all interactions). To do so, we (i) explore the dependence of the performance of ML methods on the specific selection of a training set by considering the results of all cross-validation tasks and (ii) modify the descriptors of GAP/SOAP and sGDML models.

The prediction accuracy of GAP/SOAP models is not considerably affected by the particular choice of a training set of each cross-validation task. The average energy RMSE over all cross-validation tasks (see Fig. 5) is practically the same as the RMSE of the best model (see Fig. 4), with both errors going below $1.0 \text{ kcal mol}^{-1}$ with 400 training points. The main shortcoming of GAP/SOAP models seems to be the learning of long-range interactions. This would explain the $\sim 0.3 \text{ kcal mol}^{-1}$ larger RMSE for the rotation mechanism (Fig. 5). Indeed, due to the different mutual orientations of the benzene rings [see geometries of the highest-energy structures on each transition path in Figs. 4(b) and 4(c)], the vdW energy contribution along the rotation transition pathway is, in average, larger by $\sim 0.3 \text{ kcal mol}^{-1}$ than that for the inversion one. To further explore this conclusion, we removed the vdW correction to the energy and forces from our rotation dataset (i.e., we constructed a dataset at the PBE level of theory) and retrained GAP/SOAP models using the same cross-validation tasks. Unexpectedly, the prediction accuracy of all models was the same with or without the inclusion of long-range vdW interactions. This means that GAP/SOAP might face a different problem, which could be related to an issue in the mapping from configuration space to the PES for the rotation mechanism. Since this error is comparable to the vdW energy contributions, neglecting it is comparable to ignoring non-covalent vdW interactions. Increasing the cutoff radius while keeping the same number of basis functions does not resolve this issue. Figure 6(a) shows the best energy and force RMSE as a function of the cutoff radius in GAP/SOAP models with 12 radial and six angular functions. In fact, both energy and force prediction become slightly worse. Increasing the number of basis functions to alleviate this issue would lead to computationally expensive ML models, impractical for realistic applications.

The sGDML model provides the most accurate and data-efficient FF but faces two important issues. The first one is the descriptor, as shown in Fig. 4, where we had to employ different descriptors for different transition mechanisms. Specifically, for the inversion mechanism, the default sGDML descriptor (inverse

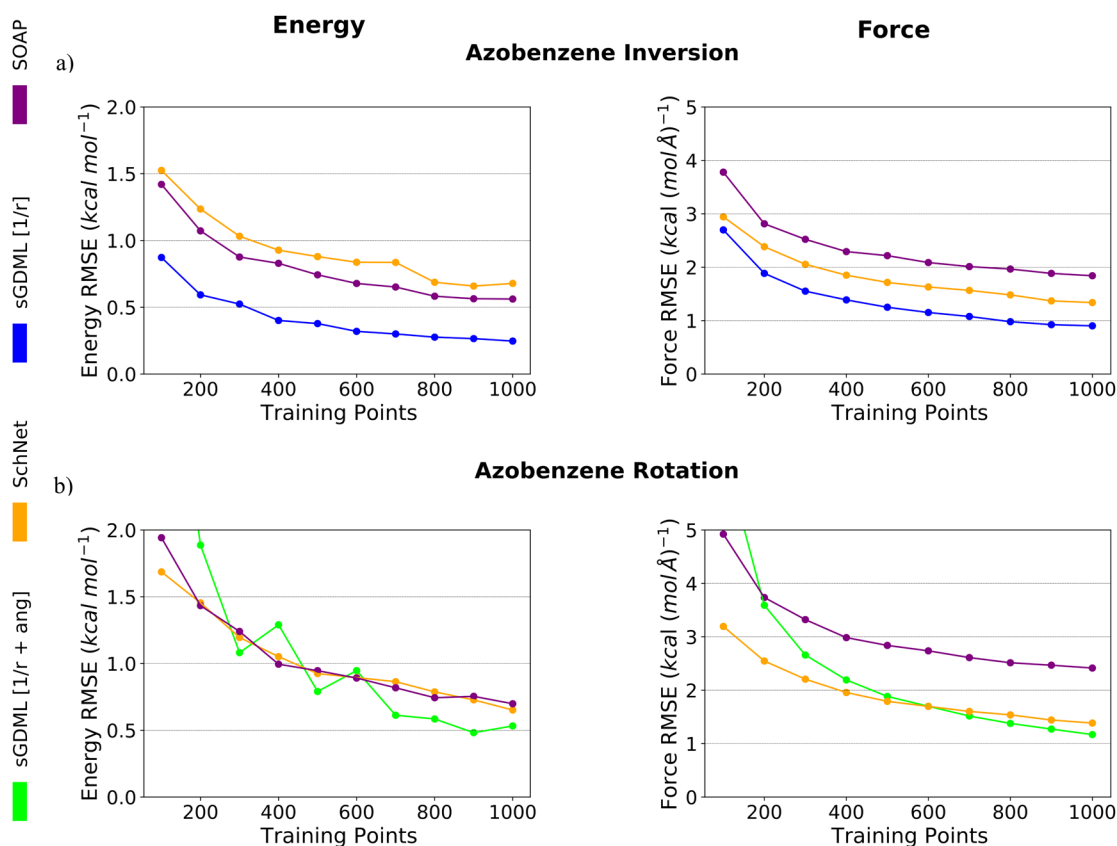


FIG. 5. Mean of the energy (in kcal mol⁻¹) and the force [in kcal (mol Å)⁻¹] prediction accuracy over all cross-validation tasks in terms of root mean squared error (RMSE) as a function of training set size for SchNet, SOAP, and sGDML with the default descriptor (sGDML [1/r]) and with the extended descriptor (sGDML [1/r + ang]) for (a) the inversion and (b) the rotation datasets of azobenzene. Only models with errors below 5.0 kcal (mol Å)⁻¹ and 2.0 kcal mol⁻¹ are shown. In the case of sGDML models, we only show the curve of the best performing descriptor for each mechanism.

pair-wise distances) is sufficient. In contrast, a reliable description of the rotation mechanism requires the inclusion of information about angles and dihedrals in the form $D_{\Theta} = (1 - e^{-\Theta})^2 - 1$ and $D_{\Phi} = 1 + \cos \Phi$, where Θ and Φ are any bonded angle and dihedral of the molecule in radians, respectively. The performance of the default and extended descriptors is shown in Fig. 6(b). The default descriptor for the rotation mechanism and the extended descriptor for the inversion mechanism present considerable oscillations in the energy error as a function of training set size [Fig. 6(b)], which is unacceptable behavior for a reliable ML model. The reason is simple: on the one hand, inverse pair-wise distances cannot correctly resolve the states along the transition path of the rotation mechanism, which are defined by changes in the dihedral angle ϕ . On the other hand, all angles and dihedrals are not equally representative of the inversion mechanism, and adding them misleads the model in this case.

We would like to remark that even when using an appropriate descriptor, the training set selection is still crucial. Out of all the methods, sGDML is the most affected by the model selection in

cross-validation tasks [compare Figs. 4(c) and 5(b)]. While for 1000 training points, the difference between the average energy RMSE over all cross-validation tasks and the energy RMSE of the best model is of only 0.2 kcal mol⁻¹, for 400 training points, this difference is as large as 0.6 kcal mol⁻¹. Thereby, one needs to be very careful when selecting the best sGDML models. Furthermore, the training set dependency often leads to models with a similar force RMSE but considerably different energy RMSEs. For instance, the average difference in energy RMSEs between the best and the worst models for the rotation mechanism using the extended descriptor (over all training set sizes considered here) is of 0.7 kcal mol⁻¹, while the average difference in force RMSEs is less than 0.2 kcal (mol Å)⁻¹. To understand this behavior, one needs to recall that sGDML models contain two hyperparameters apart from the regularization. One is the width of the kernel, which is defined by optimizing the force predictions. The other one is the constant shift for the energy, which is employed to minimize the difference between the prediction results and the energy values in the dataset. The energies of different conformers of flexible molecules

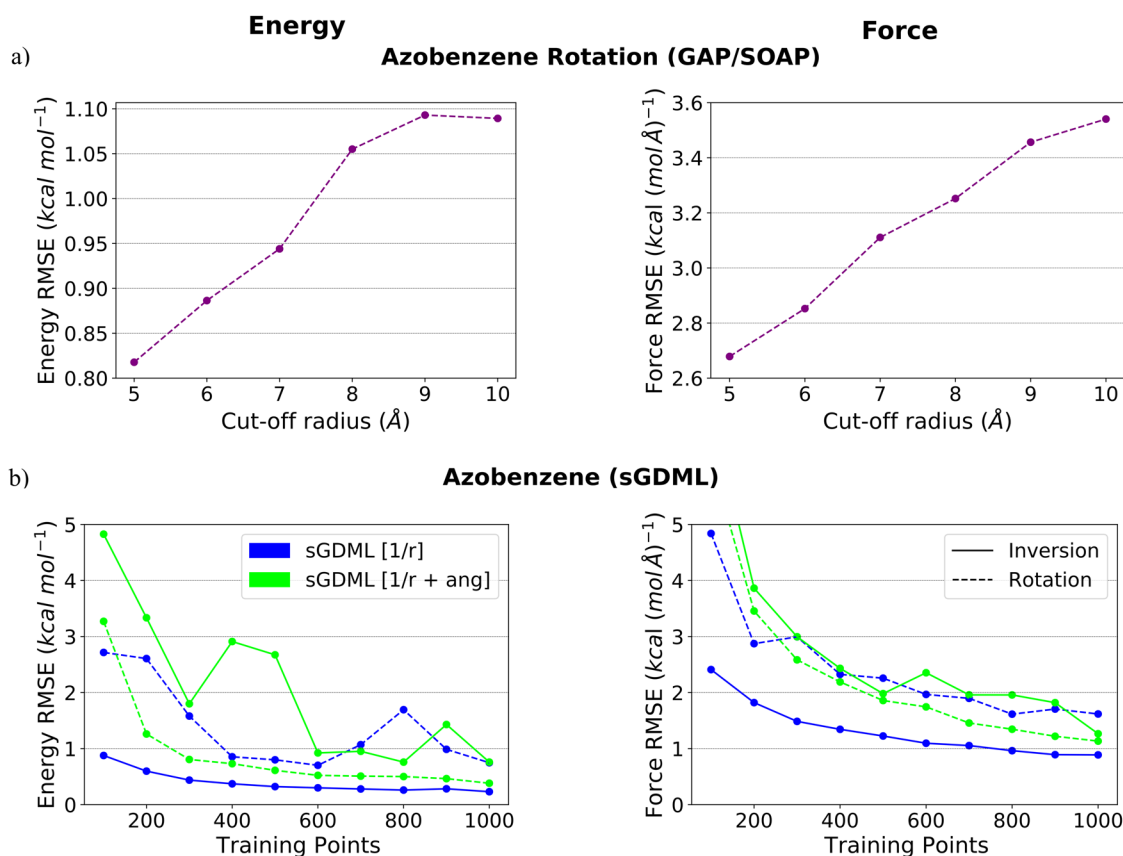


FIG. 6. Energy (in kcal mol⁻¹) and force [in kcal (mol Å)⁻¹] prediction accuracy in terms of root mean squared error (RMSE) (a) as a function of cutoff radius for GAP/SOAP models of the rotation dataset of azobenzene trained with 600 training points and (b) as a function of training set size for sGDML with the default descriptor (sGDML [1/r]) and with the extended descriptor (sGDML [1/r + ang]) for the inversion and rotation datasets of azobenzene. Only models with errors below 5.0 kcal (mol Å)⁻¹ and 5.0 kcal mol⁻¹ are shown.

are often degenerate. Consequently, even though all our training sets follow the energy distribution of the complete dataset, they represent different parts of the PES unequally. As a result, the energy shift hyperparameter obtained from a given training set can become suboptimal for the whole dataset. Hence, the force-based model selection scheme, as implemented in sGDML, may lead to large oscillation in energy prediction accuracy as a function of training set size [similar to those in the green solid line in Fig. 6(b)]. To resolve this issue, one should consider both energy and forces to select the optimal model. For each cross-validation task, the training scheme does not change and still relies only on forces, but for selecting the best model out of many possibilities, we also account for the energy prediction accuracy. Summarizing, accurate and data-efficient models are achievable with sGDML, but both descriptors and training sets must be carefully selected.

SchNet is an optimal compromise between GAP/SOAP and sGDML models. Like GAP/SOAP, it does not heavily depend on the specific selection of a training set (see Fig. 5), while being capable to learn all kind of interactions, akin to sGDML (see Fig. 4). As a

result, SchNet reproduces both transition mechanisms equally accurately using the same settings, with errors only slightly larger than those of the sGDML models. This is a consequence of the embedding of local features of different atoms through the interaction layers. To train our SchNet models, we employed six interaction layers with a 5 Å cutoff radius for local environment. This architecture guarantees that we cover all possible interatomic distances within an azobenzene molecule (Fig. 3), making our SchNet models effectively global. However, a good overall RMSE might not always mean a good ML model. Figure 7(a) shows the energy and force prediction accuracy on different clusters of the rotation dataset of the best SchNet model out of the five cross-validation tasks with 1000 training points [RMSEs of 0.6 kcal mol⁻¹ and 1.3 kcal (mol Å)⁻¹ over the complete rotation dataset]. Each cluster corresponds to different values of the dihedral angle ϕ (from the interval between 0° and 10° for index 0 to the interval between 170° and 180° for index 17). One can see that the errors for close-to-equilibrium configurations are four times larger than those for the transition regions. Increasing the training set size and having information of the two mechanisms (i.e.,

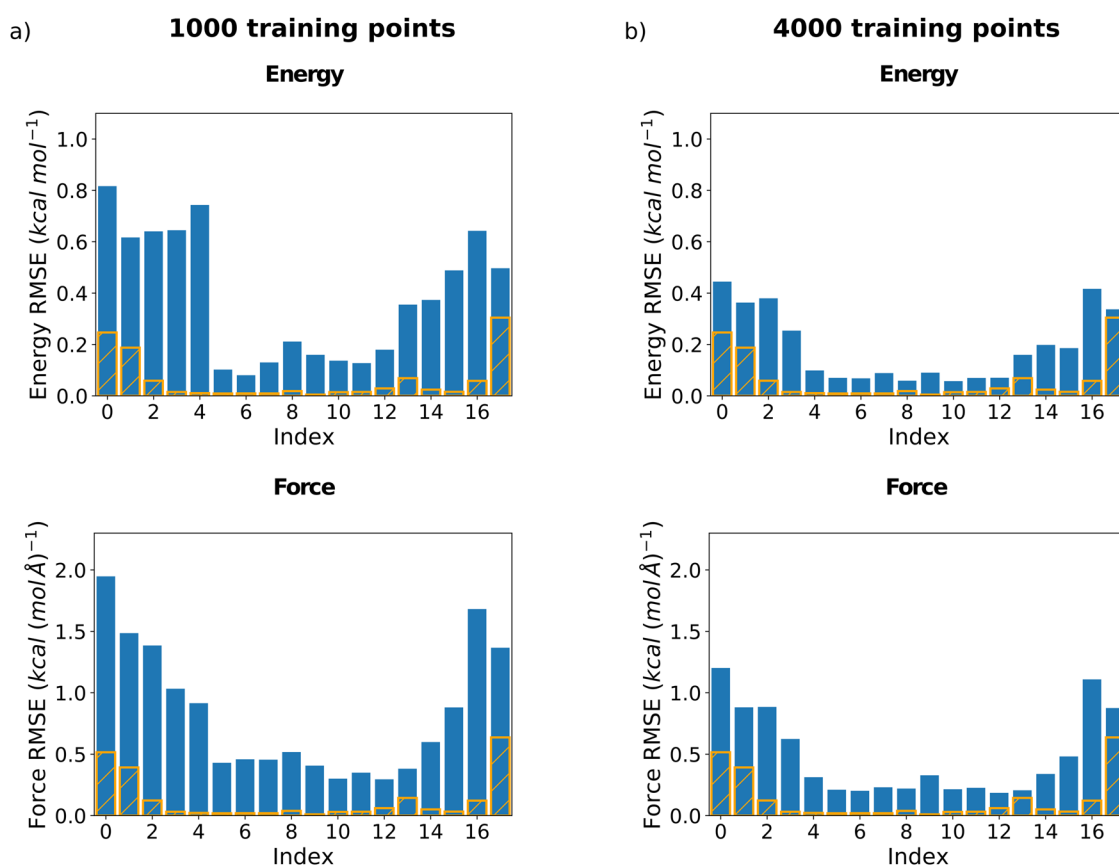


FIG. 7. Energy (in kcal mol^{-1}) and force [in $\text{kcal (mol \AA)}^{-1}$] prediction accuracy in terms of root mean squared error (RMSE) for different clusters of the rotation dataset using (a) the best SchNet model obtained with 1000 training points of the rotation dataset [RMSEs of $0.6 \text{ kcal mol}^{-1}$ and $1.3 \text{ kcal (mol \AA)}^{-1}$ over the complete rotation dataset] and (b) a SchNet model trained with 4000 configurations combining the rotation and inversion datasets [RMSEs of $0.4 \text{ kcal mol}^{-1}$ and $0.8 \text{ kcal (mol \AA)}^{-1}$ over the complete rotation dataset]. The configurations for each cluster were selected depending on the value of the dihedral angle ϕ , going from an interval between 0° and 10° for index 0 to an interval between 170° and 180° for index 17. Relative population of each cluster is also indicated (orange solid lines, arbitrary units).

adding the data of the inversion mechanism to the rotation dataset) does not change this ratio [Fig. 7(b)]. We remark that while building these models, we expect to be able to use them for accurate simulations of the transition mechanism, computing reaction rates, and lifetime of *cis* configuration. In order to achieve this goal, one needs to ensure that the MLFF is equally accurate for all relevant parts of the PES. As one can see from Fig. 7, this requirement is not fulfilled by the obtained SchNet FFs. Importantly, by using a total of only 600 training points for both *trans*- and *cis*-like configurations (300 for each isomer), one can train a single sGDML model that reproduces the performance of the 4 K SchNet model on clusters 0, 1, 2, 16, and 17. Hence, while learning the entire PES of the azobenzene molecule is possible within a SchNet model, this approach is not particularly data efficient. There are two options to solve this problem: one is to use a training set optimization technique, flattening the prediction across the configuration space.⁸² The second option would be to design schemes that combine a set of local models into a global one, finding optimal descriptors, training sets, and models for each part of the PES.

One might think of training ML models on a dataset constructed from metadynamics simulations, called ML_{meta} models hereafter (see the [supplementary material](#) for the details of the metadynamics simulation). However, as mentioned in Sec. II, metadynamics datasets of the same size would contain less information about the transition processes as compared to the rotation/inversion datasets. One expects the ML_{meta} models to be less accurate compared to ML models trained on the rotation/inversion datasets (for simplicity, henceforth, we refer to them as ML_{md} models). Figure 8 shows the energy and force prediction accuracy of $\text{SchNet}_{\text{meta}}$ models when applied to reproduce the rotation dataset from Sec. II. Each cluster corresponds to different values of the dihedral angle ϕ (from the interval between 0° and 10° for index 0 to the interval between 170° and 180° for index 17) to be compared to Fig. 7. The average RMSE of the resulting $\text{SchNet}_{\text{meta}}$ models is twice larger than that of $\text{SchNet}_{\text{md}}$ models presented in Fig. 7. In addition, the distribution of errors is drastically different. The largest errors are now in the transition region (clusters 4–13) and can be up to an order of magnitude larger compared to those of the $\text{SchNet}_{\text{md}}$ models. Moreover,

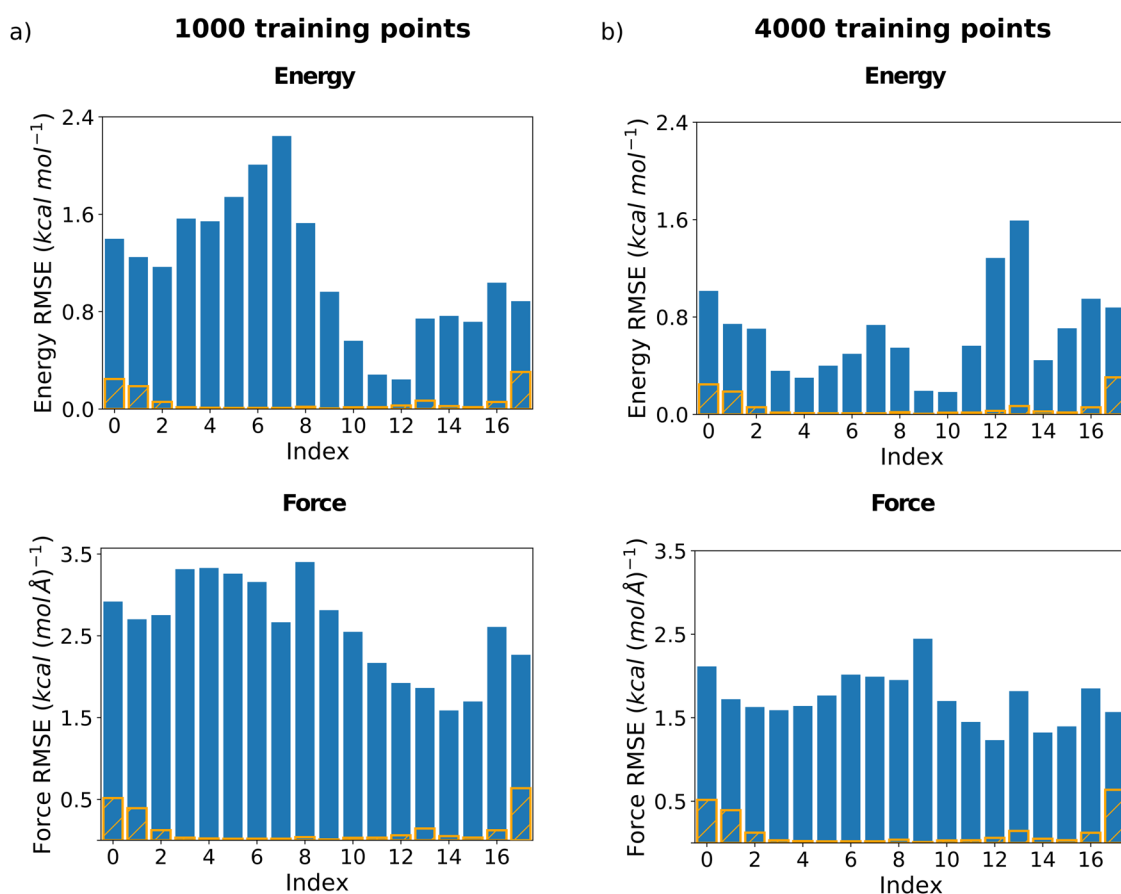


FIG. 8. Energy (in kcal mol⁻¹) and force [in kcal (mol Å)⁻¹] prediction accuracy in terms of root mean squared error (RMSE) for different clusters of the rotation dataset using (a) a SchNet_{meta} model trained with 1000 configurations and (b) a SchNet_{meta} model trained with 4000 configurations. The configurations for each cluster were selected depending on the value of the dihedral angle ϕ , going from an interval between 0° and 10° for index 0 to an interval between 170° and 180° for index 17. Relative population of each cluster is also indicated (orange solid lines, arbitrary units).

the error distribution is not uniform, which means that the different parts of the PES responsible for the transition are not equally represented within the metadynamics dataset.

It is important to mention that several approaches exist to improve ML models to large and flexible molecules, such as multiscale GAP/SOAP,⁴⁸ the recently developed LODE,⁸³ or the addition of physical potentials.⁸⁴ These are all promising directions, and our study demonstrates how more advanced techniques are necessary even for relatively small molecules such as azobenzene. Nevertheless, multiscale approaches might lead to less data-efficient models. Indeed, a single GAP/SOAP model is the most computationally expensive method among those used in this work. Therefore, although a multiscale GAP/SOAP could solve the problem for azobenzene, it could become infeasible when dealing with larger molecules. As for methods to learn long-range interactions, they need to be carefully integrated with local ML models since local properties used in long-range interaction models can vary substantially for different local chemical environments. This is supported by the results we observed when comparing the performance of

GAP/SOAP_{md} models trained on PBE and PBE+TS reference data.

The final step to confirm the reliability of our ML_{md} models is to demonstrate their applicability in MD simulations. Figure 9 shows the values of the dihedral angle ϕ and the angle θ as a function of simulation time obtained with constant-temperature MD runs starting from structures close to the transition states. We selected the best GAP/SOAP_{md}, SchNet_{md}, and sGDML_{md} models trained on 1000 training points to run with each method four MD simulations of 50k steps at 300 K with a time step of 0.5 fs (video files of the simulations can be found in the [supplementary material](#)). For the simulation that starts from step 7 of the inversion mechanism [Fig. 9(a)], GAP/SOAP_{md} and SchNet_{md} models lead to the *trans* isomer, while the sGDML_{md} model leads to the *cis* isomer. For the simulations that start from step 9 of the inversion path and steps 8 and 10 of the rotation path [Figs. 9(b)–9(d)], all models lead to the *trans*, *cis*, and *trans* isomers, respectively. To keep local geometry fluctuations under control during the relaxation process, we choose a friction coefficient for the Langevin thermostat equal to 2 fs. From Fig. 9,

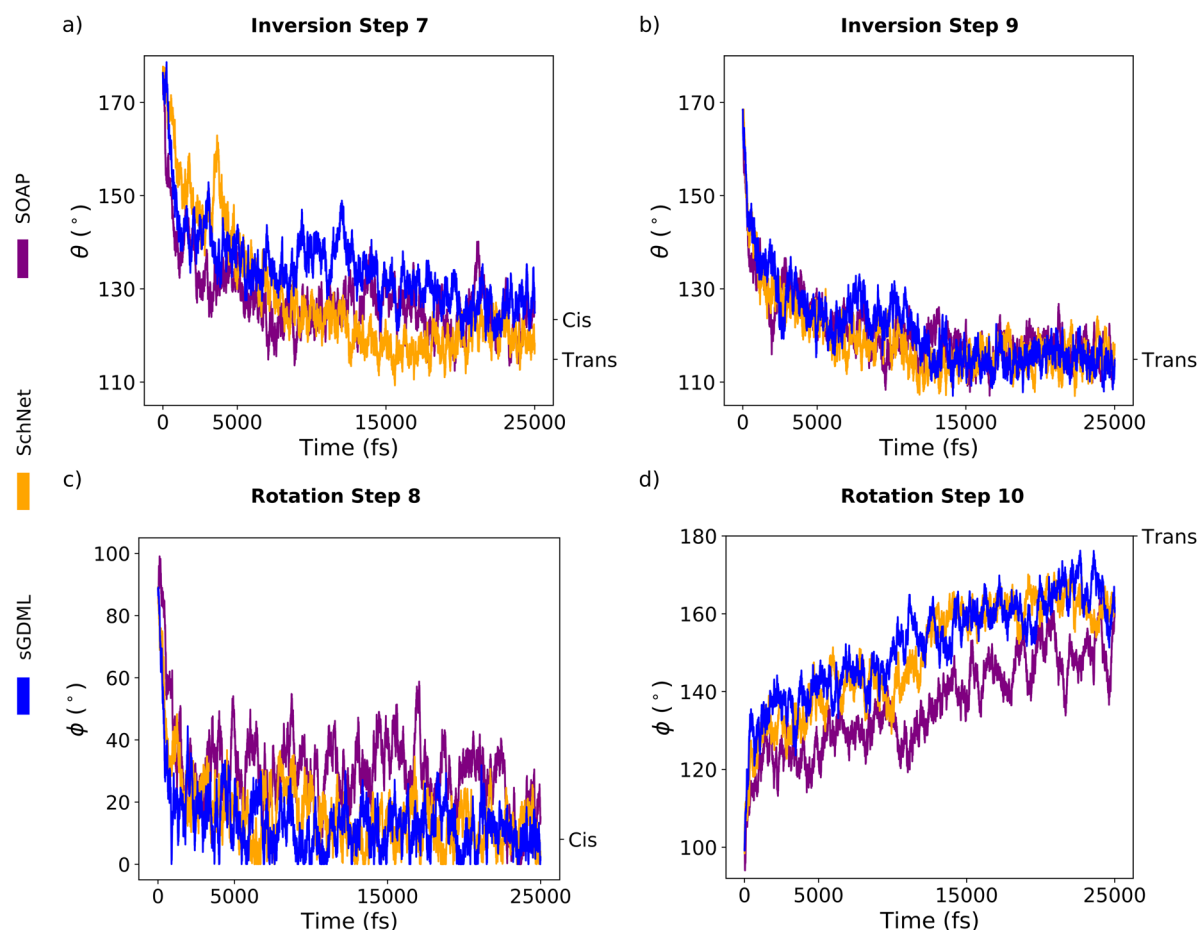


FIG. 9. (a)–(d) Change in the main degree of freedom (ϕ for the rotation mechanism and θ for the inversion one, in $^{\circ}$) as a function of the simulation time (in fs) for the four different MD simulations performed with the best GAP/SOAP_{md}, sGDML_{md}, and SchNet_{md} models trained on 1000 training points. The initial geometry used for each simulation is indicated over the graph. Tick labels on the right mark the value of the degree of freedom in the *trans* isomer and the *cis* isomer.

one can see that all ML_{md} models are able to correctly reproduce the transition processes avoiding unphysical configurations, despite the large fluctuations in the angles caused by the thermostat. This demonstrates that the ML_{md} models are reliable and stable, and the constructed datasets contain the information necessary for studying the *cis* to *trans* thermal relaxation of azobenzene.

Summarizing the results of this section, while the state-of-the-art ML models are capable of reproducing the complex PES of flexible molecules, this challenge is far from being solved in practice. Default approaches demonstrating excellent performance for small molecules or rigid systems struggle with increasing flexibility and dimensionality. Even the best performing models present difficulties to efficiently learn the PES in its entirety. Moreover, the way datasets are constructed is crucial for the performance of the resulting ML models. To overcome this challenge, one should ensure that the descriptors contain all relevant features to capture the complex geometrical transformations in the high-dimensional PES. In addition, training sets must represent all parts of configuration

space, which exhibits high energetic degeneracy, which makes purely following an energy distribution ineffective. Although MD simulations executed with the constructed ML potentials are promising, we suggest further developments of robust approaches for selecting training points, appropriate descriptors, or even using different models for different parts of the PES.

VI. CONCLUSIONS

In the present work, we discussed the challenge of modeling the PES of flexible molecules using the state-of-the-art ML models when using limited sets of training data. Our results show that methods based on local descriptors (e.g., BPNN and GAP/SOAP) saturate quickly with the increase in the number of training points, while not achieving the desired prediction accuracy. This is a consequence of the limitations imposed by the cutoff radius for the atomic environments. The ML methods based on global descriptors (e.g.,

sGDML) require careful model selection procedures when the reference datasets consist of several disconnected parts of the PES. The main challenges are the inability of simple training schemes to select appropriate training datasets in unbalanced reference data and the limitations of the standard molecular descriptors to pick up the features describing the complex geometric transformations in flexible molecules. Finally, end-to-end NNs (e.g., SchNet) do not reproduce all relevant parts of the PES with equal accuracy. Moreover, NNs require a larger amount of training data compared to kernel methods, which can result in a high computational cost when generating reference datasets. All the tested MLFFs in their current form can be further improved for quantitative studies of complex processes in flexible molecules. Important features of ML, such as descriptors and selection of training points, could and should be revised. The varying performance of most of the ML models for the rotation and inversion mechanisms in azobenzene also suggests switching from learning the entire PES within a single task to the employment of multiple local models for different parts of the PES and then combining them into a global FF.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for the following:

- Computational details and tables with the relevant geometric and energetic data of the minima and the structures on the transition paths (Supporting.pdf)
- All xyz files of the minima and the structures on the transition paths (Geometries.zip)
- The datasets used for training and testing in the extended xyz format. Energies are given in kcal mol⁻¹, and forces are given in kcal (mol Å)⁻¹ (Datasets.zip and Dataset_MetaMD.zip)
- The video files (in the mpg format) of the MD simulations performed with the ML potentials (Movies.zip)

ACKNOWLEDGMENTS

V.V.-G. acknowledges financial support from the Luxembourg National Research Fund (FNR) under the program DTU PRIDE MASSENA (Grant No. PRIDE/15/10935404). G.F. acknowledges financial support from the FNR AFR, Project No. 14593813. A.T. received support from the European Research Council (ERC-CoG BeStMo). IP acknowledges financial support from the FNR (C19/MS/13718694/QML-FLEX).

DATA AVAILABILITY

The data that support the findings of this study are available within the article and its [supplementary material](#).

REFERENCES

- ¹P. Hänggi, P. Talkner, and M. Borkovec, *Rev. Mod. Phys.* **62**, 251 (1990).
- ²M. Bachmann and W. Janke, *Phys. Rev. Lett.* **95**, 058102 (2005).
- ³D. Chakraborty and D. J. Wales, *J. Chem. Phys.* **150**, 125101 (2019).
- ⁴J. B. Cook, T. C. Lin, H.-S. Kim, A. Siordia, B. S. Dunn, and S. H. Tolbert, *ACS Nano* **13**, 1223 (2019).
- ⁵Y. Zhang, J. B. Chou, J. Li, H. Li, Q. Du, A. Yadav, S. Zhou, M. Y. Shalaginov, Z. Fang, H. Zhong, C. Roberts, P. Robinson, B. Bohlin, C. Ríos, H. Lin, M. Kang, T. Gu, J. Warner, V. Liberman, K. Richardson, and J. Hu, *Nat. Commun.* **10**, 4279 (2019).
- ⁶C. M. Dobson, *Nature* **426**, 884 (2003).
- ⁷S. Piana, K. Lindorff-Larsen, and D. E. Shaw, *Proc. Natl. Acad. Sci. U. S. A.* **109**, 17845 (2012).
- ⁸J. B. Stiller, S. Jordan Kerns, M. Hoemberger, Y.-J. Cho, R. Otten, M. F. Hagan, and D. Kern, *Nat. Catal.* **2**, 726 (2019).
- ⁹C. Peng and H. Bernhard Schlegel, *Isr. J. Chem.* **33**, 449 (1993).
- ¹⁰C. Peng, P. Y. Ayala, H. B. Schlegel, and M. J. Frisch, *J. Comput. Chem.* **17**, 49 (1996).
- ¹¹E. Weinan, W. Ren, and E. Vanden-Eijnden, *J. Chem. Phys.* **126**, 164103 (2007).
- ¹²H. Jónsson, G. Mills, and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, 1998), pp. 385–404.
- ¹³G. Henkelman and H. Jónsson, *J. Chem. Phys.* **113**, 9978 (2000).
- ¹⁴C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, *J. Chem. Phys.* **108**, 1964 (1998).
- ¹⁵P. G. Bolhuis and G. Csányi, *Phys. Rev. Lett.* **120**, 250601 (2018).
- ¹⁶S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, *Sci. Adv.* **3**, e1603015 (2017).
- ¹⁷S. Chmiela, H. E. Sauceda, K. R. Müller, and A. Tkatchenko, *Nat. Commun.* **9**, 3887 (2018).
- ¹⁸H. E. Sauceda, S. Chmiela, I. Poltavsky, K.-R. Müller, and A. Tkatchenko, *J. Chem. Phys.* **150**, 114102 (2019).
- ¹⁹S. Chmiela, H. E. Sauceda, I. Poltavsky, K. R. Müller, and A. Tkatchenko, *Comput. Phys. Commun.* **240**, 38 (2019).
- ²⁰V. Botu and R. Ramprasad, *Phys. Rev. B* **92**, 094306 (2015).
- ²¹A. S. Christensen, F. A. Faber, and O. A. von Lilienfeld, *J. Chem. Phys.* **150**, 064105 (2019).
- ²²A. S. Christensen, L. A. Bratholm, F. A. Faber, and O. A. von Lilienfeld, *J. Chem. Phys.* **152**, 044107 (2020).
- ²³M. Rupp, R. Ramakrishnan, and O. A. von Lilienfeld, *J. Phys. Chem. Lett.* **6**, 3309 (2015).
- ²⁴A. Glielmo, P. Sollich, and A. De Vita, *Phys. Rev. B* **95**, 214302 (2017).
- ²⁵M. Eickenberg, G. Exarchakis, M. Hirn, S. Mallat, and L. Thiry, *J. Chem. Phys.* **148**, 241732 (2018).
- ²⁶A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).
- ²⁷A. P. Bartók and G. Csányi, *Int. J. Quantum Chem.* **115**, 1051 (2015).
- ²⁸A. P. Bartók, R. Kondor, and G. Csányi, *Phys. Rev. B* **87**, 184115 (2013).
- ²⁹Z. Li, J. R. Kermode, and A. De Vita, *Phys. Rev. Lett.* **114**, 096405 (2015).
- ³⁰E. V. Podryabinkin and A. V. Shapeev, *Comput. Mater. Sci.* **140**, 171 (2017).
- ³¹P. O. Dral, A. Owens, S. N. Yurchenko, and W. Thiel, *J. Chem. Phys.* **146**, 244108 (2017).
- ³²F. Noé, S. Olsson, J. Köhler, and H. Wu, *Science* **365**, eaaw1147 (2019).
- ³³A. Mardt, L. Pasquali, H. Wu, and F. Noé, *Nat. Commun.* **9**, 5 (2018).
- ³⁴J. Behler, S. Lorenz, and K. Reuter, *J. Chem. Phys.* **127**, 014705 (2007).
- ³⁵J. Behler, *Phys. Chem. Chem. Phys.* **13**, 17930 (2011).
- ³⁶J. Behler, *J. Chem. Phys.* **134**, 074106 (2011).
- ³⁷J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).
- ³⁸K. V. J. Jose, N. Artrith, and J. Behler, *J. Chem. Phys.* **136**, 194111 (2012).
- ³⁹J. Behler, *J. Chem. Phys.* **145**, 170901 (2016).
- ⁴⁰M. Gastegger, J. Behler, and P. Marquetand, *Chem. Sci.* **8**, 6924 (2017).
- ⁴¹K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, A. Tkatchenko, and K.-R. Müller, *J. Chem. Phys.* **148**, 241722 (2018).
- ⁴²K. T. Schütt, P.-J. Kindermans, H. E. Sauceda, S. Chmiela, A. Tkatchenko, and K.-R. Müller, in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, (Curran Associates Inc. Long Beach, CA, 2017) p. 992.

- ⁴³K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **15**, 448 (2019).
- ⁴⁴M. Rupp, A. Tkatchenko, K. R. Müller, and O. A. von Lilienfeld, *Phys. Rev. Lett.* **108**, 058301 (2012).
- ⁴⁵K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **9**, 3404 (2013).
- ⁴⁶S. De, A. P. Bartók, G. Csányi, and M. Ceriotti, *Phys. Chem. Chem. Phys.* **18**, 13754 (2016).
- ⁴⁷N. Artrith, A. Urban, and G. Ceder, *Phys. Rev. B* **96**, 014112 (2017).
- ⁴⁸A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. Kermode, G. Csányi, and M. Ceriotti, *Sci. Adv.* **3**, e1701816 (2017).
- ⁴⁹K. Yao, J. E. Herr, and J. Parkhill, *J. Chem. Phys.* **146**, 014106 (2017).
- ⁵⁰F. A. Faber, L. Hutchison, B. Huang, J. Gilmer, S. S. Schoenholz, G. E. Dahl, O. Vinyals, S. Kearnes, P. F. Riley, and O. A. von Lilienfeld, *J. Chem. Theory Comput.* **13**, 5255 (2017).
- ⁵¹A. Glielmo, C. Zeni, and A. De Vita, *Phys. Rev. B* **97**, 184307 (2018).
- ⁵²A. Grisafi, D. M. Wilkins, G. Csányi, and M. Ceriotti, *Phys. Rev. Lett.* **120**, 036002 (2018).
- ⁵³Y.-H. Tang, D. Zhang, and G. E. Karniadakis, *J. Chem. Phys.* **148**, 034101 (2018).
- ⁵⁴W. Pronobis, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **14**, 2991 (2018).
- ⁵⁵F. A. Faber, A. S. Christensen, B. Huang, and O. A. von Lilienfeld, *J. Chem. Phys.* **148**, 241717 (2018).
- ⁵⁶K. Ryczko, K. Mills, I. Luchak, C. Homenick, and I. Tamblyn, *Comput. Mater. Sci.* **149**, 134 (2018).
- ⁵⁷L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, *Phys. Rev. Lett.* **120**, 143001 (2018).
- ⁵⁸K. Shao, J. Chen, Z. Zhao, and D. H. Zhang, *J. Chem. Phys.* **145**, 071101 (2016).
- ⁵⁹K. Yao, J. E. Herr, D. W. Toth, R. Mckintyre, and J. Parkhill, *Chem. Sci.* **9**, 2261 (2018).
- ⁶⁰F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K. R. Müller, *Nat. Commun.* **8**, 872 (2017).
- ⁶¹J. S. Smith, B. T. Nebgen, R. Zubatyuk, N. Lubbers, C. Devereux, K. Barros, S. Tretiak, O. Isayev, and A. E. Roitberg, *Nat. Commun.* **10**, 2903 (2019).
- ⁶²K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, *J. Phys. Chem. Lett.* **6**, 2326 (2015).
- ⁶³H. E. Saucedo, V. Vassilev-Galindo, S. Chmiela, K.-R. Müller, and A. Tkatchenko, *Nat. Commun.* **12**, 442 (2021).
- ⁶⁴R. Meyer, K. S. Schmuck, and A. W. Hauser, *J. Chem. Theory Comput.* **15**, 6513 (2019).
- ⁶⁵L. Pattanaik, J. B. Ingraham, C. A. Grambow, and W. H. Green, *Phys. Chem. Chem. Phys.* **22**, 23618 (2020).
- ⁶⁶P. Cattaneo and M. Persico, *Phys. Chem. Chem. Phys.* **1**, 4739 (1999).
- ⁶⁷A. Cembran, F. Bernardi, M. Garavelli, L. Gagliardi, and G. Orlandi, *J. Am. Chem. Soc.* **126**, 3234 (2004).
- ⁶⁸L. Gagliardi, G. Orlandi, F. Bernardi, A. Cembran, and M. Garavelli, *Theor. Chem. Acc.* **111**, 363 (2004).
- ⁶⁹L. Wang and X. Wang, *J. Mol. Struct.: THEOCHEM* **806**, 179 (2007).
- ⁷⁰P. Tavazde, G. Avendaño Franco, P. Ren, X. Wen, Y. Li, and J. P. Lewis, *J. Am. Chem. Soc.* **140**, 285 (2018).
- ⁷¹V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys. Commun.* **180**, 2175 (2009).
- ⁷²J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- ⁷³A. Tkatchenko and M. Scheffler, *Phys. Rev. Lett.* **102**, 073005 (2009).
- ⁷⁴V. Kapil, M. Rossi, O. Marsalek, R. Petraglia, Y. Litman, T. Spura, B. Cheng, A. Cuzzocrea, R. H. Meißner, D. M. Wilkins, B. A. Helfrecht, P. Juda, S. P. Bienvenue, W. Fang, J. Kessler, I. Poltavsky, S. Vandenbrande, J. Wieme, C. Corminboeuf, T. D. Kühne, D. E. Manolopoulos, T. E. Markland, J. O. Richardson, A. Tkatchenko, G. A. Tribello, V. Van Speybroeck, and M. Ceriotti, *Comput. Phys. Commun.* **236**, 214 (2019).
- ⁷⁵A. G. Csaszar, *J. Am. Chem. Soc.* **114**, 9568 (1992).
- ⁷⁶L. F. Pacios, O. Gálvez, and P. C. Gómez, *J. Phys. Chem. A* **105**, 5232 (2001).
- ⁷⁷G. S. Hartley, *Nature* **140**, 281 (1937).
- ⁷⁸G. S. Hartley, *J. Chem. Soc.* **1938**, 633.
- ⁷⁹E. W. G. Diau, *J. Phys. Chem. A* **108**, 950 (2004).
- ⁸⁰H. M. D. Bandara and S. C. Burdette, *Chem. Soc. Rev.* **41**, 1809 (2012).
- ⁸¹J. Dokić, M. Gothe, J. Wirth, M. V. Peters, J. Schwarz, S. Hecht, and P. Saalfrank, *J. Phys. Chem. A* **113**, 6763 (2009).
- ⁸²See <https://github.com/fonseca/MLFF> for a code to optimize the training set selection process.
- ⁸³A. Grisafi and M. Ceriotti, *J. Chem. Phys.* **151**, 204105 (2019).
- ⁸⁴T. Bereau, R. A. DiStasio, Jr., A. Tkatchenko, and O. A. von Lilienfeld, *J. Chem. Phys.* **148**, 241706 (2018).