

C.1 | Paper I

Reprinted from Schran, C., Brezina, K. and Marsalek, O.: Committee neural network potentials control generalization errors and enable active learning, *The Journal of Chemical Physics*, 153, 10, 104105, 2020, with the permission of AIP Publishing.

Committee neural network potentials control generalization errors and enable active learning

Cite as: J. Chem. Phys. 153, 104105 (2020); doi: 10.1063/5.0016004

Submitted: 31 May 2020 • Accepted: 13 August 2020 •

Published Online: 8 September 2020



Christoph Schran,^{a)} Krystof Brezina,^{b)} and Ondrej Marsalek^{b)}

AFFILIATIONS

Charles University, Faculty of Mathematics and Physics, Ke Karlovu 3, 121 16 Prague 2, Czech Republic

Note: This paper is part of the JCP Special Topic on Machine Learning Meets Chemical Physics.

^{a)} Electronic mail: christoph.schran@rub.de

^{b)} Author to whom correspondence should be addressed: ondrej.marsalek@mff.cuni.cz

ABSTRACT

It is well known in the field of machine learning that committee models improve accuracy, provide generalization error estimates, and enable active learning strategies. In this work, we adapt these concepts to interatomic potentials based on artificial neural networks. Instead of a single model, multiple models that share the same atomic environment descriptors yield an average that outperforms its individual members as well as a measure of the generalization error in the form of the committee disagreement. We not only use this disagreement to identify the most relevant configurations to build up the model's training set in an active learning procedure but also monitor and bias it during simulations to control the generalization error. This facilitates the adaptive development of committee neural network potentials and their training sets while keeping the number of *ab initio* calculations to a minimum. To illustrate the benefits of this methodology, we apply it to the development of a committee model for water in the condensed phase. Starting from a single reference *ab initio* simulation, we use active learning to expand into new state points and to describe the quantum nature of the nuclei. The final model, trained on 814 reference calculations, yields excellent results under a range of conditions, from liquid water at ambient and elevated temperatures and pressures to different phases of ice, and the air–water interface—all including nuclear quantum effects. This approach to committee models will enable the systematic development of robust machine learning models for a broad range of systems.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0016004>

I. INTRODUCTION

Machine learning has emerged in recent years as a powerful tool for the description of complex chemical systems.^{1–3} A major contribution has been the development of machine learning potentials (MLPs)—models that represent potential energy surfaces created by explicit *ab initio* calculations—which enables the study of chemical systems for long timescales and on large length scales, even with chemical reactivity included. The first method based on artificial neural networks that is, in principle, scalable to arbitrary system sizes was the high-dimensional neural network potential (NNP) methodology^{4,7} combined with atom-centered symmetry functions to describe atomic environments.⁸ Over the years, many other distinct methods have been proposed for this difficult task based on a

range of descriptors and either on artificial neural networks^{9–16} or on kernels.^{17–22} Since their introduction, NNPs have been successfully applied to solvents,^{23–25} solids,^{12,26,27} solid–liquid interfaces,²⁸ and reactive processes in solution²⁹ or at interfaces,^{30,31} and have therefore repeatedly demonstrated their reliability for the understanding of complex molecular systems and materials. However, a crucial component of any MLP is a robust and representative training set whose construction can easily become the most challenging part of the development of such a model, especially for condensed phase systems.

At the same time, it is well known in the machine learning community that the predictive power of a machine learning approach can be substantially improved by combining multiple individual models.^{32–36} Instead of a single model, multiple models are trained

31 August 2023 13:09:24

independently to form a committee that offers several benefits. Averaging over the predictions of an ensemble of committee members usually provides an improved accuracy of the prediction compared to the individual members.^{33,37–39} In addition, the disagreement of the committee, as measured by the standard deviation of the predictions of the members, provides access to an estimate of the generalization error.^{34,40,41} Moreover, the committee model can substantially reduce overfitting issues.⁴² Finally, by adding previously unlabeled data with maximal committee disagreement to the training set, the model can be systematically improved—an active learning strategy known as query by committee (QbC).^{34,43}

Despite the rise of machine learning in molecular simulations and materials science, committee models are not considered standard tools and have mainly been used in pioneering work. This includes the well-established practice to use the difference between two NNP models for the manual improvement of the training set,^{23,44,45} as first described in Ref. 46, but without combining the predictions of the two models. More recently, the use of two machine learning models has been proposed for the simulation of the infrared spectra of gas-phase molecules either in an ensemble averaging approach⁴⁷ or for the error estimation in post-processing.⁴⁸ Additionally, the disagreement of NNPs has been shown to be crucial for the automated fitting of NNPs at coupled cluster accuracy for protonated water clusters.⁴⁹ Ensemble methods were also recently proposed for uncertainty estimation in chemical machine learning.⁵⁰ Besides these examples, QbC strategies have been leveraged for the development of moment tensor potentials⁵¹ and more recently for deep potential models.⁵² QbC has also been successfully utilized for active learning in chemical space,⁵³ which additionally demonstrated the improvements in accuracy obtained by using the committee average for predictions. In the realm of Gaussian approximation potentials,¹⁷ data-driven learning strategies⁵⁴ have recently been shown to be crucial for the automated development of machine learning potentials. Finally, complementary strategies such as farthest point sampling⁵⁵ have been tested for the construction of uniform datasets.⁵⁶

In this work, we exploit the established benefits of committee models known in the field of machine learning to create robust MLPs with controlled generalization errors in an automated fashion. Conceptually, our approach is based on a combination of multiple models of the well-established NNP formalism.^{6,7} The resulting committee intentionally shares the same atom-centered symmetry functions⁸ as descriptors for the atomic environments, thus leading to a small, often negligible, computational overhead in production runs. We show that, compared to the individual NNPs, this approach results in improved accuracy for predictions and at the same time gives direct access to the committee disagreement, an estimate of the generalization error. This disagreement—being straightforward to compute during a simulation—can thus be monitored and even biased to stabilize the simulation. It also enables active learning via QbC techniques, which expands the training set, systematically improving the model. These benefits allow us to build an adaptive workflow for the development of committee NNP models and the systematic generation of their training sets. We finally illustrate the capabilities of the resulting methodology on the description of water in the condensed phase at various state points and make the resulting training set and model parameters available to the community. This methodology integrates the aspects of committee models

with new features such as our adaptive active learning workflow and the biasing of the committee disagreement into a unified framework and will yield robust MLPs for a broad range of systems.

II. COMMITTEE NEURAL NETWORK POTENTIALS

Before we introduce the principal ideas underlying committee NNPs (C-NNPs), we first briefly summarize the original Behler–Parrinello NNP methodology. To represent an interatomic potential by NNPs, the atomistic structure is first transformed using atom-centered symmetry functions⁸ into translationally and rotationally invariant descriptors of atomic environments. These serve as the input for atomic neural networks that output the auxiliary components of the total potential energy, which is then obtained as a sum of contributions from all atoms in the system. The resulting permutationally invariant structure–energy relation can be analytically differentiated to obtain forces, for example, to drive molecular dynamics, and is scalable to essentially arbitrary system sizes.⁶ The whole model is trained by optimizing the parameters (weights and biases) of the atomic neural networks, one per element, to reproduce the reference energies and optionally forces of a training set. In contrast, the network architecture and the particular choice of symmetry functions are hyperparameters that need to be specified by the user. For further details on the original NNP methodology, we refer the reader to Ref. 7.

In order to extend this methodology to a committee model, we propose to use multiple NNPs that have been optimized independently using the same training set. However, every individual NNP is trained to a slightly different subset of the full training set, while a small fraction is intentionally left out in each case. This strategy, also known as random subsampling in the machine learning community, introduces variation between the committee members as shown, for example, in Ref. 34 for artificial neural networks. Together with the intrinsic stochastic nature of the neural network optimization due to the initialization of the weights and the optimizer itself, these different contributing factors provide a sufficiently diverse committee of NNPs. Given the predictions of potential energies and atomic forces by the committee of NNPs, $\{E_i(q)\}_{i=1}^n$ and $\{-\nabla_\alpha E_i(q)\}_{i=1}^n$, as a function of the positions of all the atoms $q \equiv \{\mathbf{q}_\alpha\}_{\alpha=1}^N$, the C-NNP prediction for any given structure is obtained as an average,

$$E(q) = \frac{1}{n} \sum_{i=1}^n E_i(q), \quad (1)$$

$$\mathbf{F}_\alpha(q) = \frac{1}{n} \sum_{i=1}^n \mathbf{F}_{i\alpha}(q) = -\frac{1}{n} \sum_{i=1}^n \nabla_\alpha E_i(q),$$

where i -indexed quantities represent the predictions of the n individual committee members, the non-indexed ones represent the averaged C-NNP prediction, and α is the atomic index. As we have a set of predictions for each structure, we can define the committee disagreement as the standard deviation of the individual NNPs,

$$\sigma_E(q) = \left[\frac{1}{n} \sum_{i=1}^n (\Delta E_i)^2 \right]^{\frac{1}{2}}, \quad (2)$$

$$\sigma_{\mathbf{F}_\alpha}(q) = \left[\frac{1}{n} \sum_{i=1}^n (\nabla_\alpha \Delta E_i)^2 \right]^{\frac{1}{2}},$$

where we introduce the notation $\Delta E_i \equiv E - E_i$. These disagreements can be easily computed and monitored on-the-fly during a production run and provide an estimate of the uncertainty of the C-NNP prediction for a given configuration. The energy disagreement gives global information, while the force disagreement is locally resolved for each atom and can therefore highlight the weaknesses of the prediction for a local environment within a given configuration. Thus, access to the disagreement enables direct validation of the predictions of a C-NNP model, in particular, since it is known that the committee disagreement provides a measure of the generalization error.^{34,40,41}

At this point, it is clear that due to the correlation of disagreement and generalization error, it is beneficial to have small disagreement during a production run. This will be the case for a well-trained robust committee model, but before we obtain one, we can take steps to ensure that the disagreement is controlled. To achieve that, we define a biasing potential $E^{(b)}$ that acts on the energy disagreement, for example, using a shifted harmonic form

$$E^{(b)}[\sigma_E(q)] = \theta(\sigma_E - \sigma_0) \frac{1}{2} k^{(b)} (\sigma_E - \sigma_0)^2, \quad (3)$$

where θ is the Heaviside step function that activates the bias only upon reaching a threshold disagreement σ_0 . In principle, other functional forms of the biasing potential are possible, which will be explored in future work. The above choice makes it particularly easy to compute the associated biasing forces as

$$-\nabla_\alpha E^{(b)} = \theta(\sigma_E - \sigma_0) k^{(b)} \frac{\sigma_E - \sigma_0}{\sigma_E} \cdot \frac{1}{n} \sum_{i=1}^n -\Delta E_i \nabla_\alpha \Delta E_i, \quad (4)$$

which can be used to keep the disagreement within reasonable upper limits in a molecular dynamics run. Biasing of the committee disagreement, therefore, provides a unique way to stabilize a simulation that employs a committee model. By shifting the onset of the biasing potential to larger committee disagreements, the influence on the simulation can be fine tuned and minimized so that the biasing potential only acts as a safeguard against rare excursions of very high disagreement. Biasing the energy disagreement in this way allows the system to move freely in parts of configuration space that are well-described by the C-NNP while effectively erecting a barrier at the boundary of this region, which prevents the simulation from entering configurations with high generalization errors. As an alternative that is more local but potentially also more invasive, separate biases can be introduced on individual atomic force disagreements, as we detail in the Appendix. We also note in passing that approaches to sample *intermediate* disagreement, in the spirit of various enhanced sampling techniques, could provide a new direction to efficiently generate relevant structures to be included in training sets of MLPs as part of an active learning procedure.

In the present case, we intentionally decided to share the same set of symmetry functions for the representation of atomic environments between the C-NNP members. This has the advantage that the evaluation of the symmetry functions and their derivatives is only performed once for the whole committee, which is typically the computationally most demanding step. Then, only the atomic

neural networks are evaluated separately for each committee member, incurring only a small overhead compared to using a single NNP.

To highlight the benefits of the committee NNP approach, we illustrate some of its main features in Fig. 1. A crucial step in the development of any machine learning potential is the preparation of the training set, which we address in detail in Sec. III. The training set needs to be representative of the planned simulations and dense enough to generate reliable interpolation between the training points. Most preparations of training sets, therefore, start from simulations with the chosen reference method, typically in the spirit of *ab initio* molecular dynamics (AIMD).³⁷ A first training set can then be generated by choosing uncorrelated configurations from such a

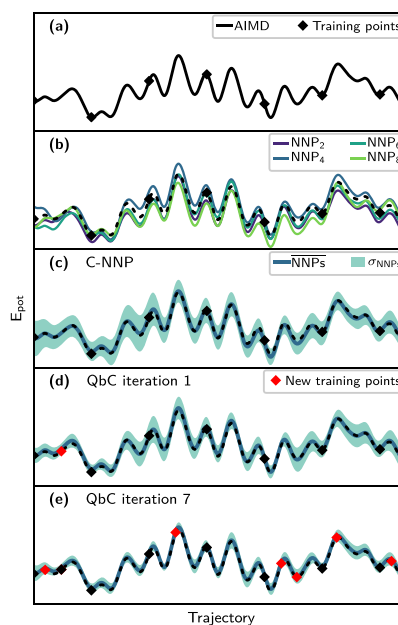


FIG. 1. Illustration of the committee model compared to the individual NNP members of the committee and of the QbC procedure for 64 water molecules in the liquid phase. (a) *Ab initio* reference trajectory and first selection of training points equally spaced along the trajectory. (b) Predicted energy along the original trajectory from eight independent NNP fits to the same training points. (c) Predicted energy along the original trajectory of the committee model composed of the eight NNPs. (d) Active improvement of the model via query by committee. The additional new training point is highlighted in red. (e) Performance of the committee model after seven query by committee iterations. The committee disagreement has been reduced, and the prediction has improved.

31 August 2023 13:09:24

trajectory, as shown in Fig. 1(a). Training multiple NNPs with different initial conditions or to different subsets of the full training set leads to varying performance between them, as highlighted in panel (b). In previous work, the user would then select the best performing NNP as the model of choice. However, if the different models are combined to a committee NNP, the prediction is substantially improved, as shown in panel (c). At the same time, the committee disagreement allows the identification of configurations for which the uncertainty of the model is high, enabling active learning strategies based on QbC techniques to iteratively improve the model. As shown in panels (d) and (e), adding selected configurations to the training set substantially reduces the committee's disagreement while improving its prediction compared to the reference data. To select new configurations for the training set, it is possible to use either the global total energy disagreement or the local information contained in the atomic force disagreement, after a suitable reduction over all atoms in the given frame. Overall, these features allow for a data-driven approach to developing C-NNP models, as presented in detail in Sec. III.

The methodology to make use of the benefits of committee models has been implemented in the CP2K simulation package³⁸ for Behler–Parrinello NNPs and will be made available in the next release. This includes the on-the-fly evaluation of the energy and force disagreement and the associated biasing of the energy disagreement. In practice, a committee NNP model can be obtained by performing individual fits with any NNP training code, for example, with the open-source n2p2 code²⁷ or the RuNNer code.³⁹ One can therefore see that the proposed concepts are straightforward to adapt for a broad range of existing MLPs while introducing benefits and additional features.

III. ACTIVE LEARNING PROCEDURE FOR COMMITTEE NEURAL NETWORK POTENTIALS

Let us now address a crucial step in the development of any MLP, the preparation of the training set. A machine learning model can only be as good as its underlying data, which needs to be representative of the situations encountered when using the final model. As discussed in the Introduction, the selection of configurations for the training of machine learning potentials has recently seen great progress toward data-driven and automated approaches.^{21,43,47,49,51–54,60,61} In a similar spirit, here, we present an adaptive active learning workflow for the construction of robust C-NNPs for classical and path integral molecular simulations. The approach developed here builds on the automated fitting of NNPs at the coupled cluster level of theory for gas-phase clusters.^{19,62}

We make use of two basic properties of the committee model to automate the development of C-NNPs. First, as shown in Sec. II, the committee disagreement can be used as an estimate of the generalization error of the model. By adding configurations to the training points that feature the highest committee disagreement, the most important points for an improvement of the model can be iteratively selected. This is the main principle behind active learning via QbC.⁴³ Second, the C-NNP is many orders of magnitude cheaper than the reference electronic structure method, and new configurations can therefore be generated rapidly using the C-NNP. These large sets of configurations can then be efficiently screened using

QbC, and expensive reference calculations are only performed for these selected points.

We organize the active learning workflow into different generations, each of them comprising multiple QbC cycles and other operations, as outlined in Fig. 2 for the condensed phase of water. Each generation includes new state points and yields a C-NNP that will be used to generate new candidate structures for the next generation. The on-the-fly monitoring and biasing of the committee disagreement provide invaluable tools to guarantee the stability of these simulations and the validity of the new configurations. Only at the beginning of the first generation, the process is seeded from an AIMD simulation in order to provide an initial set of structures. If the new conditions are not structurally drastically different from those in the previous generation and we use disagreement biasing to keep molecular dynamics stable, we can start with a single state point and gradually expand into new regions without the need to run additional expensive AIMD simulations. If the final model should be applied together with a quantum description of the nuclei, this can also be adaptively included over the generations by gradually increasing the quantum character of the nuclei in imaginary time path integral simulations.^{63,64}

Within a generation, QbC is used to adaptively extend the training set by selecting the most representative configurations separately for each state point, improving its description, as schematically shown in the bottom left panel of Fig. 2. With this procedure, multiple state points can easily be treated in parallel. At the very beginning of each QbC cycle, a small number of random configurations are chosen to train the first committee, while in subsequent iterations, new configurations are selected based on the highest committee disagreement. We chose to use the force disagreement (rather than the total energy disagreement) for this selection, since it is sensitive to the local environments within a configuration and insensitive to the global offset of the whole potential energy surface. Convergence of these individual QbC cycles can be detected by monitoring this disagreement. If the structures were generated by AIMD simulations, as is the case at the beginning of the first generation, the associated reference forces and energies are already known and the improved C-NNP model can be trained directly to the growing training set. In subsequent generations, candidate structures are generated by the molecular dynamics of the previous generation's C-NNP, and explicit electronic structure reference calculations are only needed for the small number of actively selected points. Once all QbC cycles for the selected conditions in a given generation are converged, the individual training sets are combined and a final tight optimization of that generation's resulting C-NNP is performed.

The adaptive improvement of the model and its training set is completed after several generations, when all desired conditions have been included and the final C-NNP exhibits the required accuracy in subsequent production simulations. In these simulations, the committee disagreement on energy and forces can be monitored on the fly and compared to the disagreement known from the active learning process. If this disagreement stays within the range encountered for these known conditions included in the training process, it is expected that the final model reaches the desired accuracy also in the production simulations. Thus, utilizing the properties of committee models, the data-driven workflow outlined above helps automate the development of robust machine learning

31 August 2023 13:09:24

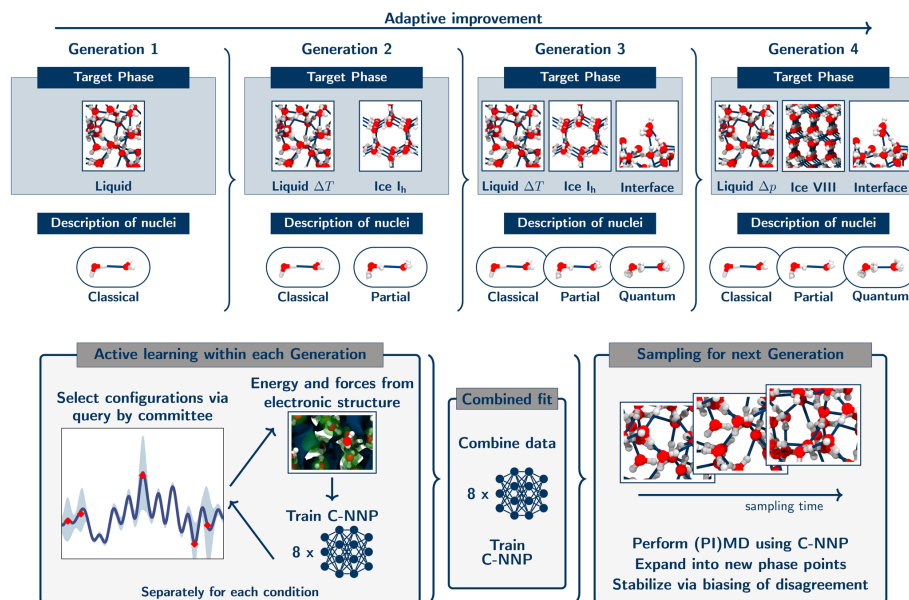


FIG. 2. Illustration of the adaptive improvement of the committee NNP over multiple generations. The top panel summarizes the expansion into new target phases and the iterative improvement of the description of the nuclei in each generation. Within each generation, the most important points for an improvement of the model are actively selected using QbC based on the highest committee disagreement, separately for each selected state point (bottom left). Afterward, the reference energy and forces, if previously unknown for these structures, are obtained from explicit electronic structure calculations. These points are added to the training set, and the committee members are trained to the expanded training set. QbC iterations are repeated until the committee disagreement converges (see the text for details). At the end of each generation, all training points are gathered in order to perform a final extended fit of the committee model (bottom middle). The resulting C-NNPs can consecutively be applied for exhaustive (PI)MD sampling at various new state points (bottom right). These simulations provide the structures for the next generation in the adaptive improvement.

31 August 2023 13:09:24

potentials and subsequent production simulations with controlled accuracy.

IV. APPLICATION OF COMMITTEE NEURAL NETWORK POTENTIALS TO WATER

A. Development of the committee model

In order to showcase the benefits of the committee NNP methodology, we develop a C-NNP model for water at various state points, including also the quantum nature of the nuclei, following the data-driven workflow described above. All specific settings used here are listed in Sec. VI, while the training input files, training set, and parameters of the final model are publicly available at <http://doi.org/10.5281/zenodo.4004590>. In the first active learning generation, we seed the procedure with 300 ps of classical AIMD simulation of liquid water (LW) at 300 K obtained at the hybrid

density functional theory (DFT) level⁶⁵ and perform a single QbC cycle targeting this state point. This QbC cycle uses a committee of 8 NNPs and is initialized with 20 structures randomly selected from the ensemble. Ten new configurations with the highest disagreement are added in each subsequent iteration. After the training of the individual members, the energies and forces of 5000 random structures from the original trajectory are predicted in order to compute the committee disagreement. We intentionally use only a subset of the large pool of candidate structures in order to make the QbC iterations computationally more efficient. For the same reason, the QbC NNPs are optimized relatively loosely (15 epochs) within each QbC iteration. In order to select the most relevant configurations for an improvement of the model, we chose to use the mean force disagreement of each configuration to rank the candidate structures. If a newly selected configuration has already been included in a previous QbC iteration, it is not added again to the training set. Such occasions indicate that the QbC process is reaching the limits of

the provided set of configurations, since structures are selected more than once. Given that the DFT energies and forces are already known in the first generation, the selected points are directly added to the training set.

Monitoring of the committee disagreement during a QbC cycle allows the user to easily gauge the convergence of the process. The evolution of the atomic force disagreement during the first QbC process is shown in the top panel of Fig. 3 separately for the structures in the training set, the newly selected structures, and the 5000 candidates, from which the next ten structures for the training set are chosen. At the beginning of the QbC process, the newly selected points feature substantially larger disagreement compared to the large set of candidate structures and the training set. As more and more points with the highest disagreement are added to the training set, the disagreement of all three sets of structures decreases monotonically. However, the disagreement of the selected points decreases faster and approaches that of the training set and the candidate structures, indicating that the newly selected points are not adding further value for an improvement of the model anymore. The disagreement for the training set and the set of candidates is similar, only slightly higher for the training structures for most of the process, which shows that the training set picks up the outliers of the ensemble, but without substantially deteriorating the quality of the model. The force disagreement for all considered sets of structures

eventually decreases more slowly, indicating that the active learning process is well converged after roughly 100 structures have been added to the training set.

Let us next focus on the actual performance of the C-NNP model for water at the chosen starting condition. As mentioned previously, the committee disagreement is an estimate of the generalization error, and so we should expect the accuracy of the model to improve over the QbC process as the disagreement decreases. In order to validate this expectation for the C-NNP approach, the evolution of the force root mean square error (RMSE) along the QbC cycle for an independently generated test set at the chosen condition is shown in the bottom panel of Fig. 3. In addition to the RMSE of the full committee with eight members, we also include the performance of all possible committees with four, two, and one members (i.e., individual NNPs) for comparison. As anticipated from the evolution of the committee disagreement, the force RMSE of the full eight-member C-NNP starts at roughly 60 meV/Å at the beginning of the QbC process and converges monotonically to a value of about 40 meV/Å after roughly 100 points have been added to the training set. At the same time, the performance of the smaller committees, and most notably the individual NNPs, is substantially worse, especially at the beginning, where the individual NNPs show an RMSE that is twice as large as that of the full C-NNP. Although the large initial differences decrease as the QbC process progresses, the difference remains clear even when convergence has been reached with roughly 100 training points, where the committee still outperforms the individual members and reduces the RMSE from an average of 48 meV/Å for the individual NNPs to 42 meV/Å for the full eight-member C-NNP. Given the slower convergence with an increasing number of structures, it is clear that it would take a much larger training set for the individual NNPs to reach the performance of the C-NNP. Thus, this analysis highlights the added accuracy of the committee approach, known from other machine learning applications.^{33,37–39}

Overall, this detailed analysis of the first QbC cycle shows that only a relatively small number of points are needed to reach convergence for the starting point of our active learning procedure. The 111 structures identified after the first 10 QbC iterations are therefore used as the final training set of the first generation C-NNP model. After stringent re-optimization of the individual NNPs—see Sec. VI for details—the C-NNP model is ready to be used for the generation of new structures at state points neighboring to the original ensemble of liquid water at 300 K. For these simulations, the on-the-fly computation of the committee disagreement is crucial in order to judge if the new configurations are physically meaningful. In addition, the biasing of the committee energy disagreement derived above can be used to prevent the system from entering regions of configuration space where the model is not well determined by the training set.

To illustrate the benefits of this feature, we used the C-NNP model of generation 1 for the simulation of the air–water interface at 300 K. In Fig. 4, we show the resulting probability distributions of the total potential energy and atomic force disagreement. The distributions from unbiased simulations feature a very long tail for the energy disagreement, which is to be expected from a model that has not been trained on gas-phase clusters or interfaces. The interface is confirmed as the culprit by inspecting the spatial distribution of the disagreement, as shown in the inset in the bottom panel of Fig. 4,

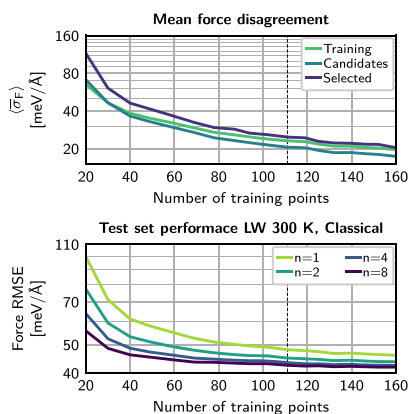


FIG. 3. Convergence of the QbC process in the first generation with respect to the number of structures in the training set. The top panel shows, in the logarithmic scale, the mean force committee disagreement averaged over the given set of structures ($\overline{\sigma_F}$) for the training set (Training), the large set of potential new candidates (Candidates), and the actual newly selected configurations (Selected). The bottom panel shows, in the logarithmic scale, the force root mean square error (RMSE) of the C-NNP model for different committee sizes from one to eight members. It was evaluated on 500 independently generated configurations for the target state point of liquid water at 300 K with classical nuclei (see Sec. VI for details). The final number of training points (111) used from the QbC cycle is marked with a vertical dashed line.

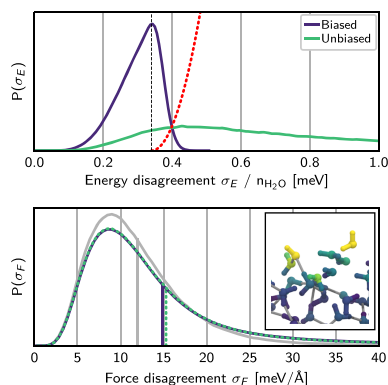


FIG. 4. Comparison of the distribution of the committee disagreement with and without biasing. The plots show the normalized probability densities of the energy (top panel) and force (bottom panel) committee disagreement for a water slab with 216 water molecules at 300 K. Two simulations were performed for the generation 1 C-NNP model—with and without applying a biasing potential acting on the energy disagreement. The offset of the biasing potential is chosen such that the bias only acts on configurations with an energy disagreement per water molecule larger than 0.33 meV per molecule. The resulting functional form is included in the top panel as a red dotted line. The force disagreement distribution obtained for the training set of the generation 1 C-NNP model is shown in gray in the bottom panel, and the averages of the respective distributions are marked as horizontal lines. The inset in the bottom panel shows a snapshot of the air–water interface with atoms color-coded by their respective force disagreement, where yellow indicates high and purple indicates low disagreement.

where individual atoms are colored by their value of the atomic force disagreement. Indeed, the highest values are found for atoms at the interface whose environments deviate from those in the bulk liquid. Compared to the force disagreement of the same model for its training set (gray distribution in the bottom panel of Fig. 4), the distribution from the slab simulation remains close but exhibits a heavier tail due to the interfacial atoms. The application of a bias to the energy disagreement suppresses the tail of its distribution and yields a more compact distribution. In contrast, it has only a very subtle effect on the distribution of force disagreement, which highlights the relatively mild influence of the biasing potential on the local behavior of the system. Therefore, the energy disagreement biasing can be understood as a global safeguard that prevents the system from moving into regions of configuration space with large disagreement while keeping local perturbations low. In light of this analysis, we chose to use a weak biasing potential for all simulations used to generate configurations under new conditions, as detailed in Sec. VI, to ensure the stability of the simulations without substantial distortion of the structures.

In order to select the target state points for the next generation in the adaptive improvement of the C-NNP model, we performed test simulations under a variety of conditions with the generation

1 model. After a careful analysis of the observed disagreement for these simulations, we chose hexagonal ice at 250 K and liquid water at increased temperatures up to 400 K as the targets for the second generation of the active learning process. Moreover, the quantum character of the nuclei is targeted by separate path integral molecular dynamics (PIMD) simulations for the same state points. To introduce quantum delocalization gradually, we use underconverged path integral discretization to stay structurally closer to the classical ensemble of generation 1. We employ a separate QbC process to select new structures for each condition to ensure that they are optimally covered by the training set independently of the others. This has the additional advantage of increased computational efficiency, as these QbC cycles can easily be run in parallel. In contrast to the first generation, the QbC iterations are seeded by choosing 20 random structures from the training set of the previous generation. The DFT reference energy and forces are unknown for the newly selected configurations and, thus, are computed during the QbC cycle. Good convergence of these QbC processes is reached after only 3–5 iterations, and generation 2, therefore, adds a total of roughly 250 new configurations to the combined training set from the eight independent QbC cycles.

In the final two generations of the active learning workflow, additional conditions were targeted. This includes high pressure liquid water, a water slab to represent the air–water interface,⁶⁶ and finally the high pressure ice phase VIII, as well as the quantum nature of the nuclei, as summarized in the top panel of Fig. 2. All details on these target conditions and the relevant simulations can be found in Sec. VI. During generation 3 and 4, 240 and 205 additional reference configurations at the various state points were added to the training set, respectively. The final result after four generations of our active learning procedure is a training set of 814 structures and the corresponding tightly optimized C-NNP, which are able to describe a broad range of conditions with classical or quantum nuclei. The whole process was originally initialized from a single classical AIMD simulation at 300 K, with no *ab initio* path integral molecular dynamics (AIPIMD) required at any point. This process could be continued to expand into additional thermodynamic regions in case they are of interest for specific scientific questions. However, we consider the diversity of the training set sufficient to showcase the ability of our approach to generate robust and accurate C-NNP models and their training sets in a data-driven and automated fashion.

B. Validation of the committee model

After presenting the details of the active learning procedure for the development of a C-NNP model for water at various state points including quantum nuclei, we analyze the improvement of the model over the different active learning generations and validate, in particular, the quality of the final generation 4 model. For this purpose, we explicitly benchmark static and dynamical thermal properties against AIMD and AIPIMD simulations available for a single state point (liquid water at 300 K) with and without nuclear quantum effects, while we compare RMSE values for all the state points considered. The RMSE analysis is performed for an independently created test set, which spans the same thermodynamic state points as targeted during the development of the model, but has been generated by separate simulations with the final generation 4 model,

31 August 2023 13:09:24




FIG. 5. Force root mean square error (RMSE) of the C-NNP models over the four generations of our active learning workflow. For each generation, we show the force RMSE relative to the revPBE0-D3 reference evaluated over an independently generated test set. The RMSE is averaged over all the state points in each case, separately for classical and quantum structures.

as detailed in Sec. VI. This test set comprises a total of 8000 configurations, split equally between classical and quantum configurations, for which DFT reference energies and forces have been calculated. It, therefore, features one order of magnitude more configurations than the final training set of generation 4 and enables a comprehensive performance analysis of the improvement of the model for the various conditions.

To summarize the progress over the different generations, we report in Fig. 5 the force RMSE for the independently generated

test set averaged over the various state points but separate for quantum and classical structures. Classical structures show only a minor improvement across the generations, as they are described already very well by the generation 1 model. On the other hand, quantum structures, which were not included at all in the training set of generation 1, improve substantially with each generation. The final C-NNP model of generation 4 reproduces the forces in the independent test set with a RMSE of 40 meV/Å and 52 meV/Å for classical and quantum structures, respectively. This is overall slightly smaller compared to previous work employing NNPs for water with a classical description of the nuclei^{23,24} and substantially smaller compared to a model including quantum nuclei²⁵ while covering diverse regions of the phase diagram and including the quantum nature of the nuclei with a training set of just ~800 configurations. We believe that the surprising robustness of the model with such a small training set is mainly due to a combination of two factors. The active learning process selects structures that are the most important for the improvement of the model while at the same time avoiding extreme structures that would distort the fit in regions relevant for simulations. Furthermore, the fact that the C-NNP is an average of multiple models lends it stability that is not available to the individual members.

To validate the performance of the final generation 4 C-NNP model for the calculation of equilibrium properties at a finite temperature, we compare against reference AIMD and AIPMD trajectories of liquid water at 300 K.⁶⁵ We used the final C-NNP

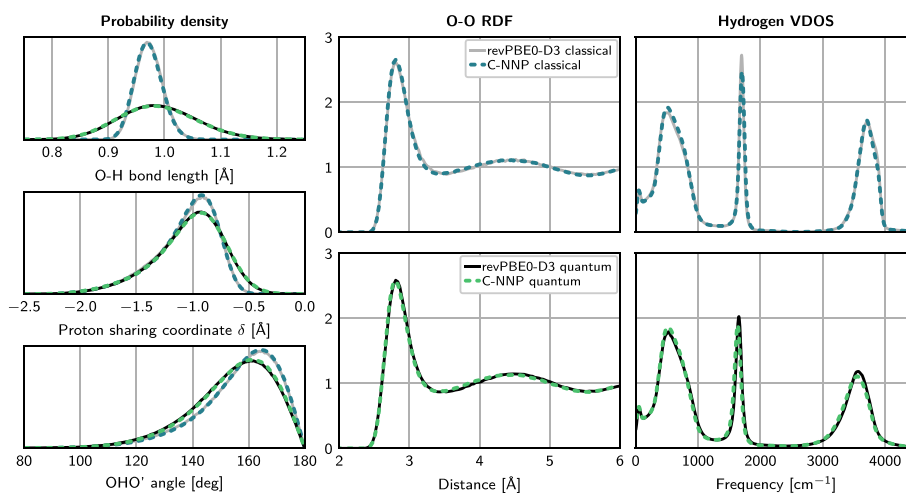


FIG. 6. Comparison of several local and global static properties as well as the vibrational density of states obtained by explicit revPBE0-D3 *ab initio* simulations and the final C-NNP generation 4 model. The three panels on the left show the normalized probability density for the O-H bond lengths (top), the proton-sharing coordinate δ (middle), and the hydrogen bond angle (bottom) for both a classical and quantum description of the nuclei. The two panels in the middle display the comparison of the O-O radial distribution functions, while the two panels on the right compare the hydrogen atom vibrational density of states for a classical (top) and a quantum (bottom) description of the nuclei.

model to obtain two sets of NVT production trajectories, one with classical nuclei and one with quantum nuclei. A comparison of both static and dynamical properties of the system is shown in Fig. 6. Of the static properties, we focus first on the local structure of individual molecules (O–H bond lengths) and their hydrogen bonds (proton-sharing coordinate δ and hydrogen bond angle). The C-NNP model exhibits excellent agreement with the reference AIMD data for these properties in both the classical and quantum cases. This includes the tail of the quantum distribution of δ that corresponds to strong proton sharing, as shown in Fig. S4. We then characterize the intermolecular structure using radial distribution functions (RDFs), which are, again, captured accurately in both cases. There is a negligible shift of the second peak of the O–O RDF to shorter distances in the quantum case (Fig. 6), while the O–H and H–H RDFs in Figs. S1 and S2 show essentially perfect agreement. Likewise, the vibrational dynamics of the system, encoded in the vibrational density of states, is reproduced reliably by the C-NNP model, as seen for hydrogen atoms in the right column of Fig. 6 and in Fig. S3 on a logarithmic scale. In the classical spectrum, the bending peak at $\sim 1700\text{ cm}^{-1}$ is ever so slightly broadened, resulting in a small decrease in the peak height, while the rest of the spectrum is essentially a perfect match, including low-intensity features only visible on the logarithmic scale. In the thermostatted ring polymer molecular dynamics (TRPMD) spectrum, the C-NNP model exhibits a very small red shift in the bending and stretching regions but shows a very good match of the spectrum as a whole. Note that a contributing factor to the already very small remaining differences in the quantum case is the incomplete convergence of ring polymer contraction in the reference AIPIMD simulations, as suggested by the comparisons of contraction to 1 and 4 replicas in the supplementary material of Ref. 65. Overall, we can see that our C-NNP model matches the reference *ab initio* method accurately in the description of the structure and vibrational dynamics of liquid water. This is particularly remarkable in the quantum case, as no explicit AIPIMD simulations were required in the parameterization of the model.

In the next step, we again widen the analysis of the performance of the developed C-NNP model for water. As just demonstrated, it is certainly possible to explicitly validate various properties with respect to AIMD or even AIPIMD simulations under selected conditions. However, this type of analysis quickly gets out of scope for all targeted state points that we considered in the present case. In addition, the whole purpose of the development of MLPs is usually to replace expensive *ab initio* sampling. We therefore come back to a detailed analysis of the RMSEs for energies and forces for the different conditions accessible in our independent test set. As shown in Fig. 7, the final model scores overall quite well for all considered conditions, usually with slightly larger RMSE values for quantum structures. The water slab performs the worst in this analysis, indicating the more complex nature of these configurations due to the presence of an interface, but still features RMSE values similar to or lower than those reported for previous work on water.^{23–25} At the same time, the two phases of ice are reproduced best, as expected from the simpler nature of these systems based on the arrangement of the molecules on a lattice. Although nuclear quantum effects lead to substantially larger and broader potential energy distributions, the performance of the final model is convincingly good across the board. Comparing the score of liquid water for classical and quantum nuclei at 300 K to the other conditions makes us confident that also the associated

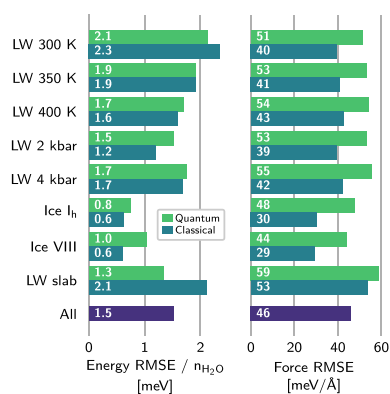


FIG. 7. Root mean square error (RMSE) comparison for the final generation 4 C-NNP model. The bar chart shows RMSEs of energies and forces for the C-NNP model with respect to the revPBE0-D3 reference for an independently generated test set. The independent test set consists of 8000 uncorrelated configurations covering the thermodynamic conditions targeted during the development of the C-NNP model and is decomposed into the individual conditions of liquid water (LW) at different temperatures and pressures, hexagonal ice I_h and ice VIII, as well as the air–water interface (LW slab), separately for a classical and quantum description of the nuclei (see Sec. VI for details on the creation of the independent test set).

properties are reproduced with similarly convincing agreement to explicit *ab initio* sampling.

Having seen the performance of the final model, it is worthwhile to step back and see how it emerges from the active learning process. To do that, we repeat the same validation procedure as above for the model from generation 1, i.e., a C-NNP trained on 111 structures selected from classical AIMD trajectories of liquid water at 300 K using QbC. The results of comparison to explicit AIMD and AIPIMD for equilibrium properties shown in Figs. S6 and S10 clearly demonstrate that even this model does very well in simulations. For classical nuclei, it is, in fact, as good as the more extensively trained final model. Remarkably, it performs well even in the quantum case, despite the fact that no path integral structures at all were used in its training. The only noticeable deviations are the red shift in the infrared region of the vibrational spectrum and the increased proton sharing, both only slightly larger than for the final model. We have already seen in Fig. 3 how the RMSE against revPBE0-D3 energies and forces evaluated on the independent test set for liquid water at 300 K converges during the QbC process that generates the training set for this model. When we look at the RMSEs also for other conditions in Fig. S5, it is clear that the generation 1 model performs worse than the final one, especially for quantum structures, which are entirely absent from the training set. However, forces for classical structures are at the level of the generation 4 C-NNP, with the exception of ice VIII and the air–water interface. These states are not only absent from the small training set of the generation 1 model but also

structurally substantially different from liquid water, and so it should be expected that they are represented less accurately. This combination of performance characteristics offers an opportunity to gauge the meaning, or possibly the limitations, of RMSEs when it comes to molecular dynamics. While the force RMSEs are roughly three times as high for quantum structures as they are for the classical ones for the generation 1 model, these forces are sufficient to generate path integral trajectories that match the reference almost as well as in the classical case. In terms of computational cost, the ten iterations of a single QbC process selecting structures from a pre-existing AIMD trajectory took 6.7 h on a single AMD EPYC computational node with 32 cores. Given that this C-NNP is rather easy and computationally cheap to create, this seems to offer an efficient way to extend existing *ab initio* simulations using an MLP from which we do not expect broad generality.

V. CONCLUSIONS

In summary, we have shown how committee models can be exploited in the context of machine learning potentials, which we demonstrated specifically using Behler–Parrinello neural network potentials, obtaining a compact training set and a robust C-NNP model with a range of advantages. While the committee prediction outperforms its individual member NNPs, the committee disagreement provides a way to monitor and control the accuracy of the model relative to its parent *ab initio* method and its training set. Using a bias of the committee disagreement that we introduced, C-NNP simulations can be stabilized by effectively erecting a barrier at the boundary of regions of configuration space described well by the training set, thus preventing such simulations from entering regions of high uncertainty and high generalization errors. The computational overhead of our approach is low compared to using a single NNP, as we intentionally share descriptors of atomic environments between the committee members, which means only the atomic neural networks need to be evaluated multiple times.

Making use of the committee disagreement and repeated query by committee processes, we designed an active learning procedure capable of generating training sets in a largely automated and data-driven fashion while also keeping the number of required reference *ab initio* calculations to a minimum. To demonstrate the benefits of this methodology, we developed a C-NNP machine learning potential—made freely available—for water under a variety of conditions with both a classical and a quantum description of the nuclei. Even though its training set comprises only ~800 configurations, the final model shows excellent agreement with the DFT reference method in comparison of energies and forces as well as in comparison of the properties calculated in classical and quantum molecular dynamics simulations.

The promising results presented in this work make us confident that the methodology introduced here can be used to systematically develop robust and general C-NNP models applicable across broad regions of the phase diagram and under a variety of conditions for systems of increasing complexity. Given the remarkable performance of the first generation C-NNP model trained on a small number of samples from a single AIMD simulation, we see great potential in the development of simpler C-NNP models for the direct extension of *ab initio* simulations for a given state point.

Thanks to the much smaller size of the resulting training set compared to previous work, it will be possible to use more advanced reference *ab initio* methods in the condensed phase. Finally, we expect that the described concepts can be applied directly to other MLPs based on artificial neural networks and can also be transferred to kernel-based MLPs after the introduction of stochastic elements into their training process. We hope that, thanks to the low additional complexity and required effort, the committee-based models can become a routine component of the development of machine learning potentials.

VI. COMPUTATIONAL DETAILS

We implemented the active learning workflow in Python, interleaving data manipulation and execution of various simulation packages to perform specific tasks, such as the optimization of individual NNPs, the evaluation of DFT reference energies and forces for the selected configurations, and the sampling of new configurations. With this code, the workflow for the development of a C-NNP model for water described in Sec. IV A was carried out over four generations.

All NNP optimizations were performed with the open-source n2p2 code,²⁷ and unless stated otherwise, all of the optimization parameters have been chosen according to the detailed benchmarking of this code for water.²⁷ We decided to use an established set of symmetry functions, which has been shown to be able to reproduce the properties of water over a large range of conditions.²³ The values of each symmetry function were centered around the respective average value of the training set and normalized to values between zero and one. These atomic environment vectors serve as the input for the atomic NNs consisting of two hidden layers of 20 neurons each with hyperbolic tangent activation functions for these two hidden layers and a linear activation function for the output neuron. The eight-member C-NNP models were constructed by random subsampling of the full set of reference data for each committee member, where 10% of the points were randomly left out in each case. After a different random initialization for each committee member, the weights and biases of the NNs were optimized using the parallel multistream version²⁷ of the adaptive global extended Kalman filter,^{67,68} as implemented in n2p2. C-NNPs used for QbC were optimized for 15 epochs with 8 streams, while the final C-NNPs for each generation, to be used for simulations, were optimized for 100 epochs with 32 streams.

The DFT reference calculations were all performed with the CP2K simulation package^{58,69} and its Quickstep module⁷⁰ using the exact same electronic structure setup for the revPBE0-D3 functional^{71–74} employed in Ref. 65. As shown therein, this choice of functional provides reliable properties of water, especially with quantum nuclei, and is therefore the ideal choice for the development of our C-NNP model.

The simulations using the C-NNP models used to generate candidate structures for the next active learning generation were all performed with a modified version of the CP2K simulation package.^{58,69} All classical simulations were propagated for 0.5 ns with a time step of 0.5 fs and a global canonical sampling through velocity rescaling (CSVR) thermostat⁷⁵ with a 30 fs time constant to sample the NVT ensemble. From these simulations, every 20th configuration

was saved to ensure uncorrelated statistics, resulting in 50 000 configurations for all classical ensembles. All PIMD simulations were propagated for 0.25 ns with a time step of 0.25 fs and a path integral Langevin equation thermostat⁶⁰ to sample the quantum canonical ensemble. Partially converged path integral simulations were realized with four replicas to discretize the path integral, which corresponds approximately to the midpoint between classical simulations and converged path integral simulations, as shown, for example, in Ref. 77 for the prototypical hydrogen bond in the Zundel cation. Full PIMD simulations were performed with 16 replicas for all simulations except at 250 K, where 32 replicas were used. In all quantum cases, every 40th configuration from four path integral replica trajectories was used to generate the set of candidate structures for the QbC cycles, resulting in a total of 100 000 configurations for every quantum ensemble. After careful testing of the influence of the energy committee disagreement biasing derived above, mild biasing with a harmonic constant $k^{(b)} = 0.95/\text{eV}$ per atom and a shift of $\sigma_0 = 0.1$ meV per atom was used for all simulations during the active learning process. The C-NNP model of generation 1 was used for the simulation of liquid water along an isochor in a cubic box of size 12.42 Å with 64 water molecules at three temperatures of 300 K, 350 K, and 400 K. The proton disordered phase ice I_h was simulated at 250 K in supercells under periodic boundary conditions including 96 water molecules, where the lattice size and initial conditions were chosen according to Ref. 78. These four state points were targeted with both a classical and partially converged quantum description of the nuclei for the generation 2 C-NNP model. The resulting C-NNP model of generation 2 was used to perform converged PIMD simulations along the same isochor of liquid water as well as for ice I_h. Besides these four simulations, a water slab was simulated with 216 molecules in a $15 \times 50 \times 15$ Å³ periodic box starting from an initial condition from Ref. 79 with both classical and partially converged quantum nuclei. Finally, the generation 3 C-NNP model was employed to simulate liquid water along an isotherm for two pressures of 2 kbar and 4 kbar, resulting in cubic boxes with sizes of 12.13 Å and 11.93 Å for 64 water molecules. In addition, ice VIII was simulated at 250 K with 64 molecules in a $9.70 \times 9.70 \times 14.11$ Å³ periodic box. These six simulations were performed for both classical and converged quantum nuclei. Finally, the water slab was simulated as in the previous generation, but now with converged quantum nuclei.

In order to generate an independent test set to validate the performance of the C-NNP models at the individual generations and the various state points, we used the final C-NNP model of generation 4 to perform independent simulations at all previously targeted conditions and state points. All classical and quantum simulations were propagated for 50 ps and 25 ps, respectively, with otherwise identical settings as for the above-mentioned simulations. The path integral has been fully converged for all these simulations by using 128 replicas for all simulations at and above 300 K and 256 replicas for the ice phases at 250 K. The biasing of the energy committee disagreement has been switched off for these production runs. From these eight classical and eight quantum ensembles in total, 500 uncorrelated structures were extracted in each case, and the reference revPBE0-D3 energies and forces were evaluated. This set of 8000 reference calculations, spanning very different conditions for both classical and quantum nuclei, was consequently used to validate the performance of the C-NNP models.

Finally, the performance of the generation 1 and 4 C-NNP models for static and dynamical properties was benchmarked against the previously published AIMD and AIPIMD results for liquid water.⁶⁵ For this purpose, 64 water molecules in a cubic box of size 12.42 Å were simulated at 300 K in the NVT ensemble to match exactly the setup of the reference *ab initio* simulations. We performed classical simulations with both C-NNP models using a time step of 0.5 fs and a global CSVN thermostat⁷⁵ with a time constant of 1 ps. Results including nuclear quantum effects were obtained from TRPMD simulations^{6,68} with 32 path integral replicas using a 0.25 fs time step. We accumulated a total simulation length of 2 ns in each case for both models. Again, no biasing of the committee disagreement was used for these production simulations.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional benchmarking results for the final generation 4 C-NNP model and detailed analysis of the performance of the C-NNP model obtained in generation 1.

ACKNOWLEDGMENTS

This work was supported by the Primus16/SCI/27/247019 grant from Charles University. This work was partially supported by the OP RDE project (No. CZ.02.2.69/0.0/0.0/18_070/0010462), International mobility of researchers at Charles University (MSCA-IF II). This work was supported by Project No. SVV 260586 of Charles University. K.B. acknowledges funding from the IMPRS for Many Particle Systems in Structured Environments. Part of this work was supported by the *Alexander von Humboldt-Stiftung* through a research grant awarded to C.S.

APPENDIX: BIASING OF FORCE DISAGREEMENT

An alternative approach to biasing due to the committee disagreement that provides sensitivity to local structural changes can be based on the force disagreement introduced in Eq. (2). Rather than bias the total energy disagreement, one can bias the disagreement of each atomic force vector separately. The atom α that enters a region with a force disagreement larger than σ_0 brings a contribution to its biasing energy into the system given again by a shifted harmonic form as

$$E_{\alpha}^{(b)} = \theta(\sigma_{F_{\alpha}} - \sigma_0) \frac{1}{2} k^{(b)} (\sigma_{F_{\alpha}} - \sigma_0)^2. \quad (\text{A1})$$

This affects all the other atoms in the system, say, β , and the corresponding biasing force exerted on them is

$$\begin{aligned} \mathbf{F}_{\beta\alpha}^{(b)} &= -\nabla_{\beta} E_{\alpha}^{(b)} \\ &= -\theta(\sigma_{F_{\alpha}} - \sigma_0) k^{(b)} (\sigma_{F_{\alpha}} - \sigma_0) \nabla_{\beta} \sigma_{F_{\alpha}}. \end{aligned} \quad (\text{A2})$$

At this point, we have to express $\nabla_{\beta} \sigma_{F_{\alpha}}$ in order to obtain a useful expression for the biasing force. Differentiating through the definition given in Eq. (2), we get

$$\begin{aligned}\nabla_{\beta}\sigma_{F_{i\alpha}} &= \nabla_{\beta}\left[\frac{1}{n}\sum_{i=1}^n(\nabla_{\alpha}\Delta E_i)^2\right]^{\frac{1}{2}} \\ &= \frac{1}{n}\frac{1}{2\sigma_{F_{i\alpha}}}\sum_{i=1}^n\nabla_{\beta}[(\nabla_{\alpha}\Delta E_i)\cdot(\nabla_{\alpha}\Delta E_i)] \\ &= \frac{1}{n}\frac{1}{\sigma_{F_{i\alpha}}}\sum_{i=1}^n(\nabla_{\alpha}\Delta E_i)\cdot(\nabla_{\beta}\nabla_{\alpha}\Delta E_i)\end{aligned}\quad (\text{A3})$$

as the desired result. A second derivative matrix operator $\nabla_{\beta}\nabla_{\alpha}$ appears, which generally requires an evaluation of analytical or numerical second derivatives for all $3N$ degrees of freedom, and thus, its applicability for practical calculations in this raw form is limited. However, we note that in this case, the second derivative is projected on the direction of the force $F_{i\alpha}$. This suggests that finite difference methods that limit the number of numerical derivative evaluations to only the one in the desired direction to evaluate the directional derivative could result in a practical computation scheme.^{31,32} Still, in addition to having to deal with the second derivatives, this biasing has to be evaluated separately for each atom α and would therefore incur a substantial computational cost. For this reason, we have not implemented it and used only the biasing of the total energy disagreement given by Eq. (3), which worked sufficiently well in practical simulations.

DATA AVAILABILITY

The data that support the findings of this study are openly available in paper-c-nnp at <http://doi.org/10.5281/zenodo.4004590>.

REFERENCES

- 1 J. Behler, "Perspective: Machine learning potentials for atomistic simulations," *J. Chem. Phys.* **145**, 170901 (2016).
- 2 A. P. Bartók, S. De, C. Poelking, N. Bernstein, J. R. Kermode, G. Csányi, and M. Ceriotti, "Machine learning unifies the modeling of materials and molecules," *Sci. Adv.* **3**, e1701816 (2017).
- 3 K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, "Machine learning for molecular and materials science," *Nature* **559**, 547–555 (2018).
- 4 V. L. Deringer, M. A. Caro, and G. Csányi, "Machine learning interatomic potentials as emerging tools for materials science," *Adv. Mater.* **31**, 1902765 (2019).
- 5 T. Mueller, A. Hernandez, and C. Wang, "Machine learning for interatomic potential models," *J. Chem. Phys.* **152**, 050902 (2020).
- 6 J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.* **98**, 146401 (2007).
- 7 J. Behler, "First principles neural network potentials for reactive simulations of large molecular and condensed systems," *Angew. Chem., Int. Ed.* **56**, 12828–12840 (2017).
- 8 J. Behler, "Atom-centered symmetry functions for constructing high-dimensional neural network potentials," *J. Chem. Phys.* **134**, 074106 (2011).
- 9 S. A. Ghasemi, A. Hofstetter, S. Saha, and S. Goedecker, "Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network," *Phys. Rev. B* **92**, 045131 (2015); [arXiv:1501.07344](https://arxiv.org/abs/1501.07344).
- 10 A. Khorshidi and A. A. Peterson, "Amp: A modular approach to machine learning in atomistic simulations," *Comput. Phys. Commun.* **207**, 310–324 (2016).
- 11 K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, "Quantum-chemical insights from deep tensor neural networks," *Nat. Commun.* **8**, 13890 (2017).
- 12 N. Artrith, A. Urban, and G. Ceder, "Efficient and accurate machine-learning interpolation of atomic energies in compositions with many species," *Phys. Rev. B* **96**, 014112 (2017); [arXiv:1706.06293](https://arxiv.org/abs/1706.06293).
- 13 J. S. Smith, O. Isayev, and A. E. Roitberg, "ANI-1: An extensible neural network potential with DFT accuracy at force field computational cost," *Chem. Sci.* **8**, 3192–3203 (2017); [arXiv:1610.08935](https://arxiv.org/abs/1610.08935).
- 14 L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, "Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics," *Phys. Rev. Lett.* **120**, 143001 (2018); [arXiv:1707.09571](https://arxiv.org/abs/1707.09571).
- 15 O. T. Unke and M. Meuwly, "PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges," *J. Chem. Theory Comput.* **15**, 3678–3693 (2019).
- 16 Y. Shao, M. Hellström, P. D. Mitev, L. Knijff, and C. Zhang, "PiNN: A python library for building atomic neural networks of molecules and materials," *J. Chem. Inf. Model.* **60**, 1184–1193 (2020); [arXiv:1910.03376](https://arxiv.org/abs/1910.03376).
- 17 A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, "Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons," *Phys. Rev. Lett.* **104**, 136403 (2010).
- 18 M. Rupp, A. Tkatchenko, K. R. Müller, and O. A. von Lilienfeld, "Fast and accurate modeling of molecular atomization energies with machine learning," *Phys. Rev. Lett.* **108**, 058301 (2012).
- 19 A. P. Thompson, L. P. Swiler, C. R. Trott, S. M. Foiles, and G. J. Tucker, "Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials," *J. Comput. Phys.* **285**, 316–330 (2015).
- 20 A. V. Shapeev, "Moment tensor potentials: A class of systematically improvable interatomic potentials," *Multiscale Model. Simul.* **14**, 1153–1173 (2015).
- 21 Z. Li, J. R. Kermode, and A. De Vita, "Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces," *Phys. Rev. Lett.* **114**, 096405 (2015).
- 22 S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller, "Machine learning of accurate energy-conserving molecular force fields," *Sci. Adv.* **3**, e1603015 (2017).
- 23 T. Morawietz, A. Singraber, C. Dellago, and J. Behler, "How van der Waals interactions determine the unique properties of water," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 8368–8373 (2016).
- 24 T. Morawietz, O. Marsalek, S. R. Pattenau, L. M. Streecker, D. Ben-Amotz, and T. E. Markland, "The interplay of structure and dynamics in the Raman spectrum of liquid water over the full frequency and temperature range," *J. Phys. Chem. Lett.* **9**, 851–857 (2018).
- 25 B. Cheng, E. A. Engel, J. Behler, C. Dellago, and M. Ceriotti, "Ab initio thermodynamics of liquid and solid water," *Proc. Natl. Acad. Sci. U. S. A.* **116**, 1110–1115 (2019); [arXiv:1811.08630](https://arxiv.org/abs/1811.08630).
- 26 N. Artrith and A. Urban, "An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for TiO₂," *Comput. Mater. Sci.* **114**, 135–150 (2016).
- 27 A. Singraber, T. Morawietz, J. Behler, and C. Dellago, "Parallel multistream training of high-dimensional neural network potentials," *J. Chem. Theory Comput.* **15**, 3075–3092 (2019).
- 28 S. K. Natarajan and J. Behler, "Neural network molecular dynamics simulations of solid-liquid interfaces: Water at low-index copper surfaces," *Phys. Chem. Chem. Phys.* **18**, 28704–28725 (2016).
- 29 M. Hellström and J. Behler, "Structure of aqueous NaOH solutions: Insights from neural-network-based molecular dynamics simulations," *Phys. Chem. Chem. Phys.* **19**, 82–96 (2017).
- 30 V. Quaranta, M. Hellström, and J. Behler, "Proton-transfer mechanisms at the water-ZnO interface: The role of presolvation," *J. Phys. Chem. Lett.* **8**, 1476–1483 (2017).
- 31 M. Hellström, V. Quaranta, and J. Behler, "One-dimensional vs. two-dimensional proton transport processes at solid-liquid zinc-oxide-water interfaces," *Chem. Sci.* **10**, 1232–1243 (2019).
- 32 R. T. Clemen, "Combining forecasts: A review and annotated bibliography," *Int. J. Forecasting* **5**, 559–583 (1989).
- 33 L. K. Hansen and P. Salamon, "Neural network ensembles," *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 993–1001 (1990).

- ³⁴A. Krogh and J. Vedelsby, "Neural network ensembles, cross validation, and active learning," *Adv. Neural Inf. Process. Syst.* **7**, 231–238 (1995).
- ³⁵Y. Zhao, J. Gao, and X. Yang, "A survey of neural network ensembles," in *Proceedings of the 2005 International Conference on Neural Networks and Brain (ICNNB'05)* (2005), Vol. 1, pp. 438–442.
- ³⁶L. Rokach, "Ensemble-based classifiers," *Artif. Intell. Rev.* **33**, 1–39 (2010).
- ³⁷S. Hashem and B. Schmeiser, "Improving model accuracy using optimal linear combinations of trained neural networks," *IEEE Trans. Neural Networks* **6**, 792–794 (1995).
- ³⁸L. Breiman, "Bagging predictors," *Mach. Learn.* **24**, 123–140 (1996).
- ³⁹S. Hashem, "Optimal linear combinations of neural networks," *Neural Networks* **10**, 599–614 (1997).
- ⁴⁰D. H. Wolpert, "Stacked generalization," *Neural Networks* **5**, 241–259 (1992).
- ⁴¹N. Ueda and R. Nakano, "Generalization error of ensemble estimators," in *IEEE International Conference on Neural Networks* (IEEE, 1996), Vol. 1, pp. 90–95.
- ⁴²P. Sollich and A. Krogh, "Learning with ensembles: How overfitting can be useful," in *Proceedings of the 1995 Conference* (MIT Press, 1996), pp. 4–10.
- ⁴³H. S. Seung, M. Oppen, and H. Sompolinsky, "Query by committee," in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory (COLT'92)* (ACM Press, New York, NY, USA, 1992), pp. 287–294.
- ⁴⁴N. Arrithr and J. Behler, "High-dimensional neural network potentials for metal surfaces: A prototype study for copper," *Phys. Rev. B* **85**, 045439 (2012).
- ⁴⁵L. Chen, I. Sukuba, M. Probst, and A. Kaiser, "Iterative training set refinement enables reactive molecular dynamics via machine learned forces," *RSC Adv.* **10**, 4293–4299 (2020).
- ⁴⁶J. Behler, "Neural network potential-energy surfaces in chemistry: A tool for large-scale simulations," *Phys. Chem. Chem. Phys.* **13**, 17930–17955 (2011).
- ⁴⁷M. Gastegger, J. Behler, and P. Marquetand, "Machine learning molecular dynamics for the simulation of infrared spectra," *Chem. Sci.* **8**, 6924–6935 (2017); M. Gastegger and P. Marquetand, "Molecular Dynamics with Neural Network Potentials," in *Machine Learning Meets Quantum Physics: Lecture Notes in Physics* (Springer, 2020), vol. 968.
- ⁴⁸N. Raimbault, A. Grisafi, M. Ceriotti, and M. Rossi, "Using Gaussian process regression to simulate the vibrational Raman spectra of molecular crystals," *New J. Phys.* **21**, 105001 (2019); [arXiv:1906.07485](https://arxiv.org/abs/1906.07485).
- ⁴⁹C. Schran, J. Behler, and D. Marx, "Automated fitting of neural network potentials at coupled cluster accuracy: Protonated water clusters as testing ground," *J. Chem. Theory Comput.* **16**, 88–99 (2020); [arXiv:1908.08734](https://arxiv.org/abs/1908.08734).
- ⁵⁰F. Musil, M. J. Willatt, M. A. Langovoy, and M. Ceriotti, "Fast and accurate uncertainty estimation in chemical machine learning," *J. Chem. Theory Comput.* **15**, 906–915 (2019); [arXiv:1809.07653](https://arxiv.org/abs/1809.07653).
- ⁵¹E. V. Podryabinkin and A. V. Shapeev, "Active learning of linearly parametrized interatomic potentials," *Comput. Mater. Sci.* **140**, 171–180 (2017).
- ⁵²L. Zhang, D.-Y. Lin, H. Wang, R. Car, and W. E, "Active learning of uniformly accurate interatomic potentials for materials simulation," *Phys. Rev. Mater.* **3**, 023804 (2019); [arXiv:1810.11890](https://arxiv.org/abs/1810.11890).
- ⁵³J. S. Smith, B. Nebgen, N. Lubbers, O. Isayev, and A. E. Roitberg, "Less is more: Sampling chemical space with active learning," *J. Chem. Phys.* **148**, 241733 (2018); [arXiv:1801.09319](https://arxiv.org/abs/1801.09319).
- ⁵⁴V. L. Deringer, C. J. Pickard, and G. Csányi, "Data-driven learning of total and local energies in elemental boron," *Phys. Rev. Lett.* **120**, 156001 (2018).
- ⁵⁵M. Ceriotti, G. A. Tribello, and M. Parrinello, "Demonstrating the transferability and the descriptive power of sketch-map," *J. Chem. Theory Comput.* **9**, 1521–1532 (2013).
- ⁵⁶F. Musil, S. De, J. Yang, J. E. Campbell, G. M. Day, and M. Ceriotti, "Machine learning for the structure-energy-property landscapes of molecular crystals," *Chem. Sci.* **9**, 1289–1300 (2018).
- ⁵⁷D. Marx and J. Hutter, *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods* (Cambridge University Press, Cambridge, 2009).
- ⁵⁸See <https://www.cp2k.org> for CP2K, released under GPL license, 2020.
- ⁵⁹J. Behler *et al.*, RuNNer: A Neural Network Code for High-Dimensional Potential-Energy Surfaces, Universität Göttingen 2020 (GPL3 license), 2020.
- ⁶⁰Y. Zhai, A. Caruso, S. Gao, and F. Paesani, "Active learning of many-body configuration space: Application to the Cs⁺-water MB-nrg potential energy function as a case study," *J. Chem. Phys.* **152**, 144103 (2020).
- ⁶¹Q. Lin, Y. Zhang, B. Zhao, and B. Jiang, "Automatically growing global reactive neural network potential energy surfaces: A trajectory-free active learning strategy," *J. Chem. Phys.* **152**, 154104 (2020); [arXiv:2002.05912](https://arxiv.org/abs/2002.05912).
- ⁶²C. Schran, F. Uhl, J. Behler, and D. Marx, "High-dimensional neural network potentials for solvation: The case of protonated water clusters in helium," *J. Chem. Phys.* **148**, 102310 (2018).
- ⁶³T. E. Markland and M. Ceriotti, "Nuclear quantum effects enter the mainstream," *Nat. Rev. Chem.* **2**, 0109 (2018); [arXiv:1803.01037](https://arxiv.org/abs/1803.01037).
- ⁶⁴M. Ceriotti, W. Fang, P. G. Kusalik, R. H. McKenzie, A. Michaelides, M. A. Morales, and T. E. Markland, "Nuclear quantum effects in water and aqueous systems: Experiment, theory, and current challenges," *Chem. Rev.* **116**, 7529–7550 (2016).
- ⁶⁵O. Marsalek and T. E. Markland, "Quantum dynamics and spectroscopy of ab initio liquid water: The interplay of nuclear and electronic quantum effects," *J. Phys. Chem. Lett.* **8**, 1545–1551 (2017).
- ⁶⁶M. D. Baer, C. J. Mundy, M. J. McGrath, L.-F. W. Kuo, J. I. Siepmann, and D. J. Tobias, "Re-examining the properties of the aqueous vapor–liquid interface using dispersion corrected density functional theory," *J. Chem. Phys.* **135**, 124712 (2011).
- ⁶⁷S. Shah, F. Palmieri, and M. Datum, "Optimal filtering algorithms for fast learning in feedforward neural networks," *Neural Networks* **5**, 779–787 (1992).
- ⁶⁸T. B. Blank and S. D. Brown, "Adaptive, global, extended Kalman filters for training feedforward neural networks," *J. Chemom.* **8**, 391–407 (1994).
- ⁶⁹J. Hutter, M. Iannuzzi, F. Schiffmann, and J. Vandevondele, "CP2K: Atomistic simulations of condensed matter systems," *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **4**, 15–25 (2014).
- ⁷⁰J. Vandevondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, "QUICKSTEP: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach," *Comput. Phys. Commun.* **167**, 103–128 (2005).
- ⁷¹J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximation made simple," *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- ⁷²Y. Zhang and W. Yang, "Comment on 'Generalized gradient approximation made simple'," *Phys. Rev. Lett.* **80**, 890 (1998).
- ⁷³C. Adamo and V. Barone, "Toward reliable density functional methods without adjustable parameters: The PBE0 model," *J. Chem. Phys.* **110**, 6158–6170 (1999).
- ⁷⁴S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, "A consistent and accurate *ab initio* parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu," *J. Chem. Phys.* **132**, 154104 (2010).
- ⁷⁵G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *J. Chem. Phys.* **126**, 014101 (2007); [arXiv:0803.4060](https://arxiv.org/abs/0803.4060).
- ⁷⁶M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, "Efficient stochastic thermostatting of path integral molecular dynamics," *J. Chem. Phys.* **133**, 124104 (2010); [arXiv:1009.1045](https://arxiv.org/abs/1009.1045).
- ⁷⁷C. Schran, F. Briec, and D. Marx, "Converged colored noise path integral molecular dynamics study of the Zundel cation down to ultralow temperatures at coupled cluster accuracy," *J. Chem. Theory Comput.* **14**, 5068–5078 (2018).
- ⁷⁸C. Schran and D. Marx, "Quantum nature of the hydrogen bond from ambient conditions down to ultra-low temperatures," *Phys. Chem. Chem. Phys.* **21**, 24967–24975 (2019).
- ⁷⁹R. Vácha, O. Marsalek, A. P. Willard, D. J. Bonhuis, R. R. Netz, and P. Jungwirth, "Charge transfer between water molecules as the possible origin of the observed charging at the surface of pure water," *J. Phys. Chem. Lett.* **3**, 107–111 (2011).
- ⁸⁰M. Rossi, M. Ceriotti, and D. E. Manolopoulos, "How to remove the spurious resonances from ring polymer molecular dynamics," *J. Chem. Phys.* **140**, 234116 (2014); [arXiv:1406.1074v1](https://arxiv.org/abs/1406.1074v1).
- ⁸¹V. Kapil, J. Behler, and M. Ceriotti, "High order path integrals made easy," *J. Chem. Phys.* **145**, 234103 (2016).
- ⁸²M. Buchowiecki and J. Vaníček, "Monte Carlo evaluation of the equilibrium isotope effects using the Takahashi-Imada factorization of the Feynman path integral," *Chem. Phys. Lett.* **588**, 11–16 (2013).

C.2 | Paper II

Brezina, K., Beck, H. and Marsalek, O.: Reducing the cost of neural network potential generation for reactive molecular systems, *Journal of Chemical Theory and Computation*, 19, 19, 6589–6604, 2023

Reducing the Cost of Neural Network Potential Generation for Reactive Molecular Systems

Krystof Brezina, Hubert Beck, and Ondrej Marsalek*

Cite This: *J. Chem. Theory Comput.* 2023, 19, 6589–6604

Read Online

ACCESS |


Metrics & More

Article Recommendations

Supporting Information

Downloaded via CHARLES UNIV IN PRAGUE on March 20, 2024 at 12:18:24 (UTC).
See <https://pubs.acs.org/sharingguidelines> for options on how to legitimately share published articles.

ABSTRACT: Although machine learning potentials have recently had a substantial impact on molecular simulations, the construction of a robust training set can still become a limiting factor, especially due to the requirement of a reference ab initio simulation that covers all the relevant geometries of the system. Recognizing that this can be prohibitive for certain systems, we develop the method of transition tube sampling that mitigates the computational cost of training set and model generation. In this approach, we generate classical or quantum thermal geometries around a transition path describing a conformational change or a chemical reaction using only a sparse set of local normal mode expansions along this path and select from these geometries by an active learning protocol. This yields a training set with geometries that characterize the whole transition without the need for a costly reference trajectory. The performance of the method is evaluated on different molecular systems with the complexity of the potential energy landscape increasing from a single minimum to a double proton-transfer reaction with high barriers. Our results show that the method leads to training sets that give rise to models applicable in classical and path integral simulations alike that are on par with those based directly on ab initio calculations while providing the computational speedup we have come to expect from machine learning potentials.



1. INTRODUCTION

Owing to the detailed atomistic insight into the structure and dynamics of molecular systems and materials, the relevance of computer simulations of molecular dynamics (MD) in current research is undeniable. MD simulations represent a valuable analytic and predictive tool in multiple fields of both basic and applied research including physical chemistry, materials science, or drug design.^{1–5} They also provide a way to explain and corroborate experimental data that might be difficult to interpret otherwise. For many systems of interest, MD simulations can be routinely performed under the Born–Oppenheimer approximation in the electronic ground state, with the nuclei being treated either classically or quantum-mechanically within the imaginary time path integral formalism. This makes the choice of the potential energy surface (PES) a key decision that determines the accuracy of the resulting simulation. Out of the available options, ab initio molecular dynamics⁶ (AIMD) is a state-of-the-art methodology that relies on a full, on-the-fly quantum electronic structure calculation^{7,8} at every step of the simulation to evaluate the potential energy and forces. This is most commonly performed at the level of density functional theory^{8–10} (DFT), which provides correlated electronic energies at a computational cost accessible in practice, but for smaller systems, the use of correlated wave function methods is feasible as well.^{11–13} In any case, the computational

cost of AIMD simulations is typically large—especially so for advanced hybrid DFT functionals in the condensed phase—and can easily become prohibitive in the light of the ever-growing demand for larger time and length scales of the relevant simulations.

This issue can be mitigated by the recent development of the so-called machine learning potentials (MLPs).^{14,15} These use various machine learning approaches^{16–20} to faithfully approximate the desired ab initio PES by training on a reference data set consisting of a relatively modest number of ab initio geometries and their corresponding energies and, optionally, forces. As such, they indeed combine the best of the two worlds: they are able to maintain the accuracy of the parent ab initio method, but they also circumvent the need for explicit electronic structure calculations at each step of the MD simulation. Thus, they evaluate the potential energy and forces at a significantly reduced computational cost.²¹ One particular flavor of MLPs of major practical importance is represented by

Received: April 6, 2023

Published: September 25, 2023



neural network potentials (NNPs), which rely on artificial neural networks combined with a set of appropriate atomic descriptors to accurately represent the molecular geometry-to-energy relationship, including all its symmetries.^{16,22} NNPs have repeatedly proved their worth in modeling a plethora of various molecular systems ranging from liquids and solutions to interfaces and solids.^{20,23–26} Our recent study,^{26,27} building on the findings of previous studies focusing on NNPs,^{28–30} shows that rather than using a single NNP to represent the PES, it is advantageous to build a model as a committee^{31,32} of NNPs (C-NNP) that comprises a small number of NNPs, each trained individually to a subset of the main training set. The advantage is twofold: first, the energy prediction obtained as the committee average is known to be a better approximation of the ab initio energy than the estimates of the individual members.^{29,33} Second, the committee disagreement,³⁴ represented by the standard deviation of the individual member estimates of energies or forces, serves as a valuable indicator of prediction reliability and can be used to monitor and optionally ensure the stability of the simulation.²⁷ Crucially, this disagreement can be used as the key ingredient of the active learning process called query by committee³⁵ (QbC) that systematically builds the training set in a data-driven way.^{27,34,36}

An accurate and stable NNP can only be obtained on a foundation of robust, high-quality training data. This is typically based on a reference AIMD trajectory, from which geometries are selected for the training set, together with the corresponding energies and forces. However, the trajectory is highly correlated in time and thus most of the expensive AIMD data do not contribute useful information for the training of the model. This selection has been approached in different ways from random sampling and manual selection to more data-driven procedures,^{25,27,30,37–43} with QbC being a particularly efficient method. QbC considers a set of candidate structures, in this case, the whole AIMD trajectory, and iteratively builds up the training set. It starts by training a C-NNP on a very small set of initial configurations and using its disagreement to screen the candidate configurations for those with the most uncertain prediction. A small number of these configurations are then added to the training set, a new C-NNP is trained, and the process is iteratively repeated until some convergence criteria are met. In comparison to random selection, this approach is known to generate more compact training sets that give rise to robust models of similar accuracy.^{30,39} Even though the initial AIMD trajectory is typically the most expensive part of the procedure, numerous successful MLPs have been generated on top of reasonably short AIMD simulations.²⁶

However, for many purposes, this process involving AIMD is still too expensive to be practical. For instance, the requirements on a high-level electronic structure method can raise the computational demands above a reasonable threshold. One might also be interested in a system that features rare events, such as chemical or conformational changes, which will happen quickly and occur infrequently or not at all in a direct AIMD simulation. In turn, these crucial configurations are under-represented in the set of candidates and enhanced sampling simulations would be required in order to construct a robust training set, which typically raises the computational cost further by one or more orders of magnitude.

In case such a situation occurs, one needs to adhere to an approximate method of candidate generation that relieves

some of the computational expenses while maintaining the quality of the resulting candidate set. For simple systems with a single potential energy minimum, the solution is fairly straightforward. In this case, one can benefit from a random sampling of displacements in the directions of a fixed set of normal modes to obtain a set of distorted configurations. This approach, sometimes called normal mode sampling (NMS) in the literature,¹⁸ avoids the cost of a full AIMD simulation by replacing it with a more manageable combination of a Hessian matrix evaluation and a number of single-point electronic structure calculations for the generated geometries. The sampling of the known normal mode distribution itself yields uncorrelated samples by definition and requires no ab initio calculations; therefore, its cost is negligible. Various versions of this approach were successfully used to generate structures for the training of MLPs. Using a scaled uniform random sampling of the normal modes, the method was used to obtain auxiliary structures used in model validation⁴⁴ and with approximate thermal distortions in NNP training set generation around configurational minima¹⁸ as well as to construct an NNP model for a gas-phase ammonia molecule.⁴² Clearly, the utility of NMS is limited when the harmonic approximation becomes insufficient. This can be the case if individual modes are strongly anharmonic or coupled, or if the system features conformational changes or reactions, where multiple local minima come into play. The need for reactive training data sets was recognized in a recent work introducing the Transition-1x data set,³⁵ which includes training points along a converged minimum energy path (MEP) obtained through a nudged elastic band calculation⁴⁶ and its surrounding arising from prior unconverged iterations of the optimization.

In this work, we propose transition tube sampling (TTS), a robust and general approach to the generation of training sets and models that are able to accurately describe processes that feature transitions over potential energy barriers, which includes both conformational flexibility and chemical reactivity. We achieve this by generating thermally distorted candidate geometries along a reaction pathway with the help of multiple normal mode expansions and screening these candidates using QbC. The role of the minimum geometry in NMS is taken by the MEP that describes the course of the reaction through configuration space. Local harmonic expansions are performed in a small number of relevant configurations along the MEP and physically relevant candidate configurations are generated with uniform distribution along the MEP and with classical or quantum thermal weights in all perpendicular directions based on one of the sets of normal modes. An arbitrary number of these candidate configurations can be generated at a negligible computational cost and submitted to the QbC process, which selects the most important ones to have ab initio calculations performed and to be included in the training set. This results in compact and robust training sets and models that maintain consistent accuracy along the reaction path, making them suitable for MD simulations of the reactive process, including enhanced sampling simulations, while no AIMD trajectories are required as part of this process. We test this method on three different molecules in the gas phase to illustrate its capabilities.

The rest of the paper is organized as follows. In Section 2, we begin by formalizing thermal NMS, which samples the exact classical or quantum canonical distribution under the harmonic approximation. With the obtained framework, we then proceed to introduce the MEP into the picture and

describe the technical details of TTS. In Section 3, we describe how we used TTS to create C-NNP models, the simulations performed with these models, and other related computational details. In Section 4, we apply this approach to three different gas-phase systems of increasing complexity represented by the molecules of benzene, malonaldehyde, and 2,5-diaminobenzoquinone-1,4-dimine (DABQDI) and discuss its successes and possible pitfalls. Section 5 concludes the paper and offers outlooks concerning the generalization and the limitations of the method beyond the gas phase.

2. THEORY

In this section, we discuss the theoretical basis of the TTS method. In this approach, we rely on the harmonic approximation and the vibrational normal mode formalism to obtain ab initio training data for the construction of C-NNPs for reactive systems without the need to run expensive sampling simulations, such as full AIMD. First, we present the simple key idea behind NMS which relies on the harmonic approximation to describe the underlying PES and thus is expected to work well for systems that are close to harmonic around a single given minimum geometry at the temperature of interest. Clearly, this does not yield a flexible and general method, since the harmonic approximation is readily challenged by many realistic systems, notably those that exhibit more pronounced configurational flexibility or chemical reactivity. Therefore, we propose a more general approach to sampling candidate geometries based on NMS which is applicable even to systems described by multiple minima separated by barriers. This is achieved using the harmonic expansion locally along an MEP in a way that eventually generates a balanced training set.

2.1. NMS for Thermal Sampling around Minimum Geometries. To open the discussion of the theory behind TTS, we first turn our attention to the simple case represented by a PES with a single minimum geometry \mathbf{R}_0 on which the nuclear motion is described by classical mechanics. Assuming a reasonable extent of validity of the harmonic approximation to capture the thermally accessible potential energy landscape, the classical thermal probability density ρ_c at temperature T is approximated by

$$\rho_c(\Omega_1, \dots, \Omega_{N_{\text{int}}}) \propto \prod_{i=1}^{N_{\text{int}}} \exp\left(-\frac{1}{2}\beta\omega_i^2\Omega_i^2\right) \quad (1)$$

In this expression, ω_i and Ω_i denote, respectively, the natural frequency and the normal coordinate corresponding to the i -th normalized vibrational normal mode vector $\mathbf{\Omega}_i$, and β is the inverse temperature equal to $1/k_{\text{B}}T$ (with k_{B} representing the Boltzmann constant). N_{int} is the total number of internal degrees of freedom of the species, typically $3N - 6$ for N atoms. Hence, in the harmonic approximation, the thermal density is described as a multivariate, yet uncoupled normal distribution where each i -th orthogonal degree of freedom has the standard deviation of $\sigma_i = 1/\sqrt{\beta\omega_i^2}$.

As such, it is straightforward to generate completely uncorrelated thermal geometries \mathbf{R} by distorting the minimum geometry \mathbf{R}_0 independently in the direction of each of the normal mode vectors. The appropriate magnitude of the distortions is given by a randomly generated value of the corresponding normal coordinate Ω_i from the distribution in eq 1. The instrumental prescription for this procedure is the

inverse coordinate transformation from normal modes back to Cartesian coordinates

$$\mathbf{R} = \mathbf{R}_0 + \mathbf{M}^{-1/2} \sum_{i=1}^{N_{\text{int}}} \Omega_i \mathbf{\Omega}_i \quad (2)$$

where \mathbf{M} represents the diagonal mass matrix. Thus, by drawing samples of normal coordinates and transforming them, we obtain correctly distributed thermal samples in Cartesian coordinates.

We can now perform thermal NMS by sampling from this auxiliary harmonic ensemble as a source of candidate geometries to be potentially included in the training set of an MLP. The auxiliary ensemble is thus never used directly and no expectation values are calculated over it. It only needs to provide good coverage of the thermally accessible region of the PES, which will be the case as long as the harmonic approximation is reasonably accurate for the system of interest. Specifically, we construct a training set in our active learning procedure by generating a large number of these NMS candidate geometries and screening them in a QbC process using a C-NNP model. In each QbC iteration, electronic structure calculations are performed only for a small number of selected structures to obtain their potential energies and possibly forces, which then comprise the final training set once the process converges. The computational cost is thus determined primarily by the geometry optimization procedure, the Hessian calculation, the C-NNP prediction required for screening, and the electronic structure calculations for the selected geometries. The cost of the sample generation is negligible. This approach is substantially less computationally demanding when compared to the more conventional approach of sampling the candidate geometries for QbC from an AIMD trajectory, which requires a large number of electronic structure calculations for very similar geometries that do not contribute diversity to the training set. In contrast to that, NMS generates fully decorrelated geometries by construction, and electronic structure calculations are only needed for the relatively small number of the most important geometries selected by the subsequent QbC process.

So far, we have focused on the situation where NMS is used to sample a classical distribution on the studied PES. However, since the harmonic approximation describes a molecule as a set of independent one-dimensional harmonic oscillators, we can readily generalize the above classical case to a quantum one as the analytic solution of the quantum harmonic oscillator is known. Specifically, it is straightforward to show (see Section S1 of the Supporting Information) that the canonical thermal density of a quantum harmonic oscillator at a given temperature is Gaussian just as its classical counterpart, but broader. This broadening is encoded in the quantum effective inverse temperature⁴⁷

$$\beta^*(\beta, \omega) = \frac{2}{\hbar\omega} \tanh\left(\frac{\beta\hbar\omega}{2}\right) \quad (3)$$

at which a classical harmonic oscillator would have the same thermal width as a quantum harmonic oscillator at a reference inverse temperature β . Since β^* is by definition a frequency-dependent quantity, one cannot describe the whole molecule by a single quantum effective temperature but instead has to assign one to each individual mode. In turn, the quantum thermal density is given by

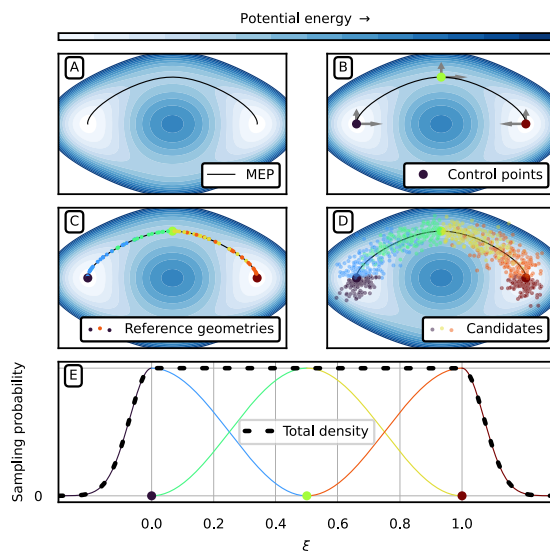


Figure 1. Schematic depiction of the TTS approach proposed for reactive systems. Panel A: an illustrative MEP winding through a model two-dimensional configuration space given by $V(x,y) = (1/6)\{4(1-x^2-y^2)^2 + 2(x^2-2)^2 + [(x+y)^2 - 1]2 + [(x-y)^2 - 1]^2 - 2\}$. Panel B: control points are selected and their local normal modes (gray arrows) are calculated. Here, the control points are taken as the two end-point minima on the MEP (purple and brown dots) and the transition state (yellow-green). Panel C: a much denser set of uniformly distributed reference geometries is generated along the MEP. Panel D: each of the reference geometries is distorted using the local modes of their assigned control point (as detailed in panel E) to become a candidate geometry. Panel E: a detailed view showing how each reference geometry is assigned to a set of local modes. A set of reference geometries on the MEP is assigned to each control point following eq 5. At the MEP edges, standard Gaussian thermal NMS is performed outside of the reaction coordinate (decaying purple and brown tails of the total density).

$$\rho_q(\Omega_1, \dots, \Omega_{N_m}) \propto \prod_{i=1}^{N_m} \exp\left[-\frac{1}{2}\beta^*(\beta, \omega_i)\omega_i^2\Omega_i^2\right] \quad (4)$$

This simple modification allows one to generate an auxiliary quantum ensemble at practically the same cost as the classical one that would otherwise need to be approached from a significantly more demanding perspective, perhaps based on sampling techniques using the imaginary time path integral formalism.

2.2. Transition Tube Sampling. Up to this point, we have relied on the ability of the harmonic expansion around a single minimum to approximate the real PES so that the generated samples cover sufficiently all the relevant regions for the purpose of generating an MLP. Arguably, this is a reasonable requirement for most stable molecules with a single minimum geometry where the onset of the anharmonic region connected to the dissociation of the molecule is not thermally accessible. On the other hand, it is a stringent requirement for molecular systems which display conformational changes or chemical reactivity and, therefore, are represented by multiple PES minima connected by MEPs: features not captured by a single harmonic expansion. However, in such cases, it is desirable for the resulting model to be able to describe the potential energy landscape not only around local potential energy minima but also in the transition regions. This is vital in the case of low- $k_B T$ barriers, where spontaneous transitions occur during direct

MD. Nonetheless, it cannot be omitted even in the case of high- $k_B T$ barriers where an enhanced sampling simulation would be required to cross the barrier. Even if the transition does not actually occur, the presence of the transition state may introduce substantial anharmonicity within the original PES basin that an eventual MLP should learn. However, in the case of barrier transitions it is not desirable to attempt to build the C-NNP model starting from a candidate set representing the true thermal ensemble, even if we could obtain it, since this would lead to a possibly detrimental under-representation of the high-energy configurations close to the transition state in the resulting candidate set and, in turn, to poor performance of the resulting model in the transition regions.

Therefore, we propose the TTS method: a generalization of NMS for systems with transitions that employ local normal modes along an MEP to sample uniformly along the path and with proper thermal weights in all perpendicular directions. This leads to an auxiliary harmonic ensemble that differs significantly from the true thermal one but enables the construction of MLPs with uniform accuracy along the whole transition. The TTS method naturally reduces to thermal NMS as described above for single-minimum systems in the zero MEP length limit. The process, illustrated in Figure 1, starts by finding the MEP $R(\xi)$ on the given PES (panel A). Here, ξ is a dimensionless reaction coordinate along the MEP curve through configuration space normalized to the interval from

0 to 1. In the following, we shall assume that the MEP is available as a continuous, differentiable function of the parameter ξ . In practice, this can be achieved by spline fitting of the discretized representation of the MEP originating from, for instance, a nudged elastic band calculation.⁴⁶ Note that by definition, the MEP is a minimum of the PES in all directions perpendicular to it. Once the relevant MEP is known, we continue by selecting a sparse set of control points \mathbf{R}_c , $c = 1, \dots, N_p$ along the MEP at positions ξ_c for which the Hessian matrices are calculated and diagonalized to give the set of local normal mode vectors $\Omega_{c,i}$ and their corresponding frequencies. For instance, this can be the two end-point minima and the transition state between them (Figure 1, panel B), although there is no constraint on how densely one might select the control points along the MEP other than the limiting computational expense of the Hessian matrix calculation. The selection of the control points is performed by hand by the user, ensuring a homogeneous coverage of the MEP. Formally, the expansion of the PES along the MEP becomes exact under the harmonic approximation in the limit of a large number of control points N_p . Since we want to achieve uniform sampling along the MEP, we now proceed to the generation of reference geometries on the MEP that do have this property. Specifically, to each control point \mathbf{R}_c , we first assign a probability distribution $p_c(\xi)$ defined on the interval $[\xi_{c-1}, \xi_{c+1}]$ (Figure 1, panel E, interval between purple and brown control points) as

$$p_c(\xi) = \begin{cases} \sin^2 \left[\frac{\pi}{2|\xi_c - \xi_{c-1}|} (\xi - \xi_{c-1}) \right] & \xi_{c-1} \leq \xi \leq \xi_c \\ \cos^2 \left[\frac{\pi}{2|\xi_c - \xi_{c+1}|} (\xi - \xi_c) \right] & \xi_c \leq \xi \leq \xi_{c+1} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Once this is done for all N_p control points, the identity

$$p(\xi) = \sum_{c=1}^{N_p} p_c(\xi) = 1 \quad (6)$$

holds over the whole length of the MEP (Figure 1, panel E over the range of ξ). Note that the choice of squares of harmonic functions is only one out of many possibilities, and any other pair of complementary functions that sum up to unity would work in this case. Next, we generate an arbitrary number of reference geometries $\mathbf{R}_0(\xi)$ at a chosen linear density by drawing random values of ξ from the above distributions and passing them to the continuous prescription of the MEP, all while keeping track of the parent c -th control point (Figure 1, panel C). Analogously to the distortion of the minimum geometry in the single-minimum case through eq 2, we distort each of these reference geometries using the normal modes and frequencies of its parent control point using

$$\mathbf{R} = \mathbf{R}_0(\xi) + \mathbf{M}^{-1/2} \sum_i \Omega_{c,i} [1 - \mathbb{P}(\xi)] \Omega_{c,i} \quad (7)$$

where the normal coordinate values $\Omega_{c,i}$ are sampled thermally according to eq 1 or 4 (Figure 1, panel D); the prime indicates that modes with imaginary frequencies are omitted from the sum. The matrix $\mathbb{P}(\xi)$ is the projector on the tangent direction at the point ξ which can be constructed analytically from

$d\mathbf{R}(\xi)/d\xi$. This is used to obtain distortions strictly perpendicular to the MEP and thus to correct for the approximate validity of the normal mode expansion calculated at ξ_c for all the displaced geometries. However, the use of the decaying probability distributions (eq 5) favors the use of the local modes close to their origin. Through this procedure, one obtains a set of candidate geometries distributed inside a tube around the MEP the width of which is given thermally. At this point, this tube still has open ends cut sharply by the planes defined by normal vectors equal to the MEP tangent vector at the end points of the MEP. Since these points are (usually) also well-defined minima on the PES, the presence of these sharp edges is easily sanitized by appending the usual thermal NMS samples at these minima, although only adding the configurations away from the MEP (Figure 1, decaying purple and brown lines). In other words, just one-half of the multivariate Gaussian is appended to the tube that does not overlap with it. In our TTS implementation, we ensure that the uniform density of the sampling along the MEP and the one at the peak of the half-Gaussian are seamlessly matched (as described in Section S1 of the Supporting Information).

Using the described sampling approach leads to an auxiliary ensemble of candidates that does not correspond to the true thermal ensemble, but contains a balanced selection of geometries distributed uniformly along the MEP with classical or quantum thermal displacements around it. Just like with plain NMS, we submit these samples as candidates to the QbC procedure, where in each iteration a large number of them is screened and a small number of those is selected to be included in the training set. Ab initio calculations are only required for these selected geometries. This enables the building of diverse training sets in which all representative structures that might be encountered in a future simulation are contained so that the resulting model is, in fact, able to accurately describe the PES along the whole MEP, even in regions that have negligible thermal populations. Similar to NMS, the computational cost of TTS is determined primarily by the MEP optimization procedure, the Hessian calculation, the C-NNP prediction required for screening, and the electronic structure calculations for the selected geometries. In general, this can be expected to be substantially less computationally demanding than executing direct, or even enhanced sampling, classical, or path integral AIMD simulations and sampling from their trajectories.

3. COMPUTATIONAL DETAILS

3.1. Ab Initio Electronic Structure. Two different levels of electronic structure theory were used in the simulations presented in this work. In both cases, we used the implementation provided by the CP2K software package⁴⁸ with its Quickstep DFT module.^{49,50} We described the electronic structure of the benzene molecule in the gas phase at the self-consistent charge density-functional tight binding⁵¹ (SCC-DFTB) level with third-order diagonal corrections. The system was enclosed in a 10 Å wide cubic box with open boundary conditions. For malonaldehyde and DABQDI systems, we used the revPBE0-D3 hybrid density functional^{52–55} combined with the TVZ2P Gaussian basis set^{49,56,57} to represent the molecular orbitals and a plane wave basis with a 600 Ry cutoff to represent the density. The core electrons of the heavy atoms were represented using Goedecker–Tetter–Hutter pseudopotentials.⁵⁸ In addition, we used the auxiliary density matrix method⁵⁹ with the cpFIT3

fitting basis set for the DABQDI molecule. Both systems using hybrid DFT were centered in a 15 Å wide cubic box with open boundary conditions and the wavelet Poisson equation solver.

3.2. C-NNP Model Generation. Throughout all of our investigations, we used committees consisting of eight different Behler–Parrinello NNPs,¹⁶ where for each of them, a different initialization of weights and a different 90% subset of the full training data set was used to ensure a diverse committee. The models consisted of two hidden layers of 20 nodes each and were trained using the multistream⁶⁰ adaptive extended Kalman^{61,62} filter with 32 streams. The input features were a standard set of atom-centered symmetry functions.²⁶ The training of the individual models was done using the n2p2 package⁶⁰ and the selection of training structures by QbC was done with a development version of our AML package⁶³ following the procedures outlined in refs 26 and 27. For benzene, 20 structures were randomly sampled initially and in each of the 40 QbC iterations, 10 new structures with the highest committee force disagreement were added to the data set, for a total of 420 training geometries. The final NNPs were trained for 2000 epochs. For the first generation of malonaldehyde and DABQDI models, 20 initial structures were sampled randomly and then 15 structures were selected and added to the data set in each of the 40 QbC iterations, for a total of 620 training geometries. The selection process is based on the standard deviation of the force prediction of the individual committee members and in each iteration, a predetermined number of structures with the highest disagreement of the candidate structures is selected. In our case, adding 15 structures at each iteration forms a compromise between choosing only the optimal structure and limiting the number of QbC iterations (which each require training a new committee) necessary to get to a sufficient training set size. For malonaldehyde, where additional generations of models were required (as detailed in Section 4), the original training set was supplemented by additional structures QbC-sampled from an MD trajectory which was produced using the previous C-NNP model. Here, 15 structures were added in each iteration until the force committee disagreements for the selected structures and the remaining candidate structures were similar. Like in our previous work, we chose this stopping criterion over a predetermined force disagreement threshold because it works well on its own, whereas an additional calibration against the force error would have been necessary to determine a suitable threshold. All reference calculations were done using CP2K and the electronic structure settings described above.

3.3. Geometry Optimization and Vibrational Analysis. The optimization of the minimum reference geometries for benzene and DABQDI was executed natively in the CP2K software. It was performed using the BFGS optimizer⁶⁴ combined with threshold criteria of 0.07 eV Å⁻¹ for the maximum change in force components, 0.009 Å for the change in atomic positions, and 0.13 eV for the change in total energy. For the malonaldehyde molecule, we employed the Atomic Simulation Environment (ASE)⁶⁵ together with CP2K and performed the optimization using the FIRE optimizer⁶⁶ while specifying only a force criterion of 0.01 eV Å⁻¹. Additional constrained optimizations in the case of DABQDI needed for the relaxed PES scan were performed using the constraint functionality provided by ASE together with its FIRE optimizer. The Hessian matrix evaluation on the optimized structures was performed in each case using CP2K and a Cartesian atomic displacement of 0.0005 Å.

3.4. Nudged Elastic Band Calculations. The relevant MEPs needed for the TTS procedure were obtained through the climbing-image⁶⁷ nudged elastic band⁴⁶ (CI-NEB) optimization procedure as implemented in CP2K. The initial band geometries in this work consisted of 15 replicas of the molecule in question including the two fixed, preoptimized endpoints, and were obtained through linear interpolation. The spring constant of the harmonic links between the neighboring replicas was kept constant at the value of 4.86 eV Å⁻². We used a force convergence criterion of 0.007 eV Å⁻¹ and the minimization of the band energy was performed using a DIIS optimizer.

3.5. MD Simulations. All MD simulations involving both ab initio as well as C-NNP potentials³³ were run using the CP2K package. The simulations with the classical representation of the nuclei were propagated at a temperature of 300 K using a time step of 0.5 fs to numerically integrate the Langevin equation with the friction coefficient γ of 0.02 fs⁻¹ to achieve canonical sampling. The path integral simulations that include nuclear quantum effects were performed using imaginary-time ring polymers consisting of 64 replicas using the RPMD propagator. The canonical distribution at 300 K was sampled using the local path integral Langevin equation thermostat⁶⁸ (PILE-L) with the time constant for the centroid motion of 200 fs while the integration time step was kept at 0.25 fs.

3.6. Umbrella Sampling. The initial conditions for each umbrella sampling window were extracted from a steered MD trajectory, which was performed in the CP2K v2022.1 software package combined with the PLUMED plugin.^{69–71} In this case, the value of the proton-sharing coordinate δ_1 (as detailed in Section 4) was biased from -1.2 to 1.2 Å during a 10 ps long simulation using a moving harmonic restraint with the force constant κ of 500.0 kJ mol⁻¹ Å⁻². The simulation was performed classically with an integration time step of 0.5 fs in the canonical ensemble at 300 K using a local CSVR thermostat⁷² with a time constant of 50 fs.

30 equidistant umbrella sampling windows separated by 0.08 Å were set up from the above steered MD simulation. Individually in each window, the value of δ_1 was biased by a static harmonic restraint of 500.0 kJ mol⁻¹ Å⁻² and simulated for 50 ps using the same setup as for the steered MD simulation above. The overlap of the corresponding histograms of δ_1 values observed in each simulation window is shown in Section S2 of the Supporting Information. The value of δ_2 was kept unbiased in each simulation window but was monitored for use in the following analysis. The biased configurations were reweighted to the unbiased ensemble using a Python implementation of the multistate Bennet acceptance ratio^{73,74} (MBAR) procedure to obtain both a one-dimensional free energy profile for the proton-transfer along δ_1 as well as a two-dimensional free energy surface showing the dependence on both proton-sharing coordinates. This was done by determining the thermal weight associated with each configuration in the biased simulations and using these to obtain the probability distribution in the δ_1, δ_2 subspace, and from that the corresponding free energy surface.

4. RESULTS AND DISCUSSION

To showcase the performance of the TTS procedure in the creation of models for realistic potentials, we select three different gas-phase molecules with an increasing complexity of their PES. We begin with benzene, which represents a single-minimum system with a close-to-harmonic potential at room

temperature and thus allows us to illustrate the simple thermal NMS procedure. This is followed by a study of the enol form of 1,3-propanedial (malonaldehyde), which exhibits reactivity by sharing the acidic proton between the two oxygen atoms spontaneously at ambient conditions.⁷⁵ Finally, we focus on a more involved proton-sharing system represented by 2,5-diaminobenzoquinone-1,4-diimine (DABQDI), which has two proton-sharing sites.⁷⁶ Spontaneous proton transfer is hindered by a barrier thermally insurmountable at room temperature, and an enhanced sampling simulation is necessary to determine the free energy profile.

4.1. Benzene. To lead off the discussion of the ability of TTS to seed a training set for the creation of C-NNP models for realistic systems in the gas phase, we focus on the benzene molecule. It represents an ideal example to illustrate the basic idea of thermal NMS using a single normal mode expansion at an optimal geometry since it features a single configurational minimum and the surrounding PES exhibits almost no anharmonic effects.

To prepare the ground for comparison with the relevant C-NNP data, we initially performed one 250 ps AIMD simulation of gas-phase benzene at 300 K at the DFTB level using a classical representation of the atomic nuclei as well as a 100 ps PIMD simulation using 64 replicas to approximate the imaginary time path. Two C-NNP models were then based on candidate sets obtained from a thermal NMS of gas-phase benzene using a Hessian matrix calculated at the same DFTB level of theory as the (PI)-AIMD simulations at 300 K for the classical model and with the appropriate effective temperatures at 300 K for the quantum one. The resulting models were evaluated on test sets consisting of 1000 structures sampled from the two AIMD trajectories. Both models performed very well with an energy root-mean-square error (RMSE) of 1.66 and 5.90 meV for the model constructed for the use without and with path integral structures, respectively. The RMSE for a single force component was 14.9 and 30.4 meV \AA^{-1} . Subsequently, the models were used to obtain new 500 ps long MD and 100 ps PIMD simulations at 300 K.

The comparison of the C-NNP models to the corresponding (PI)-AIMD trajectories in terms of molecular geometry properties is summarized in Figure 2. In general, we can see the expected broadening of probability distributions due to nuclear quantum effects when we compare the left and right columns of Figure 2. In both the classical and quantum case, we observe a perfect match between the ab initio (green shading) and C-NNP distributions (blue dashed lines) in C–C bond lengths (panels A and B), C–H bond lengths (panels C and D), C–C–C angles (panels E and F), and C–C–C–C dihedrals (panels G and H). The two types of covalent bonds have expected distributions; the mean of the C–C–C angle is located at 120° which shows the average hexagonal arrangement of the aromatic ring subject to planarity, which is, in turn, demonstrated by the (signed) C–C–C–C dihedral angle peaking at 0° as expected. This level of agreement suggests that the final models used for production MD represent excellent approximations of the original DFTB PES. The negligible deviations between the C-NNP and the (PI)-AIMD results are quantified by the differences shown in the small sub-panels in Figure 2 in blue. Additionally, we show the distributions of the NMS structures (orange dotted lines) alongside the anharmonic distributions. These exhibit significant overlap with both the (PI)-AIMD and C-NNP data. This suggests that the harmonic approximation to the original ensemble is

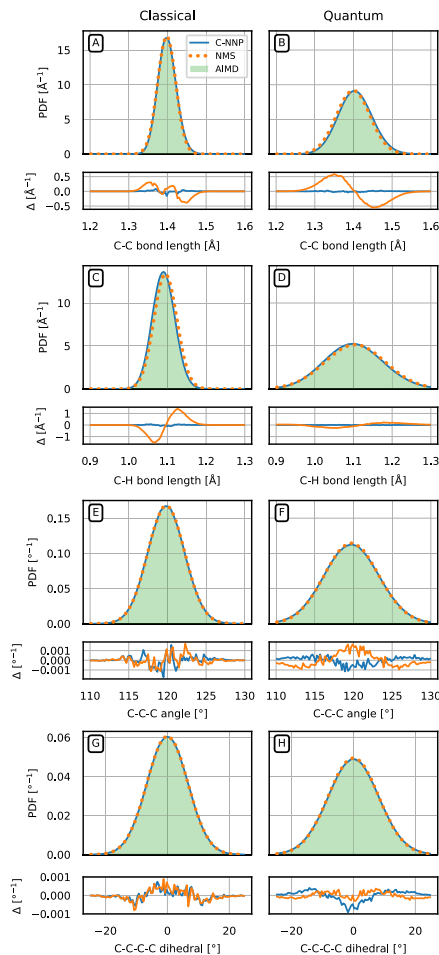


Figure 2. Thermal geometry properties of benzene in the gas phase at 300 K from classical MD (left column) and path integral MD (right column) compared between simulations using the reference DFTB potential, the harmonic TTS ensemble, and simulations using a C-NNP model building on the thermal NMS geometries. Panels A and B show the distribution of C–C bond lengths, panels C and D the distribution of C–H bond lengths, panels E and F the distribution of C–C–C angles, and, finally, panels G and H the distribution of the C–C–C–C dihedral angles. The smaller panels below each labeled panel show the deviations of the NMS and C-NNP data from the DFTB reference, using the same color coding as in panel A.

relatively good and confirms the assumed high degree of harmonicity of the 300 K gas-phase benzene PES, even in the quantum case. However, note that the match of the NMS data

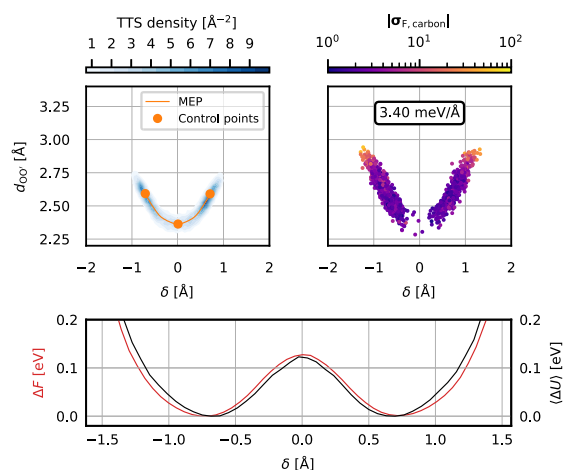


Figure 3. TTS sampling of the malonaldehyde proton transfer and the MD simulation with the resulting C-NNP model. The top left panel shows the relevant MEP (orange) with the three selected control points in the two configurational minima and the transition state highlighted and the distribution of the TTS geometries (blue). The top right panel shows a scatter plot of a subset of geometries (selected with a stride of 37.5 fs) obtained during a 250 ps MD simulation using the C-NNP model built on top of the TTS ensemble. Each point is colored by the norm of the force committee disagreement on the carbon atoms and the mean of the quantity is shown in the box. Note the high force disagreement in the high δ tails of the distribution. The bottom panel shows the Boltzmann-inverted free energy profile (red) and the corresponding binned average potential energy of the system (black, aligned to zero) along the proton transfer reaction as observed in the MD simulation. The error of the free energy, obtained by block averaging, is ≤ 2 meV, which roughly corresponds to the thickness of the red line.





Figure 4. Evolution of the force disagreement of the carbon atoms through multiple instances of QbC. The left panel shows a subset of configurations originating from an MD simulation using a model trained on data selected directly from the TTS candidate set (identical data as in the top right panel of Figure 3 are shown). The force disagreement (depicted using the color scale) in the vicinity of the two configurational minima and along the proton-sharing reaction is adequately low; however, for structures with a high absolute δ , it is more than 1 order of magnitude higher. The central panel shows configurations and disagreements obtained from an MD simulation using a model trained on the initial training data augmented by QbC-selected high-disagreement configurations from the data in the left panel. In turn, the right panel shows data obtained by improving the model using the new data sampled in the simulation shown in the central panel. Most notably, structures at the tails of the populated configuration space are substantially improved. The mean force disagreement over all configurations in each data set is shown in the framed box in each panel. The δ and $d_{\text{OO}'}$ coordinates are illustrated in the snapshot to the left of the panels.

with the (PI)-AIMD data is not nearly as perfect as that of the C-NNP data and certain deviations are, in fact, present. As discussed in Section 2, these are to be expected since the NMS ensemble is only auxiliary and its goal is to provide sufficient coverage of the accessible PES which ultimately leads to an accurate C-NNP model. The differences of the NMS ensemble from the ab initio reference are again quantified as differences in Figure 2. Using thermal NMS, we were thus able to construct a C-NNP that accurately describes the original PES

of benzene based on a single Hessian evaluation and 420 single-point electronic structure calculations.

4.2. Malonaldehyde. The enol form of malonaldehyde is a simple organic molecule that has been used in MD simulations⁵ to illustrate a simple intramolecular proton-transfer reaction



in which the proton is moved from the enol oxygen to the aldehyde oxygen with a simultaneous electronic rearrangement which altogether causes the reactant and the product to become symmetrically mirrored, chemically identical structures (eq 8). As such, malonaldehyde is a convenient molecule to demonstrate the ability of TTS to describe a simple reaction.

With the aim to describe the proton-sharing process accurately, we decided to model the original ab initio PES at the hybrid DFT level using the revPBE0-D3 functional. Thanks to the ability to use quantum normal coordinate distributions in the TTS method, we produced both classical and quantum models and classical and PIMD trajectories for malonaldehyde to test and showcase this functionality. However, we focus mostly on the classical case in the main text and discuss the complementary quantum results, for which qualitatively similar conclusions arise, in the [Supporting Information](#), Section S2. The starting point of the TTS procedure is the proton-sharing MEP, which was discretized into 15 replicas and optimized using the CI-NEB procedure and the revPBE0-D3 density functional. For illustration purposes, we decided to use the full MEP with both the reactant and product explicitly represented: this is strictly speaking not necessary since the symmetry of the reaction allows the use of only one nontrivial half of the MEP for TTS. Out of the optimized full-length MEP, three control points were selected in the two minima (reactant and product) and in the transition state. For the visualization of the multidimensional configuration data, we choose the reduction into a 2D space of two geometric parameters: the proton-sharing coordinate $\delta(\mathbf{R}) = |\mathbf{R}_{\text{O}} - \mathbf{R}_{\text{H}}| - |\mathbf{R}_{\text{O}'} - \mathbf{R}_{\text{H}}|$ and the oxygen–oxygen $d_{\text{OO}'}(\mathbf{R}) = |\mathbf{R}_{\text{O}} - \mathbf{R}_{\text{O}}'|$, where O and O' denote the two oxygen atoms that share the proton H. The obtained optimized MEP and the selected control points in this representation are shown in the top left panel of [Figure 3](#) in orange. The chosen parameters, illustrated in the snapshot on the left of [Figure 4](#), are not relevant for the execution of TTS itself, which takes place in the full dimension, but allow to conveniently show the results of the sampling in a reduced-dimensionality parameter space that is physically meaningful and suitable for the characterization of a proton transfer process. The TTS classical candidate structures were then generated using the procedure outlined in [Section 2](#) at the temperature of 300 K, linear sampling density along the MEP of $1 \times 10^5 \text{ \AA}^{-1}$, and matched-density sampling at the minima. The distribution of the obtained configurations is shown in the top left panel of [Figure 3](#) as blue contours. The same distribution colored by the assignment of each candidate to the individual control points (corresponding to the situation shown in panel D of [Figure 1](#)) is shown in [Section S2](#) of the [Supporting Information](#). Note that this particular presentation of the data does not do justice to the uniformity of the sample distribution along the MEP as the regions around the minima seem to be more populated than the transition state. This is an effect of the deformation of the configuration space by the projection on the selected subspace; the samples are distributed uniformly in the full dimension. After passing the resulting set of candidates through the QbC selection and training a C-NNP model on the obtained training set, the model was used to run a direct 250 ps MD simulation of gas-phase malonaldehyde at 300 K. A subset of the obtained configurations is shown in the form of a scatter plot in the top right panel of [Figure 3](#) colored by the decadic logarithm of the norm of the force committee disagreement on carbon atoms in

the usual δ and $d_{\text{OO}'}$ representation. Additionally, a 1D free-energy profile obtained by a Boltzmann inversion of the probability density of configurations along the δ -axis is shown in the bottom panel of [Figure 3](#); the height of the barrier is approximately 120 meV which corresponds to roughly $5 k_{\text{B}}T$ at 300 K. This accounts for the expected low, but existing population surrounding the transition state at $\delta = 0 \text{ \AA}$. We estimated the error of the free energy by the block-averaging method followed by extrapolation to infinite block size. This gives errors lower than 2 meV over the studied range of δ , which corresponds to the fact that the transition is sampled often during the simulation and the fact that the raw profile is already symmetric. Alongside the free energy profile, we show the corresponding average potential energy as a function of δ .

An important observation can be made from the presented data. By comparison of the two distributions in the top panels of [Figure 3](#), it is clear that the TTS distribution populates a smaller volume of the configuration space than the data obtained from the MD simulation. In this particular case, it means that TTS does not directly provide good enough coverage and the resulting model is undertrained in the missing, yet thermally accessible regions. Specifically, the C-NNP model performs poorly in the large $d_{\text{OO}'}$ tails of the shown distribution, as quantified by the larger force disagreement values. On the contrary, in regions around the proton-sharing MEP, the coverage is good and the force disagreement remains small, despite the tiny thermal population. Regardless of the elevated model uncertainty for some configurations, these MD simulations remain stable. The observed increased disagreement can be interpreted in the following way: going in the opposite direction from the minima as the proton-sharing MEP, the true anharmonic PES has a potential wall that grows slower than the harmonic wall captured by the TTS distribution and, therefore, the thermal coverage of the TTS configurations cannot reach far enough. In principle, this behavior is caused by either the strongly anharmonic character of the chemical bonds leading to bond dissociation or, more likely in this case, the presence of another reactive process leading to a new transition state. In the following two paragraphs, we present two possible solutions to the issue. The first one relies on an active-learning-based iterative improvement of the model which has the advantage of requiring no knowledge of the origin of the anharmonicity but is tedious to perform since several intermediate simulations and model generations need to be created. Meanwhile, the other solution relies on the chemical intuition of the user to identify the reactive nature of the issue with the aim to extend the initial TTS, which leads to a fully capable C-NNP model straight away.

The QbC process can be used to fill in an already existing training data set that has gaps, perhaps due to an incomplete TTS candidate set in the QbC selection for the initial model. We illustrate this process in [Figure 4](#), where the left panel shows the same data as the top right panel of [Figure 3](#) as a starting point. Regions of configuration space not covered well in the training set of this generation 1 model can be easily identified by the high committee disagreement, as can be seen in the tails of the distribution. Hence, one can start a new QbC using the existing training data set augmented by structures from an MD simulation performed with the initial C-NNP. Depending on the size of the gaps in the initial training data set, only a few iterations of QbC are typically necessary. However, adding these structures to the training data set can

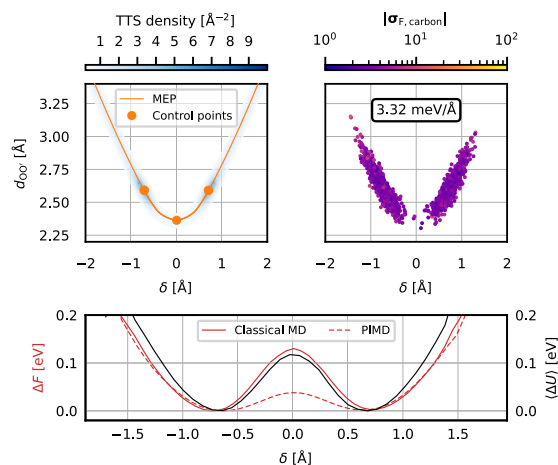



Figure 5. TTS sampling of the extended malonaldehyde MEP containing the proton transfer reaction as well as the single C–C bond torsion and the MD simulation with the resulting C-NNP model. Identical quantities as in Figure 3 are shown with the free energy being plotted here for both classical (solid line) and path-integral (dashed line) data. The orange curve shown in the top left panel is a union of the MEPs corresponding to the proton transfer and the C–C bond torsion; the projection of the latter into the δ , $d_{\text{OO}'}$ subspace is not, strictly speaking, a physically meaningful concept, but clearly visualizes the fact that MEP is the continuation in the desired direction. As such, the C-NNP model trained on the combined TTS structures has no further deficiencies as shown by the overall uniform force disagreement in the top right panel.

lead to substantial changes in the previously inaccurate regions of the PES, resulting in an MD simulation that again reaches new regions of the configuration space where the shape of the PES is yet unknown to the model and the committee disagreement is high. This can be seen in the generation 2 model shown in the middle panel of Figure 4. Therefore, multiple repetitions of the MD–QbC cycle might be necessary until a highly accurate model that exhibits low and uniform disagreements over the sampled data is reached, as is the case for the generation 3 model in the right panel of Figure 4. As such, the approach could become practically cumbersome when the regions of high disagreement coincide with regions of high free energy and long MD simulations are needed to uncover these structures, but nonetheless represents a general solution to the anharmonicity problem. Overall, repeating the cycle of sampling MD configurations with a given generation of a C-NNP model followed by training a new generation on a training set enhanced by high-disagreement QbC-selected structures from the previous MD simulations leads to a force disagreement that is lower in the problematic PES regions and, therefore, more uniform overall. In addition, we observe a decrease in the mean of the force disagreement of the sampled configurations due to the fact that the size of the training set increases in each generation. Specifically, 620 structures were used for the original model in the left panel of Figure 4, 800 structures for the model in the middle panel, and 950 structures for the last model in the right panel. This approach can be beneficial in systems where it is difficult to identify the origin of the anharmonicity of the original PES but is rather demanding from the point of view of both computational requirements and user involvement due to the need for the semi-supervised iterative procedure.

However, in the case of malonaldehyde, the general active learning iterative procedure to improve the model might be excessive. The possible reasons for the softer-than-harmonic wall due to conformational flexibility are few and the particular direction against which the first generation of the model is pushing can be easily identified with the *s*-cis and *s*-trans torsional isomerism



which is mediated by rotation around the C–C single bond in the propane backbone (eq 9). The optimized MEP corresponding to this torsion displays a perfect continuation in the correct direction when projected into the δ and $d_{\text{OO}'}$ subspace, as shown in the top left panel of Figure 5 in orange. Although this pair of descriptors is not appropriate for the whole torsion MEP, which entails a more complicated motion, it is accurate enough at small deviations from the equilibrium geometry. To include structures along this MEP into the initial (first generation) proton-sharing TTS, two control points were chosen in the minimum (shared by the two MEPs) and in the new transition state (not shown in Figure 5, as it is around $d_{\text{OO}'} = 3.8$ Å). We do not need to use the *s*-trans minimum at all, as we are not interested in including the transition itself, only the shape of the PES on the side of the global minimum. A new TTS was performed between these control points with the same parameters as the initial one and the resulting distribution of the combined sets of configurations is shown in blue in the top left panel of Figure 5. Running a 250 ps long MD simulation with a new C-NNP model trained on the QbC-selected training set from this combined candidate set leads to the distribution shown in the top right panel in Figure 5.




Figure 6. Comparison between the reference ab initio revPBE0-D3 and the C-NNP proton-sharing PESs of the DABQDI molecule. The left panel shows the projection of the reference DFT PES into the δ_1 , δ_2 subspace using a color scale as well as individual isoenergetic contours. Furthermore, the minimal nontrivial MEP which describes a single proton transfer is depicted in white with the selected control points highlighted. A sparse subset of the TTS geometries sampled around the MEP at 300 K is shown in green. The middle panel shows the corresponding PES projection calculated with the resulting C-NNP. Finally, the right panel shows the difference between the two PESs aligned to the global minima. The black contours range from $-1.5 k_B T$ to $1.5 k_B T$ and are spaced by $0.5 k_B T$ at 300 K. The two δ coordinates are illustrated in the snapshot of the DABQDI molecule to the left of the panels.


Clearly, the distribution reaches all the expected thermally accessible regions, the force disagreement is evened out across the configurations, and the high-disagreement tails are no longer present. The mean value of the disagreement is comparable to that of the first-generation model, as these two models are based on training sets of the same size. The bottom panel of Figure 5 shows the potential energy and the free energy profile along δ , where the latter is shown for both classical and path-integral data. In all three curves, one can observe a softening of the barrier in the high $|\delta|$ regions in comparison to the data shown in Figure 3 resulting from the present C-NNP being aware of the anharmonic nature of the PES in these regions. The transition-state free energy of 43 meV for the quantum simulation is ~ 4 times lower than in the classical case, which is a manifestation of proton tunneling through the barrier. This is qualitatively consistent with existing literature but quantitatively different from the results reported for the BLYP functional, where the classical free energy barrier is somewhat higher and the quantum effect is substantially smaller, decreasing the barrier by a factor of only ~ 2 .⁷⁵

For further insights, a test set independent of the training set data was created by generating 500 structures using TTS and sampling 500 structures from the classical MD trajectory shown in Figure 5, and evaluating their energies and forces with the revPBE0-D3 reference method. The generation 3 C-NNP of the iterative approach as well as the extended TTS C-NNP performs well with an energy RMSE of 1.80 and 3.44 meV and a force component RMSE of 18.3 and 24.4 meV \AA^{-1} , respectively. The slightly elevated RMSE of the extended TTS approach is due to the broader coverage of the training set. It includes a range of geometries along the C–C single bond torsion, even in regions that are not populated during MD, leading to a less dense coverage of the rest of the configuration space. The validation errors of the intermediate models of the iterative approach and the distribution of errors within configuration space are discussed in more detail in the Supporting Information, Section S2.

Both of the approaches above thus yield highly accurate models for the description of the proton-sharing reaction in

malonaldehyde for classical and quantum nuclei. The advantage of one over the other therefore depends mostly on the specific situation in which the need for any of them should arise: if the reaction coordinate of the complementary transition can be identified, then the latter approach using the extended TTS reaches the desired result with higher efficiency. Note that this approach can also be used when multiple different transitions are to be included in a single model. Finally, it is worth noting that an MLP trained on structures selected only from the quantum formulation of TTS and PIMD trajectories performs well for classical MD simulations of malonaldehyde, too, as detailed in Section S2 of the Supporting Information.

4.3. DABQDI. The most complex reactive system used to demonstrate the performance of the TTS method is represented by the DABQDI molecule. This nitrogenated benzoquinone derivative can exchange two protons between the neighboring amine and imine groups



again accompanied by an electronic rearrangement that maintains the π -electron conjugation throughout the process (eq 10). However, this time, the proton-sharing does not take place at ambient conditions, which suggests high barriers to the process.

The corresponding PES reduced to the relevant δ_1 and δ_2 subspace (illustrated in the snapshot in Figure 6), where each proton-sharing coordinate describes a single proton-sharing site, was obtained at the revPBE0-D3 level of electronic structure theory through a relaxed scan of the molecular potential energies while applying appropriate constraints and is shown in the left panel of Figure 6. The shown data was aligned so that the global minimum of the PES corresponds to the zero-energy level. The typical structure of the PES featuring four distinct configurational minima and four transition states corresponds to a sequential double proton transfer at the level of an MEP. Here, one proton is first fully exchanged to reach an intermediate state located at a higher

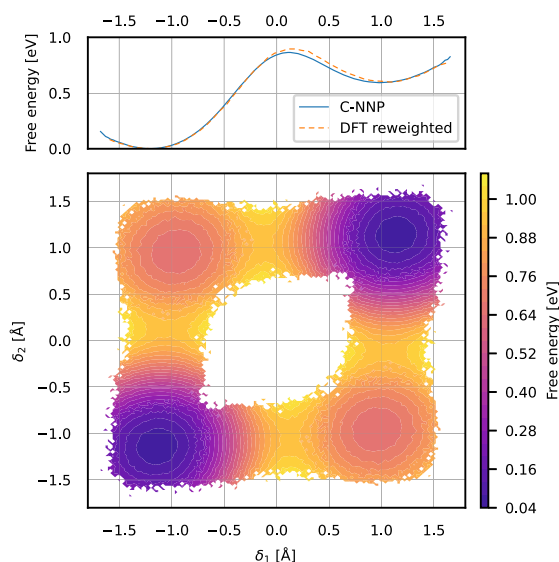


Figure 7. Umbrella sampling simulation of the single proton transfer in the DABQDI molecule along the δ_1 collective variable using the C-NNP potential. The top panel shows the obtained free energy profile in blue. For validation purposes, we also show the DFT free energy profile (orange, dashed) obtained by reweighting the C-NNP configurations as described in the text. Note that no umbrella sampling using the DFT potential was performed to obtain the DFT free energy profile. The bottom panel shows the full 2D free energy surface obtained by weighting the distribution in the two proton-sharing coordinates by the thermal Boltzmann factors extracted from the biased simulation and symmetrizing the resulting histogram.

potential energy and only then the other proton follows. The height of the potential energy barrier for this sequential process of roughly 0.8 eV indicates that its thermal rate should be negligible. The alternative concerted proton transfer path that is seen in other species including carboxylic acid dimers⁷⁷ is classically disallowed in this case by a tall (>1.2 eV) potential barrier in the middle of the presented PES which represents a second-order saddle point and, as such, no MEP can go through it. The symmetry of the DABQDI PES allows us to explicitly address only a single proton transfer: unlike in the previous example, we exploit this feature here for the C-NNP model generation. The relevant nontrivial MEP connecting the two chemically distinct minima was obtained using the CI-NEB optimization at the revPBE0-D3 level of theory and is shown in white in the left panel of Figure 6. From there, the straightforward TTS candidate generation was performed using the three usual control points in the two minima and in the transition state at 300 K with the linear sampling density of $1 \times 10^3 \text{ \AA}^{-1}$. A subset of the candidate geometries is shown in the left panel of Figure 6 as green points. The obtained C-NNP model was used to recreate the 2D proton transfer PES which is shown in the middle panel of Figure 6 with the energies aligned in the same way as in the DFT case. Qualitatively, the C-NNP model captures all the features of the original ab initio PES including the position of the minima, the transition state, and the barriers, as well as the potential energy values. Note that the good agreement in the representation of the central barrier in spite of the lack of corresponding

geometries in the TTS candidates is due to the successful extrapolation by the model. The quantitative difference between the ab initio and C-NNP potential energy landscape is captured in the right panel of Figure 6 which shows the difference between them relative to the global minima, with contours in multiples of $0.5 k_B T$ for reference. It is important to view this deviation in the context of the height of the barrier, which is 0.71 eV ($27.3 k_B T$, deviation of $\sim -1.25 k_B T$), and the relative energy of the two minima, which is 0.45 eV ($17.3 k_B T$, deviation of $\sim 0.5 k_B T$). A test set calculated over the TTS auxiliary ensemble quantifies the error in the thermal vicinity of the reaction coordinate to 0.88 meV/atom for energies and 47.6 meV/Å for forces. These errors are comparable to test set errors of accurate nonreactive models for water using the same NNP architecture.^{27,60} This difference could be decreased further, if desired, by optimizing hyperparameters of the model, by completely changing the architecture of the MLP, or by increasing the size of the training set beyond the current intentionally rather small set of 620 structures.

Since DABQDI features barriers that are not practically accessible by direct MD, it serves as a useful example to illustrate the power of the TTS-based C-NNP model to perform an enhanced sampling calculation to correctly estimate the free energy profile of the double proton transfer at 300 K. This was obtained using an umbrella sampling simulation in the coordinate δ_1 with the C-NNP PES (as described in Section 3) followed by a multistate Bennett acceptance ratio (MBAR) reweighting of the biased config-

urations. The obtained 1D free energy profile in δ_1 is shown in blue in the top panel of Figure 7. The transition state is located at roughly 0.8 eV above the global minimum. Comparing this with the value of the corresponding potential energy suggests that the entropic contribution in the gas-phase system is small and that the population at the barrier is clearly negligible at 300 K. The maximum error of the free energy profile is ± 2 meV, which was estimated by the error analysis infrastructure of the PyMBAR implementation of MBAR⁷³ (shown in Figure S6 of the Supporting Information). To validate the obtained free energy profile, we perform a reweighing of a subset of the obtained configurations in each umbrella window to the original DFT ensemble. This is achieved by additionally multiplying each MBAR-obtained unbiased weight by the corresponding factor $e^{-\beta\Delta E}$, where the energy difference ΔE is the difference between the C-NNP and DFT potential energy for each configuration. For this purpose, we used a total of 3000 configurations obtained by selecting 100 geometries evenly spaced in time from each of the 30 umbrella sampling windows. The resulting profile, which is a good approximation to the full-DFT free energy profile, is displayed as the orange dashed line in Figure 7 and shows very good correspondence with the profile obtained using the C-NNP model alone. This procedure thus at the same time validates the C-NNP model and provides DFT data for a fraction of the cost of the hypothetical purely ab initio enhanced sampling simulation. Monitoring the values of the collective variable δ_2 along the umbrella sampling simulation and using the thermal weights obtained from the MBAR treatment of the biased simulations also allows for recovering the 2D free energy surface in δ_1 and δ_2 . Its symmetrized version is shown in the bottom panel of Figure 7.

4.4. Computational Efficiency. The computational speedup due to using TTS rather than the traditional sampling based on AIMD simulations is not straightforward to measure as it depends on the particular system in question. At least, one can give a rough estimate of the order of magnitude of the difference between the number of single-point evaluations required by both approaches. For TTS, its dominant computational requirements arise due to the parts of the algorithm that require ab initio calculations: the MEP optimization, the calculation of the Hessians at the chosen control points, and the evaluation of training energies and forces during QbC. Therefore, the total number of single-point energy and force evaluations is

$$N_{\text{SP}} = N_{\text{SP}}^{\text{MEP}} + N_{\text{SP}}^{\text{Hessian}} + N_{\text{SP}}^{\text{QbC}} \quad (11)$$

where

$$N_{\text{SP}}^{\text{Hessian}} = 6NN_{\text{p}} \quad (12)$$

when analytic forces are available. For the MEP optimization, a reasonable expectation is an NEB with low tens of replicas that converges to the desired path on the order $\sim 10^2$ steps, which means the total number of single-point evaluations $N_{\text{SP}}^{\text{MEP}}$ on the order of 10^3 . For the Hessians, we typically require three control points (maybe five in more challenging cases not encountered in this work), which yields $N_{\text{SP}}^{\text{Hessian}}$ on the order of 10^2 . Similarly, for the QbC we find that a converged training set size is as small as several hundred thermal geometries. In total, the number of single-point evaluations required for TTS is thus on the order of 10^3 . For simple NMS, where the MEP is replaced by optimization of a single structure to potential

energy minimum and where only one Hessian is required, we expect the number of single points to be an order of magnitude less than TTS. The number of single-point evaluations for the AIMD alternative is again not clear, as it depends on the system-dependent correlation time scale. For a small gas-phase molecule with a single minimum such as benzene, a reasonable (although still optimistic) estimate is that a 20 ps long classical direct AIMD trajectory with a 0.5 fs integration time step, which allows us to extract 400 geometries with a 50 fs stride, should be sufficient to provide decorrelated data for the training of a C-NNP. This already requires $\sim 10^4$ single-point evaluations. However, more complex systems typically require longer simulations to provide sufficient sampling. Crucially, reactivity increases the computational cost of a hypothetical reference AIMD simulation dramatically. Already in the reactive case of malonaldehyde, adequately sampling the barrier regions (Figure 5) required $\sim 10^5$ MD steps of direct sampling, which would have to be performed at the ab initio level without TTS. Systems with higher barriers that will not be crossed by direct AIMD on reasonable time scales would require the use of enhanced sampling techniques, and the construction of a model compatible with quantum nuclei would require path integral simulations. These would, again, raise the cost of the simulations by at least an order of magnitude each—think, for instance, low tens of umbrella sampling windows to cover a reaction coordinate and tens to hundreds of path integral replicas to converge quantum properties. All in all, we expect TTS to be always computationally superior to running AIMD simulations. Especially when enhanced sampling and path integral simulations are needed for an appropriate description of the studied system, we expect the difference between the methods to be three to 4 orders of magnitude.

Another facet of the considerations of computational efficiency is the simulation performance of the resulting C-NNP in comparison to the original ab initio electronic structure method itself, which we illustrate on the DABQDI enhanced sampling simulation, where the acceleration of the C-NNP umbrella simulation in comparison to the naive execution with the original DFT method is substantial. To illustrate the computational savings, we can compare the times required for one MD step with the implementations in CP2K used in this work. With the hybrid functional, one MD step takes 272 s on a single core or 17 s on 32 cores (a full node) of our EPYC-based cluster. With the C-NNP, one step takes 0.006 s and does not scale meaningfully to more cores due to the small system size. This yields a speedup of $\sim 45,000\times$ on identical resources or $\sim 2800\times$ with more resources given to the DFT calculation. Obviously, the specific numbers will depend on the details of the electronic structure setup and the MLP architecture used, as well as the specific implementations and hardware used, but this behavior of our particular setup should provide a general idea.

5. CONCLUSIONS

In this work, we have introduced the TTS method to sample thermal geometries around MEPs that describe barrier-crossing transitions in molecular systems. The goal of the method is to provide a physically meaningful set of candidate structures for the creation of MLPs without the need to run computationally demanding ab initio simulations. In our case specifically, we submit these geometries to QbC and construct a C-NNP model, but the same candidates could be used for

other types of models as well. The execution of the TTS protocol as a whole entails a relatively modest computational cost with respect to the original *ab initio* method that is given by the MEP optimization, several Hessian evaluations, and a small number of single-point *ab initio* calculations for the QbC-selected geometries. In terms of application to realistic systems, the TTS method yields highly accurate C-NNP models in all studied cases. This was achieved either by using the generated candidate set directly or by letting the resulting C-NNP model undergo additional active learning generations to compensate for a pronounced anharmonic effect as seen in the case of malonaldehyde. As such, the performance of TTS in the presented test systems demonstrates its robustness and efficiency and suggests applicability in most gas-phase systems, including highly anharmonic cases.

A noteworthy feature of the TTS method is its ability to provide thermal geometries sampled from the quantum thermal distribution at essentially the computational cost of the classical case. As such, models that are appropriate for use in path integral simulations are made readily available without the need to run expensive PI-AIMD simulations at all. However, it is important to recognize that although the present formulation of TTS can address quantum behavior, it has limitations in this regard that derive from the fundamentally classical nature of the MEP. Nuclear quantum effects, in particular quantum tunneling through the potential barrier, can cause the configuration-space probability density of the system to deviate from the transition tube around the MEP in a way that renders the coverage by TTS samples insufficient.

To account for this, the above formulation of TTS can be straightforwardly generalized from sampling around classical MEPs to ring-polymer instantons,⁷⁸ which represent the paths of optimal tunneling. While this modification requires essentially no adaptation of the TTS theory and implementation itself, one can anticipate an elevated computational cost due to the required instanton optimization at the explicit *ab initio* level. The approach will find applications beyond the gas phase, in systems where vibrational normal modes are a meaningful concept, such as in the study of materials, molecular crystals, or in surface science for the description of growth and molecular adsorption. Disordered condensed phase, including liquids, represents a more challenging case in which TTS alone is not applicable for efficient thermal sampling of geometries. However, the auxiliary use of the TTS protocol in obtaining more diverse thermal structures of liquids, for instance with the help of local normal modes, should be explored. Our research anticipates the need to address some of these condensed-phase systems in the near future and we expect TTS to be a valuable tool in the creation of accurate, yet computationally accessible potentials that will enable the accurate description of these more complex systems at unprecedented sizes and simulation time scales.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.3c00391>.

Additional theory details, additional results, and C-NNP validation (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Ondrej Marsalek – Charles University, Faculty of Mathematics and Physics, 121 16 Prague 2, Czech Republic; orcid.org/0000-0002-8624-8837; Email: ondrej.marsalek@mff.cuni.cz

Authors

Krystof Brezina – Charles University, Faculty of Mathematics and Physics, 121 16 Prague 2, Czech Republic; orcid.org/0000-0003-0285-1282

Hubert Beck – Charles University, Faculty of Mathematics and Physics, 121 16 Prague 2, Czech Republic; orcid.org/0000-0002-7191-2192

Complete contact information is available at: <https://pubs.acs.org/10.1021/acs.jctc.3c00391>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

The authors thank Mariana Rossi for her helpful comments and discussions. K.B. and H.B. acknowledge funding from grant schemes at Charles University, reg. n. CZ.02.2.69/0.0/0.0/19_073/0016935, and from the IMPRS for Quantum Dynamics and Control. O.M. acknowledges support from the Czech Science Foundation, project no. 21-27987S.

■ REFERENCES

- (1) Tian, P. Molecular dynamics simulations of nanoparticles. *Annu. Rep. Prog. Chem., Sect. C: Phys. Chem.* **2008**, *104*, 142–164.
- (2) De Vivo, M.; Masetti, M.; Bottegoni, G.; Cavalli, A. Role of Molecular Dynamics and Related Methods in Drug Discovery. *J. Med. Chem.* **2016**, *59*, 4035–4061.
- (3) Hollingsworth, S. A.; Dror, R. O. Molecular Dynamics Simulation for All. *Neuron* **2018**, *99*, 1129–1143.
- (4) Deringer, V. L.; Caro, M. A.; Csányi, G. Machine Learning Interatomic Potentials as Emerging Tools for Materials Science. *Adv. Mater.* **2019**, *31*, 1902765.
- (5) Yao, N.; Chen, X.; Fu, Z. H.; Zhang, Q. Applying Classical, *Ab Initio*, and Machine-Learning Molecular Dynamics Simulations to the Liquid Electrolyte for Rechargeable Batteries. *Chem. Rev.* **2022**, *122*, 10970–11021.
- (6) Marx, D.; Hutter, J. *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods*; Cambridge University Press: New York, 2009.
- (7) Szabo, A.; Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*; Dover Publications, Inc.: Mineola, N.Y., 1996.
- (8) Parr, R. G.; Yang, W. *Density-Functional Theory of Atoms and Molecules*; Oxford University Press: New York, 1989.
- (9) Hohenberg, P.; Kohn, W. Inhomogeneous Electron Gas. *Phys. Rev.* **1964**, *136*, B864–B871.
- (10) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133–A1138.
- (11) Marsalek, O.; Chen, P. Y.; Dupuis, R.; Benoit, M.; Méheut, M.; Bačić, Z.; Tuckerman, M. E. Efficient calculation of free energy differences associated with isotopic substitution using path-integral molecular dynamics. *J. Chem. Theory Comput.* **2014**, *10*, 1440–1453.
- (12) Spura, T.; Elgabarty, H.; Kühne, T. D. On-the-fly coupled cluster path-integral molecular dynamics: impact of nuclear quantum effects on the protonated water dimer. *Phys. Chem. Chem. Phys.* **2015**, *17*, 14355–14359.
- (13) Kapil, V.; VandeVondele, J.; Ceriotti, M. Accurate molecular dynamics and nuclear quantum effects at low cost by multiple steps in

real and imaginary time: Using density functional theory to accelerate wavefunction methods. *J. Chem. Phys.* **2016**, *144*, 054111.

(14) Behler, J. Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.* **2011**, *13*, 17930–17955.

(15) Friederich, P.; Häse, F.; Proppe, J.; Aspuru-Guzik, A. Machine-learned potentials for next-generation matter simulations. *Nat. Mater.* **2021**, *20*, 750–761.

(16) Behler, J.; Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.* **2007**, *98*, 146401.

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

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

(19) Schütt, K. T.; Kindermans, P. J.; Sauceda, H. E.; Chmiela, S.; Tkatchenko, A.; Müller, K. R. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 992–1002.

(20) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smid, T. E.; Kozinsky, B. E. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **2022**, *13*, 2453.

(21) Dral, P. O. Quantum Chemistry in the Age of Machine Learning. *J. Phys. Chem. Lett.* **2020**, *11*, 2336–2347.

(22) Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.* **2011**, *134*, 074106.

(23) Natarajan, S. K.; Behler, J. Neural network molecular dynamics simulations of solid-liquid interfaces: Water at low-index copper surfaces. *Phys. Chem. Chem. Phys.* **2016**, *18*, 28704–28725.

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

(25) Sivaraman, G.; Krishnamoorthy, A. N.; Baur, M.; Holm, C.; Stan, M.; Csányi, G.; Benmore, C.; Vázquez-Mayagoitia, A. Machine-learned interatomic potentials by active learning: amorphous and liquid hafnium dioxide. *npj Comput. Mater.* **2020**, *6*, 104.

(26) Schran, C.; Thiemann, F. L.; Rowe, P.; Müller, E. A.; Marsalek, O.; Michaelides, A. Machine learning potentials for complex aqueous systems made simple. *Proc. Natl. Acad. Sci. U.S.A.* **2021**, *118*, No. e2110077118.

(27) Schran, C.; Brezina, K.; Marsalek, O. Committee neural network potentials control generalization errors and enable active learning. *J. Chem. Phys.* **2020**, *153*, 104105.

(28) Artrith, N.; Behler, J. High-dimensional neural network potentials for metal surfaces: A prototype study for copper. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2012**, *85*, 045439.

(29) Gastegger, M.; Behler, J.; Marquetand, P. Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.* **2017**, *8*, 6924–6935.

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

(31) Hansen, L. K.; Salamon, P. Neural Network Ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.* **1990**, *12*, 993–1001.

(32) Cooper, A. M.; Hallmen, P. P.; Kästner, J. Potential energy surface interpolation with neural networks for instanton rate calculations. *J. Chem. Phys.* **2018**, *148*, 094106.

(33) Schran, C.; Behler, J.; Marx, D. Automated Fitting of Neural Network Potentials at Coupled Cluster Accuracy: Protonated Water Clusters as Testing Ground. *J. Chem. Theory Comput.* **2020**, *16*, 88–99.

(34) Krogh, A.; Vedelsby, J. Neural Network Ensembles, Cross Validation, and Active Learning. *Adv. Neural Inf. Process. Syst.* **1995**, *7*, 231–238.

(35) Seung, H. S.; Opper, M.; Sompolinsky, H. Query by committee. *Proceedings of the Fifth Annual Workshop on Computational Learning Theory-COLT'92*: New York, New York, USA, 1992; pp 287–294.

(36) Chen, L.; Sukuba, I.; Probst, M.; Kaiser, A. Iterative training set refinement enables reactive molecular dynamics via machine learned forces. *RSC Adv.* **2020**, *10*, 4293–4299.

(37) Morawietz, T.; Singraber, A.; Dellago, C.; Behler, J. How van der Waals interactions determine the unique properties of water. *Proc. Natl. Acad. Sci.* **2016**, *113*, 8368–8373.

(38) Musil, F.; De, S.; Yang, J.; Campbell, J. E.; Day, G. M.; Ceriotti, M. Machine learning for the structure-energy-property landscapes of molecular crystals. *Chem. Sci.* **2018**, *9*, 1289–1300.

(39) Vandermause, J.; Torrisi, S. B.; Batzner, S.; Xie, Y.; Sun, L.; Kolpak, A. M.; Kozinsky, B. On-the-fly active learning of interpretable Bayesian force fields for atomistic rare events. *npj Comput. Mater.* **2020**, *6*, 20.

(40) Wang, W.; Yang, T.; Harris, W. H.; Gómez-Bombarelli, R. Active learning and neural network potentials accelerate molecular screening of ether-based solvate ionic liquids. *Chem. Commun.* **2020**, *56*, 8920–8923.

(41) Miksch, A. M.; Morawietz, T.; Kästner, J.; Urban, A.; Artrith, N. Strategies for the construction of machine-learning potentials for accurate and efficient atomic-scale simulations. *Mach. Learn.: Sci. Technol.* **2021**, *2*, 031001.

(42) Schwalbe-Koda, D.; Tan, A. R.; Gómez-Bombarelli, R. Differentiable sampling of molecular geometries with uncertainty-based adversarial attacks. *Nat. Commun.* **2021**, *12*, 5104–5112.

(43) Vandenhoute, S.; Cools-Ceuppens, M.; DeKeyser, S.; Verstraelen, T.; Van Speybroeck, V. Machine learning potentials for metal-organic frameworks using an incremental learning approach. *npj Comput. Mater.* **2023**, *9*, 19.

(44) Rupp, M.; Ramakrishnan, R.; Von Lilienfeld, O. A. Machine Learning for Quantum Mechanical Properties of Atoms in Molecules. *J. Phys. Chem. Lett.* **2015**, *6*, 3309–3313.

(45) Schreiner, M.; Bhowmik, A.; Vegge, T.; Busk, J.; Winther, O. Transition1x - a dataset for building generalizable reactive machine learning potentials. *Sci. Data* **2022**, *9*, 779.

(46) Jónsson, H.; Mills, G.; Jacobsen, K. W. Nudged elastic band method for finding minimum energy paths of transitions. *Classical and Quantum Dynamics in Condensed Phase Simulations*; World Scientific, 1998; pp 385–404.

(47) Ceriotti, M.; Bussi, G.; Parrinello, M. Nuclear Quantum Effects in Solids Using a Colored-Noise Thermostat. *Phys. Rev. Lett.* **2009**, *103*, 030603.

(48) Hutter, J.; Iannuzzi, M.; Schiffmann, F.; Vandevondele, J. CP2K: Atomistic simulations of condensed matter systems. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2014**, *4*, 15–25.

(49) Vandevondele, J.; Krack, M.; Mohamed, F.; Parrinello, M.; Chassaing, T.; Hutter, J. Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach. *Comput. Phys. Commun.* **2005**, *167*, 103–128.

(50) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khalullin, R. Z.; Schütt, O.; Schiffmann, F.; et al. CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **2020**, *152*, 194103.

(51) Porezag, D.; Frauenheim, T.; Köhler, T.; Seifert, G.; Kaschner, R. Construction of tight-binding-like potentials on the basis of density-functional theory: Application to carbon. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1995**, *51*, 12947–12957.

(52) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **1996**, *77*, 3865–3868.

(53) Zhang, Y.; Yang, W. Comment on “generalized gradient approximation made simple”. *Phys. Rev. Lett.* **1998**, *80*, 890.

(54) Adamo, C.; Barone, V. Toward reliable density functional methods without adjustable parameters: The PBE0 model. *J. Chem. Phys.* **1999**, *110*, 6158–6170.

(55) Goerigk, L.; Grimme, S. A thorough benchmark of density functional methods for general main group thermochemistry, kinetics,

and noncovalent interactions. *Phys. Chem. Chem. Phys.* **2011**, *13*, 6670–6688.

(56) Guidon, M.; Schiffmann, F.; Hutter, J.; Vandevondele, J. Ab initio molecular dynamics using hybrid density functionals. *J. Chem. Phys.* **2008**, *128*, 214104–214115.

(57) Guidon, M.; Hutter, J.; Vandevondele, J. Robust periodic Hartree-Fock exchange for large-scale simulations using Gaussian basis sets. *J. Chem. Theory Comput.* **2009**, *5*, 3010–3021.

(58) Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space Gaussian pseudopotentials. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1996**, *54*, 1703–1710.

(59) Guidon, M.; Hutter, J.; Vandevondele, J. Auxiliary density matrix methods for Hartree-Fock exchange calculations. *J. Chem. Theory Comput.* **2010**, *6*, 2348–2364.

(60) Singraber, A.; Morawietz, T.; Behler, J.; Dellago, C. Parallel Multistream Training of High-Dimensional Neural Network Potentials. *J. Chem. Theory Comput.* **2019**, *15*, 3075–3092.

(61) Shah, S.; Palmieri, F.; Datum, M. Optimal filtering algorithms for fast learning in feedforward neural networks. *Neural Network.* **1992**, *5*, 779–787.

(62) Blank, T. B.; Brown, S. D. Adaptive, global, extended Kalman filters for training feedforward neural networks. *J. Chemom.* **1994**, *8*, 391–407.

(63) AML public GitHub repository. <https://github.com/MarsalekGroup/aml> (accessed March 20, 2023).

(64) Broyden, C. G. The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations. *IMA J. Appl. Math.* **1970**, *6*, 76–90.

(65) Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dulak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; et al. The atomic simulation environment - A Python library for working with atoms. *J. Phys. Condens. Matter.* **2017**, *29*, 273002.

(66) Bitzek, E.; Koskinen, P.; Gähler, F.; Moseler, M.; Gumbusch, P. Structural relaxation made simple. *Phys. Rev. Lett.* **2006**, *97*, 170201.

(67) Henkelman, G.; Uberuaga, B. P.; Jónsson, H. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *J. Chem. Phys.* **2000**, *113*, 9901–9904.

(68) Ceriotti, M.; Parrinello, M.; Markland, T. E.; Manolopoulos, D. E. Efficient stochastic thermostating of path integral molecular dynamics. *J. Chem. Phys.* **2010**, *133*, 124104.

(69) Bonomi, M.; Branduardi, D.; Bussi, G.; Camilloni, C.; Provasi, D.; Raiteri, P.; Donadio, D.; Marinelli, F.; Pietrucci, F.; Broglia, R. A.; Parrinello, M. PLUMED: A portable plugin for free-energy calculations with molecular dynamics. *Comput. Phys. Commun.* **2009**, *180*, 1961–1972.

(70) The PLUMED consortium; et al. Promoting transparency and reproducibility in enhanced molecular simulations. *Nat. Methods* **2019**, *16*, 670–673.

(71) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. *Comput. Phys. Commun.* **2014**, *185*, 604–613.

(72) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **2007**, *126*, 014101.

(73) Shirts, M. R.; Chodera, J. D. Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **2008**, *129*, 124105.

(74) WHAM. <https://github.com/apalath/WHAM> (accessed March 20, 2023).

(75) Tuckerman, M. E.; Marx, D. Heavy-Atom Skeleton Quantization and Proton Tunneling in "Intermediate-Barrier" Hydrogen Bonds. *Phys. Rev. Lett.* **2001**, *86*, 4946–4949.

(76) Cahlik, A.; Hellerstedt, J.; Mendieta-Moreno, J. I.; Švec, M.; Santhini, V. M.; Pascal, S.; Soler-Polo, D.; Erlingsson, S. I.; Výborný, K.; Mutombo, P.; Marsalek, O.; Siri, O.; Jelinek, P. Significance of Nuclear Quantum Effects in Hydrogen Bonded Molecular Chains. *ACS Nano* **2021**, *15*, 10357–10365.

(77) Smedarchina, Z.; Siebrand, W.; Fernández-Ramos, A.; Meana-Pañeda, R. Mechanisms of double proton transfer. Theory and applications. *Z. Phys. Chem.* **2008**, *222*, 1291–1309.

(78) Kästner, J. Theory and simulation of atom tunneling in chemical reactions. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2014**, *4*, 158–168.

C.3 | Paper III

Reprinted (adapted) with permission from Brezina, K., Jungwirth, P. and Marsalek, O.: Benzene radical anion in the context of the Birch reduction: When solvation is the key, *Journal of Physical Chemistry Letters*, 11, 15, 6032–6038, 2020. Copyright 2020 American Chemical Society.

Benzene Radical Anion in the Context of the Birch Reduction: When Solvation Is the Key

Krystof Brezina, Pavel Jungwirth, and Ondrej Marsalek*

Cite This: *J. Phys. Chem. Lett.* 2020, 11, 6032–6038

Read Online


ACCESS |

Metrics & More


Article Recommendations

Supporting Information

ABSTRACT: The benzene radical anion is an important intermediate in the Birch reduction of benzene by solvated electrons in liquid ammonia. Beyond organic chemistry, it is an intriguing subject of spectroscopic and theoretical studies due to its rich structural and dynamical behavior. In the gas phase, the species appears as a metastable shape resonance, while in the condensed phase, it remains stable. Here, we approach the system by *ab initio* molecular dynamics in liquid ammonia and demonstrate that the inclusion of solvent is crucial and indeed leads to stability. Beyond the mere existence of the radical anion species, our simulations explore its condensed-phase behavior at the molecular level and offer new insights into its properties. These include the dynamic Jahn–Teller distortions, vibrational spectra in liquid ammonia, and the structure of the solvent shell, including the motif of a π -hydrogen bond between ammonia molecules and the aromatic ring.



One of the most prominent roles of the benzene radical anion ($C_6H_6^{\bullet-}$) in organic chemistry is that of a reactive intermediate in the Birch process^{1–6} (eq 1) used to reduce benzene or related derivatives into 1,4-cyclohexadienes. Such reduction is realized through the action of excess electrons. Specifically, the reaction⁷



is initiated by introducing the aromatic substrate together with a proton source (usually a simple alcohol) into the dark blue solution of solvated electrons⁸ in liquid ammonia,^{9,10} created by dissolution of an alkali metal.¹¹ Already in the reaction's first step, the benzene radical anion is formed when a neutral benzene molecule accepts an excess electron from the surrounding solution, allowing the substrate to enter the reduction pathway. The Birch reduction has proved to be an important tool for selective reduction of organic substrates both industrially and in organic synthesis.^{12–14} However, interest in the anion reaches beyond organic chemistry. Its open-shell electronic structure defined by a highly symmetric molecular geometry gives rise to phenomena that have been the subject of numerous theoretical and spectroscopic studies since the 1950s.^{15–19} Computationally, the anion presents a challenge due to its need for both a large basis set and a sufficient level of electronic structure theory to accurately capture its properties in the gas phase. Recently, Bazante and co-workers²⁰ have published a detailed gas-phase *ab initio* study focusing on the isolated species, specifically its molecular

geometry, including the Jahn–Teller (JT) distortions, electron paramagnetic resonance (EPR) spectroscopy parameters, and on confirming its previously suggested¹⁵ character of a metastable shape resonance in vacuum. Experimentally, the character of the gas-phase benzene radical anion as a metastable resonance was studied using electron transmission spectroscopy^{21,22} where the results even determine the extremely short lifetime of the species on the order of 10 fs. In contrast, the species has been observed in the condensed phase using multiple experimental approaches. This includes spectroscopic studies of film-like co-deposits of an alkali metal with benzene using vibrational spectroscopy¹⁸ and in the presence of a polar organic solvent using ESR,^{16,23} or electronic spectroscopy.¹⁷ The issue of electronic stability in solvated environments was directly addressed experimentally in the context of small aqueous clusters,²⁴ and the thermodynamic equilibrium with solvated electrons was measured at various temperatures in a tetrahydrofuran solution.²⁵ These studies point to the presence of a radical species and suggest that its gas-phase metastability is removed in the condensed phase: a stable bound state is observed. Electrochemical properties and the reactivity of the radical anion in solution have also been studied by cyclic voltammetry in dimethoxyethane with a proton source.²⁶ The condensed-phase-induced

Received: May 15, 2020

Accepted: July 6, 2020

Published: July 6, 2020



stability is also implied by the high experimental yields^{1,27–29} of the Birch reduction that would not be achievable in case an intermediate was electronically unstable. This experimental evidence shows that this stabilization is not solvent specific but rather general to a broad class of polar solvents, implying that the stabilizing effect of the solvent is due to the dielectric environment that it creates around the anion. It is therefore clear that the solvent environment plays a crucial role in the stabilization of the benzene radical anion in solution.

Our aim in this work is to shed new light on this species using a computational approach. To this end, we use *ab initio* molecular dynamics (AIMD) to gain insight into and understand in detail the static and dynamic properties of the benzene radical anion in solution, in contrast with previous gas-phase studies, because experimentally the inclusion of the solvent appears to be the key to obtaining an electronically and thermodynamically stable species. Specifically, we perform our condensed-phase simulations in explicit liquid ammonia, as this is the natural environment for excess electrons in the context of the Birch reduction. Using these simulations, we characterize the JT distortions of the radical anion, explore the solvent structure of liquid ammonia around the anionic solute, and predict its contribution to vibrational spectra, comparing to its neutral counterpart in all cases.

The subsequent discussion of our results is based on AIMD simulations of the benzene radical anion and neutral benzene in liquid ammonia at 223 K under periodic boundary conditions using a hybrid density functional theory (DFT) electronic structure calculation on the fly. AIMD methodology at the hybrid DFT level is required to obtain an electronically stable anionic species with most of the excess electron density localized on the aromatic ring [around 95% as determined by the Hirshfeld population analysis³⁰ (see Figures S1 and S2)], as expected for an electron occupying a bound state in solution. This is necessary only due to the presence of the anionic radical species: liquid ammonia itself is already well described at lower levels of DFT theory.³¹ A detailed discussion of the employed methodology^{32,33} is presented in the Supporting Information; here, we show a snapshot illustrating the localization of the excess electron density (green and violet contours of the spin density) on the aromatic ring in Figure 1 and a trajectory video file (detailed in section S6) that shows also the time evolution of the system, including the spin density.




Figure 1. Representative snapshot from the hybrid DFT AIMD simulation. On the left, we show a snapshot of the whole simulation box containing one benzene, 64 ammonia molecules, and an excess electron with the spin density of the system colored green (positive part) and violet (negative part). The close-up shows a detail of the localization of the spin density distribution over the benzene radical anion.

With regard to the molecular structure of the radical anion, quantum theory predicts that a symmetry breaking of the initially perfectly hexagonal molecule of benzene is expected via a JT distortion after the addition of the excess electron to a degenerate lowest unoccupied electronic state. More precisely, the initial E_{2u} electronic state in the D_{6h} point group distorts to two possible distinct states labeled A_u and B_{1u} in the D_{2h} nearest lower-symmetry point group. A geometric relaxation of the molecular structure follows, resulting in two possible shapes that can be roughly described as contracted (B_{1u}) and elongated (A_u) hexagons that have been previously predicted by gas-phase calculations^{15,20} and experimentally.²³ For comparison, the related but computationally simpler benzene radical cation, which has almost identical distortions, does exist as a stable species already in the gas phase and has been studied by AIMD before.³⁴ The energy barriers between the two JT structures of the radical anion are expected to be low enough to allow for a dynamic JT effect where the structure is not rigid but undergoes a pseudorotation²⁰ (Figure 2 and the




Figure 2. JT pseudorotation of the benzene radical anion. The central hexagonal structure distorts to attain one of the lower-symmetry and lower-energy structures around the pseudorotational path around which it then circles.

video file detailed in section S6). Below, we analyze the structural distortions of the anion as observed in our condensed-phase simulations and discuss whether and how the JT behavior manifests in solution.

Initially, we analyze the structure of the solute by correlating pairs of carbon–carbon bond lengths around the aromatic ring. Specifically, we first correlate all pairs of directly neighboring bonds (1–2 correlation), all bond pairs with one additional bond in between (1–3 correlation), and finally all bond pairs neighboring over two extra bonds (1–4 correlation). The resulting bivariate probability densities are shown in Figure 3 for both neutral benzene (top row, blue) and benzene radical anion (bottom row, green).

The neutral benzene reference data show unimodal distributions in all three panels positioned on the main diagonal. This implies a hexagonal symmetry of the neutral solute and together with the mean carbon–carbon bond length of 1.39 Å suggests that its structure is very similar to that in the gas phase: this is expected as no symmetry-lowering distortions are anticipated at this point. In contrast, adding the excess electron to the solute changes these results dramatically. Bimodal distributions appear, and moreover, the 1–2 and 1–3 correlations now clearly show a negative correlation; on the contrary, the remaining 1–4 correlation displays a positive one. These observations are consistent with the predicted A_u and




Figure 3. Correlations of carbon–carbon bond lengths (d) between bonds of different neighboring schemes. The uncharged system data can be found in the top row (blue), whereas the bottom row contains the anionic data (green).

B_{1u} JT structures that require a bond length heterogeneity in the 1–2 and 1–3 pairs (i.e., having a shorter and a longer bond in the pair, respectively) and a homogeneity in the 1–4 pair where the bond lengths have to match due to the presence of a center of symmetry in the D_{2h} point group.

The analysis presented above using bond length correlations unarguably shows the presence of symmetry-lowering molecular distortions of the anion that agree with the previously predicted JT structures. However, as it averages over all bond pairs around the ring, it is in fact insensitive to distinguishing the individual pseudorotamers. Moreover, the bond lengths do not provide a meaningful reaction coordinate for evaluating the energetics of the distortions. Therefore, we present a complementary analysis where we project the immediate distortions of the solute onto the harmonic vibrational normal modes of the neutral benzene gas-phase reference and observe the correlations of degenerate pairs of modes of the JT active e_{2g} symmetry. Only one of the four possible JT active normal mode pairs yields a correlation that rises above thermal fluctuations, and thus, we track the JT distortions using this particular pair of modes, averaging over all remaining degrees of freedom (Figure 4). The pair in question describes the vertical elongation of the carbon ring and its sideways-skewing motion,³⁵ is located at a frequency of 1654 cm^{-1} in the gas phase, and will be denoted further as Q_+ and Q_- . Linear combinations of these two modes can reconstruct the crucial




Figure 4. Correlations of carbon ring-deforming JT active normal mode amplitudes for the neutral (left, blue) and anionic (right, green) systems.

features of the distortions sketched in Figure 2. A video file detailed in section S6 visualizes the propagation of the benzene ring along a circular trajectory in the subspace spanned by these two modes.

Once again, a trivial, origin-centered peak appears in the neutral system, which corresponds to a symmetric structure with respect to the two studied modes. Its narrowness implies a certain level of structural rigidity. In contrast, the anion displays a typical “Mexican hat” shape with the area around the center being depopulated and the main population residing in a ridge around it. This ridge directly corresponds to the pseudorotational path sketched in Figure 2. Such an observation is in agreement with the prediction of the JT effect and indeed shows that the structure of the radical anion is not rigid, but rather flexible and dynamic. As we discuss in the Supporting Information and show in Figure S5, the JT distortion changes on a very short time scale of no more than 100 fs. While the symmetric structure with respect to the two studied modes is penalized by a free energy barrier of around 2 kT, the free energy variation along the pseudorotational path turns out to be small relative to thermal fluctuations, as detailed in Figure S4 and the corresponding discussion in the Supporting Information.

Having addressed the features of the intrinsic structure of the solute, we now turn our attention to the structure of the solvent. The specific solvation structure is captured in the nitrogen spatial distribution functions (SDFs) around each of the two (anionic and neutral) solutes shown in Figure 5. The




Figure 5. Spatial distribution functions of ammonia nitrogen atoms around the center of mass of the aromatic ring. Anionic data are shown on the left (green solute), and neutral data on the right (blue solute). The blue–purple contour shows regions where the density is 4.5 times that of a non-interacting solvent. For the purpose of presentation, the three-dimensional histograms were convolved with a $\sigma = 1\text{ Å}$ Gaussian filter.

SDFs take on a cage-like form with maxima in strips around the carbon–carbon midbond regions and in the areas above and under the ring. Intuitively, these areas, clearly separated from the rest of the solvation patterns, can be interpreted as a manifestation of the π -hydrogen bonding phenomenon, as discussed below. The SDFs of both species look similar to the naked eye: the only visible difference is a more blurred character in the neutral benzene case, suggesting that the diffuse negative charge on the anion does provide a structuring effect on the solvent arrangement, but not a very substantial one. The JT distortions of the anion are not expected to alter the symmetry of the solvation structure: the distortions happen symmetrically in three different directions, which would lead to a cancellation of their effect on the shape of the SDF in the long time limit. Moreover, the solvent does not have enough time to rearrange and accommodate the immediate JT

distortions of the solute anyway due to their previously mentioned very short time scale.

In Figure 6, we plot the plane-restricted radial distribution functions (RRDFs) of ammonia nitrogen atoms around the

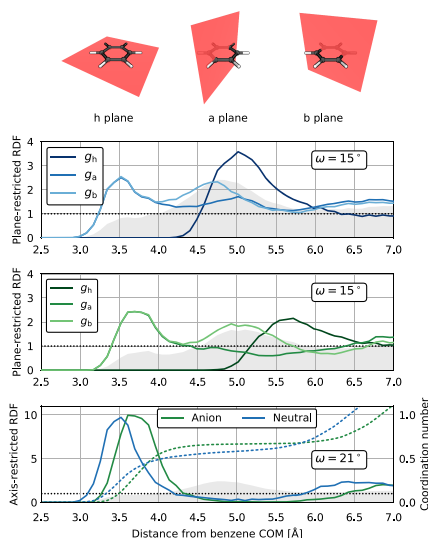


Figure 6. RRDFs of ammonia nitrogen atoms around the COM of the aromatic ring. RRDFs restricted to significant planes of the solute geometry (top row) under a restriction angle of $\omega = 15^\circ$ are shown in the top (neutral benzene) and middle (radical anion) panels. The gray shading represents the unrestricted RDF ($\omega = 90^\circ$) in both cases. In the bottom panel, axis RRDFs focus on the π -hydrogen cloud with a restriction angle of $\omega = 21^\circ$ for the radical anion (green) and neutral benzene (blue) with the gray shaded neutral benzene unrestricted RDF for size comparison. Dashed lines show the per-site running coordination numbers of the axis RRDFs.

center of mass (COM) of both the neutral and anionic solutes, obtained formally by partial angular integration of their respective SDFs (see the Supporting Information for the specifics of the procedure and Figure S3 for its convergence). These RRDFs describe the radial solvent distribution close to certain given planes instead of the whole sphere (as in the usual plain RDF). Therefore, they provide more detailed information about the anisotropic solvation environment than a plain unrestricted RDF does. In total, there are three categories of perpendicular significant planes with respect to the geometry of the studied solutes: the horizontal (h) plane in which the aromatic ring of ideal geometry would lie, three equivalent vertical planes (a) perpendicular to h cutting through a pair of distal carbon atoms, and three more equivalent vertical planes (b) perpendicular to both h and a and cutting through two opposing carbon-carbon bonds (Figure 6, top row). A comparison between the RRDFs in both presented panels further quantitatively confirms the previous qualitative assessment based on the SDFs that the solvent

structure around the anionic and the neutral structures is, despite minor differences, similar. We shall discuss mainly the anionic data (Figure 6, middle panel) while keeping in mind that a similar discussion applies to neutral benzene. In the g_h RRDF (dark green), which addresses the shape of the solvation structure around the horizontal plane, one can observe a rather late onset due to steric shielding by the solute itself followed by a single peak. The maximum of the first solvation shell of the benzene anion in the horizontal direction is located approximately 5.6 Å (5.0 Å for benzene) from the center of mass. The g_a RRDF (medium shade of green) shows a similar shape, even though the onset is at a much shorter distance in this case as there is no steric shielding by the solute in the vertical direction. The observed peak at 3.7 Å (3.5 Å for benzene) corresponds to the π -hydrogen cloud and is followed by a wide shallow minimum as there is no significant density of solvent around the carbon atoms. In contrast, the g_b RRDF (light green) shows a second additional peak at 5.0 Å (4.7 Å for benzene) arising from the SDF areas not belonging to the π -hydrogen cloud. Both peaks of the RRDF thus belong to the first solvation shell. This splitting slightly transfers to the g_h plane of neutral benzene (medium blue), which can be considered a sign of its more diffuse (weaker) solvation structure in comparison to that of the anion. However, we note in passing that the solvation cavity of the anion is systematically larger in all directions. The onset of the second solvation shell can be noticed in the vertical RRDFs above 6 Å but is cut short by the employed simulation box size of approximately 14 Å. The second solvation shell is not captured in the horizontal RRDF at all.

To close the discussion, we turn briefly to the phenomenon of π -hydrogen bonding, i.e., bonding between the polar solvent and the π -electron cloud of the solute. This behavior has been found in aqueous solutions of benzene both experimentally³⁶ and computationally.³⁷ In our liquid ammonia simulations, we observe a well-separated peak in both SDFs, symmetrically located above and under the h plane, indicative of the presence of a π -hydrogen bond. Visual inspection of the solvent molecules bound at the site reveals a solvent orientation consistent with π -hydrogen bonding. These occurrences are found to have a temporary character: molecules do bind and unbind on the simulated time scales. Once again, we use the procedure of RDF restriction to quantify the π -hydrogen cloud of the SDFs: this time we restrict the RDF to an axis in a manner complementary to the planar restrictions above (again, see the Supporting Information for additional details). A natural choice of this axis is the normal of the benzene horizontal plane, i.e., its C_6 axis. In this case, the restriction angle was set to $\omega = 21^\circ$ purely on the basis of the SDF geometry to completely isolate the whole π -hydrogen cloud from the rest of the SDF. The obtained RRDFs can be found in the bottom panel of Figure 6 for both the anionic solute and the neutral solute. The π -hydrogen peak here is massive, reaching a magnitude of 9–10 (cf. the gray shaded unrestricted RDF reference). In the same panel, we also present the corresponding coordination numbers to quantify the solvent population of one of the two π -hydrogen binding sites located above and under the planar solute. Even here, it appears that the two studied solutes behave rather similarly and the excess charge on the anion only slightly enhances the interaction: the specific populations are ~ 0.65 and ~ 0.70 molecule per binding site for the neutral and anionic solutes, respectively.

To be able to relate to spectroscopic measurements, we predict the vibrational spectra of both species in liquid ammonia in the form of the vibrational density of states (VDOS) and infrared intensity (IR). Theoretical calculations of vibrational spectra allow, unlike their experimental measurement, for a division of the overall spectrum into atomic (or molecular) contributions. Figure 7 shows, for both systems, the




Figure 7. Solute-associated vibrational spectra of neutral benzene (blue) and benzene radical anion (green) in liquid ammonia. The top panel shows the VDOS, while the bottom panel shows the IR intensity. Gas-phase calculated harmonic vibrational frequencies of neutral benzene are shown in the form of black sticks (doubled in length for degenerate modes). The red triangle points to the group of peaks discussed in the text.

atomic VDOS of the solute and the solute-associated IR spectrum, i.e., only terms including a contribution by the solute dipole. Methodological details and a further decomposition of the IR spectrum into the solute-only and solute–solvent terms can be found in Figure S6. In comparison to its harmonic gas-phase frequencies (stick spectrum in Figure 7), neutral benzene in solution manifests vibrational patterns that can be clearly related to this reference by peak positions. For the anion, this remains obvious only in the well-separated C–H stretch region ($\sim 3200\text{ cm}^{-1}$). In the aromatic fingerprint region, the JT distortions nontrivially alter the spectrum; substantial peak broadening and shifts mean that peak assignment is no longer straightforward. A possible way to identify the anionic species in IR could be based on comparing the regions around 700 and 1500 cm^{-1} , where the pronounced peaks of the neutral species are no longer present in the spectrum of the anion. Consistent with experimental findings in metal–benzene films,¹⁸ we observe a slight red shifting in a majority of the anionic peaks that can be related to their neutral counterparts. Interestingly, we observe additional peaks in the IR that are not present in the VDOS of the solute. Namely, this is the group of peaks around 3500 cm^{-1} whose presence suggests that a dipole response exists in the solute without a corresponding change in its atomic positions. These peaks likely arise as a consequence of the perturbation of the electronic structure of the solute caused by the presence of π -hydrogen bonds. This is supported by several facts such as the high frequency of the vibration comparable to the N–H stretch in ammonia,³⁸ the fact that isolated benzene contains no normal modes at the corresponding frequencies, the fact that the solute–solvent cross-term spectrum (see the Supporting Information) shows a contribution to these peaks, and the fact that benzene solvated in nonpolar

environments (such as liquid benzene itself) does not show these additional peaks in the IR spectra.³⁹

The presented condensed-phase AIMD simulations bridge the gap between experimental observations and gas-phase calculations of the benzene radical anion, while providing crucial insights into the structure and dynamics of the species at the molecular level. Using a combination of advanced computational methodology and an explicit representation of the solvent environment, we were able to simulate a radical anion that is electronically stable in solution over extended periods of time, in agreement with its real-world behavior. Our observation of the localization of the spin density over the aromatic ring points to the existence of a bound state. Explicit confirmation that the excess electron on the radical anion is bound with respect to the vacuum level is a possible subject of a follow-up study that would build on the simulations presented here. Although our results do not directly address the thermodynamic stability of the benzene radical anion relative to solvated electrons, the species persists in our simulations for tens of picoseconds, enough to be characterized in a meaningful way. Furthermore, its thermodynamic stability is implied, at least in the sense of equilibrium with solvated electrons, by the fact that the radical anion is experimentally measurable. The most prominent structural feature of our condensed-phase benzene radical anion is the presence of JT distortions that are structurally consistent with previous experimental studies and theoretical studies of the gas-phase metastable species in finite basis sets. We also identified clear differences between the vibrational spectra of the radical anion and its neutral counterpart. Because of the explicit representation of the solvent environment, our results also offer information about the structure and dynamics of the solvation shell. A particularly intriguing aspect of the behavior of the solvent arises in the form of π -hydrogen bonding. This interaction appears clearly in our solvent distributions and leaves a specific imprint in the IR spectrum associated with the solute, enabled by the dipole response of the solute to the vibrational motion of the solvent. Our simulations do not contain the proton source, such as an aliphatic alcohol, required for the Birch reduction to proceed. We can, however, hypothesize that should it be introduced, its participation in this π -hydrogen bond might constitute the initial protonation step of the Birch reduction mechanism sketched schematically in eq 1. Thus, the demonstrated ability of AIMD to treat this solute has the potential to bring us closer to explicit modeling of the chemical reactivity in the Birch reduction as a whole. Such a study would provide a complement to the treatment of reaction mechanisms in organic chemistry that is often empirical and simplified, as illustrated, for instance, by the “organic” representation of the radical anion with a localized dot in eq 1. This is in contrast with the highly complex delocalized distribution of spin density over the anion observed in our simulations that follows the dynamic JT distortions of the ring. While experimental data are currently not available for some of the properties calculated in this work, we anticipate that this will change in the future so that a more detailed understanding of the system can arise from a combination of theoretical and experimental work.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jpclett.0c01505>.

6036

<https://dx.doi.org/10.1021/acs.jpclett.0c01505>
J. Phys. Chem. Lett. 2020, 11, 6032–6038

Description of the computational methodology, details of data analysis, additional information about the energetics of the JT distortion of the benzene radical anion, and additional computational treatment of the calculated solute-associated vibrational spectra in Figure 7 (PDF)

Video file visualizing the pseudorotational motion of the benzene ring (MP4)

Video file visualizing the simulated trajectory of the benzene radical anion, including the evolution of the spin density (MP4)

■ AUTHOR INFORMATION

Corresponding Author

Ondrej Marsalek – Charles University, Faculty of Mathematics and Physics, 121 16 Prague, Czech Republic; orcid.org/0000-0002-8624-8837; Email: ondrej.marsalek@mff.cuni.cz

Authors

Krystof Brezina – Charles University, Faculty of Mathematics and Physics, 121 16 Prague, Czech Republic; Institute of Organic Chemistry and Biochemistry of the Czech Academy of Sciences, 166 10 Prague, Czech Republic; orcid.org/0000-0003-0285-1282

Pavel Jungwirth – Institute of Organic Chemistry and Biochemistry of the Czech Academy of Sciences, 166 10 Prague, Czech Republic; orcid.org/0000-0002-6892-3288

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.jpclett.0c01505>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

This work was supported by the Primus16/SCI/27/247019 grant from Charles University. K.B. acknowledges funding from the IMPRS for Many Particle Systems in Structured Environments. This work was supported by Project SVV 260586 of Charles University. P.J. is thankful for support from the European Regional Development Fund (Project ChemBioDrug no. CZ.02.1.01/0.0/0.0/16_019/0000729). This work was supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations Project "IT4Innovations National Supercomputing Center LM2015070". The authors thank Christoph Schran, Tomáš Martinek, and Vojtěch Košťál for helpful comments on the manuscript.

■ REFERENCES

- (1) Birch, A. J. Reduction by dissolving metals. Part I. *J. Chem. Soc.* **1944**, 430–436.
- (2) Birch, A. J. Reduction by dissolving metals. Part II. *J. Chem. Soc.* **1945**, 809–813.
- (3) Birch, A. J. Reduction by dissolving metals. Part III. *J. Chem. Soc.* **1946**, 593–597.
- (4) Birch, A. J. Reduction by dissolving metals. Part IV. *J. Chem. Soc.* **1947**, 102–105.
- (5) Birch, A. J. Reduction by dissolving metals. Part V. *J. Chem. Soc.* **1947**, 1642–1648.
- (6) Birch, A. J.; Mukherji, S. M. Reduction by dissolving metals. Part VI. Some applications in synthesis. *J. Chem. Soc.* **1949**, 0, 2531–2563.
- (7) Clayden, J.; Greeves, N.; Wothers, P. *Organic Chemistry*, 2nd ed.; Oxford University Press, 2012.
- (8) Buttersack, T.; et al. Photoelectron spectra of alkali metal-ammonia microjets: From blue electrolyte to bronze metal. *Science* **2020**, *368*, 1086–1091.
- (9) Boese, A. D.; Chandra, A.; Martin, J. M.; Marx, D. From ab initio quantum chemistry to molecular dynamics: The delicate case of hydrogen bonding in ammonia. *J. Chem. Phys.* **2003**, *119*, 5965–5980.
- (10) Buttersack, T.; Mason, P. E.; McMullen, R. S.; Martinek, T.; Brezina, K.; Hein, D.; Ali, H.; Kolbeck, C.; Schewe, C.; Malerz, S.; Winter, B.; Seidel, R.; Marsalek, O.; Jungwirth, P.; Bradforth, S. E. Valence and Core-Level X-ray Photoelectron Spectroscopy of a Liquid Ammonia Microjet. *J. Am. Chem. Soc.* **2019**, *141*, 1838–1841.
- (11) Zurek, E.; Edwards, P. P.; Hoffmann, R. A molecular perspective on lithium-ammonia solutions. *Angew. Chem., Int. Ed.* **2009**, *48*, 8198–8232.
- (12) Birch, A. J. Steroid hormones and the Luftwaffe. A venture into fundamental strategic research and some of its consequences: The Birch reduction becomes a birth reduction. *Steroids* **1992**, *57*, 363–377.
- (13) Birch, A. J. The Birch reduction in organic synthesis. *Pure Appl. Chem.* **1996**, *68*, 553–556.
- (14) Joshi, D. K.; Sutton, J. W.; Carver, S.; Blanchard, J. P. Experiences with commercial production scale operation of dissolving metal reduction using lithium metal and liquid ammonia. *Org. Process Res. Dev.* **2005**, *9*, 997–1002.
- (15) Hinde, A. L.; Poppinger, D.; Radom, L. Ab Initio Study of the Benzene Radical Anion. *J. Am. Chem. Soc.* **1978**, *100*, 4681–4685.
- (16) Tuttle, T. R.; Weissman, S. I. Electron Spin Resonance Spectra of the Anions of Benzene, Toluene and the Xylenes. *J. Am. Chem. Soc.* **1958**, *80*, 5342–5344.
- (17) Shida, T.; Iwata, S. Electronic Spectra of Ion Radicals and Their Molecular Orbital Interpretation. III. Aromatic Hydrocarbons. *J. Am. Chem. Soc.* **1973**, *95*, 3473–3483.
- (18) Moore, J. C.; Thornton, C.; Collier, W. B.; Devlin, J. P. Vibrational spectra, Jahn-Teller distortion, and the structure of the benzene radical anion. *J. Phys. Chem.* **1981**, *85*, 350–354.
- (19) Birch, A. J.; Hinde, A. L.; Radom, L. A Theoretical Approach to the Birch Reduction. Structures and Stabilities of the Radical Anions of Substituted Benzenes. *J. Am. Chem. Soc.* **1980**, *102*, 3370–3376.
- (20) Bazante, A. P.; Davidson, E. R.; Bartlett, R. J. The benzene radical anion: A computationally demanding prototype for aromatic anions. *J. Chem. Phys.* **2015**, *142*, 204304.
- (21) Sanche, L.; Schulz, G. J. Electron transmission spectroscopy: Resonances in triatomic molecules and hydrocarbons. *J. Chem. Phys.* **1973**, *58*, 479–493.
- (22) Jordan, K. D.; Michejda, J. A.; Burrow, P. D. The Relative Stability of Alkyl-Substituted Benzene Anions in the Gas Phase. *J. Am. Chem. Soc.* **1976**, *98*, 1295–1296.
- (23) Hasegawa, A.; Shiotani, M.; Hama, Y. ESR studies on Jahn-Teller distortion in the radical anions and cations of hexafluorobenzene. *J. Phys. Chem.* **1994**, *98*, 1834–1839.
- (24) Maeyama, T.; Oikawa, T.; Seguchi, K.; Mikami, N. Considerable stability and visible absorption of the benzene - water hexamer cluster anion. *J. Phys. Chem. A* **1997**, *101*, 8371–8373.
- (25) Marasas, R. A.; Iyoda, T.; Miller, J. R. Benzene radical ion in equilibrium with solvated electrons. *J. Phys. Chem. A* **2003**, *107*, 2033–2038.
- (26) Mortensen, J.; Heinze, J. The Electrochemical Reduction of Benzene—First Direct Determination of the Reduction Potential. *Angew. Chem., Int. Ed. Engl.* **1984**, *23*, 84–85.
- (27) Wilds, A. L.; Nelson, N. A. A Superior Method for Reducing Phenol Ethers to Dihydro Derivatives and Unsaturated Ketones. *J. Am. Chem. Soc.* **1953**, *75*, 5360–5365.
- (28) Hook, J. M.; Mander, L. N. Recent developments in the birch reduction of aromatic compounds: Applications to the synthesis of natural products. *Nat. Prod. Rep.* **1986**, *3*, 35–85.
- (29) Brandsma, L.; van Soelingen, J.; Andringa, H. A procedure for the preparation of pure 1, 4-cyclohexadiene from benzene. *Synth. Commun.* **1990**, *20*, 2165–2168.

- (30) Hirshfeld, F. L. Bonded-atom fragments for describing molecular charge densities. *Theor. Chim. Acta* **1977**, *44*, 129–138.
- (31) Diraison, M.; Martyna, G. J.; Tuckerman, M. E. Simulation studies of liquid ammonia by classical ab initio, classical, and path-integral molecular dynamics. *J. Chem. Phys.* **1999**, *111*, 1096–1103.
- (32) Vandevondele, J.; Krack, M.; Mohamed, F.; Parrinello, M.; Chassaing, T.; Hutter, J. Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach. *Comput. Phys. Commun.* **2005**, *167*, 103–128.
- (33) Hutter, J.; Iannuzzi, M.; Schiffrin, F.; Vandevondele, J. Cp2k: Atomistic simulations of condensed matter systems. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2014**, *4*, 15–25.
- (34) Tachikawa, H. Jahn-Teller Effect of the Benzene Radical Cation: A Direct ab Initio Molecular Dynamics Study. *J. Phys. Chem. A* **2018**, *122*, 4121–4129.
- (35) Preuss, M.; Bechstedt, F. Vibrational spectra of ammonia, benzene, and benzene adsorbed on Si (001) by first principles calculations with periodic boundary conditions. *Phys. Rev. B: Condens. Matter Mater. Phys.* **2006**, *73*, 155413.
- (36) Gierszal, K. P.; Davis, J. G.; Hands, M. D.; Wilcox, D. S.; Slipchenko, L. V.; Ben-Amotz, D. π -hydrogen bonding in liquid water. *J. Phys. Chem. Lett.* **2011**, *2*, 2930–2933.
- (37) Allesch, M.; Schwegler, E.; Galli, G. Structure of hydrophobic hydration of benzene and hexafluorobenzene from first principles. *J. Phys. Chem. B* **2007**, *111*, 1081–1089.
- (38) Bromberg, A.; Kimel, S.; Ron, A. Infrared spectrum of liquid and crystalline ammonia. *Chem. Phys. Lett.* **1977**, *46*, 262–266.
- (39) Thomas, M.; Brehm, M.; Kirchner, B. Voronoi dipole moments for the simulation of bulk phase vibrational spectra. *Phys. Chem. Chem. Phys.* **2015**, *17*, 3207–3213.

C.4 | Paper IV

Reprinted from Brezina, K., Kostal, V., Jungwirth, P. and Marsalek, O.: Electronic structure of the solvated benzene radical anion, *Journal of Chemical Physics*, 156, 1, 14501, 2022, with the permission of AIP Publishing.

Electronic structure of the solvated benzene radical anion

Cite as: J. Chem. Phys. 156, 014501 (2022); doi: 10.1063/5.0076115

Submitted: 21 October 2021 • Accepted: 29 November 2021 •

Published Online: 3 January 2022

Krystof Brezina,^{1,2}  Vojtech Kostal,²  Pavel Jungwirth,²  and Ondrej Marsalek^{1,a)} 

AFFILIATIONS

¹Charles University, Faculty of Mathematics and Physics, Ke Karlovu 3, 121 16 Prague 2, Czech Republic²Institute of Organic Chemistry and Biochemistry of the Czech Academy of Sciences, Flemingovo nám. 2, 16610 Prague 6, Czech Republic

Note: This paper is part of the JCP Special Topic on Beyond GGA Total Energies for Solids and Surfaces.

^{a)}Author to whom correspondence should be addressed: ondrej.marsalek@mff.cuni.cz

ABSTRACT

The benzene radical anion is a molecular ion pertinent to several organic reactions, including the Birch reduction of benzene in liquid ammonia. The species exhibits a dynamic Jahn–Teller effect due to its open-shell nature and undergoes pseudorotation of its geometry. Here, we characterize the complex electronic structure of this condensed-phase system based on *ab initio* molecular dynamics simulations and GW calculations of the benzene radical anion solvated in liquid ammonia. Using detailed analysis of the molecular and electronic structure, we find that the spatial character of the excess electron of the solvated radical anion follows the underlying Jahn–Teller distortions of the molecular geometry. We decompose the electronic density of states to isolate the contribution of the solute and to examine the response of the solvent to its presence. Our findings show the correspondence between instantaneous molecular structure and spin density; provide important insights into the electronic stability of the species, revealing that it is, indeed, a bound state in the condensed phase; and offer electronic densities of states that aid in the interpretation of experimental photoelectron spectra.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0076115>

I. INTRODUCTION

Liquid ammonia is particularly well-known as a solvent that sustains long-lived solvated electrons formed by the dissolution of alkali metals.¹ Recently, we used the flexible combination of refrigerated liquid microjet x-ray photoelectron spectroscopy (XPS)^{2–3} and advanced *ab initio* calculations to characterize the electronic structure of neat liquid ammonia⁴ as well as the alkali metal solutions.⁵ In the latter, the hallmark XPS feature of the solvated electron is located at the electron binding energy of -2.0 eV relative to the vacuum level and its concentration dependence was used to experimentally map the electrolyte-to-metal transition.⁵ Solvated electrons, essentially localized electrons bound in cavities formed within the solvent structure, act as powerful chemical reducing agents and, as such, find applications in numerous organic reductions. Arguably, the best known example is the Birch reduction of benzene in the environment of solvated electrons with the addition of an aliphatic alcohol.⁶ During the course of the reaction, the solvated electron binds to the benzene molecule, forming the benzene radical anion as the first reactive intermediate. This chemical

role of the benzene radical anion and its prominent position as the simplest example of an aromatic anion has prompted several experimental^{7–10} and theoretical^{11,12} studies of the species in the past. A particularly intriguing conclusion that arises from these studies is that the stability of the species is environment-dependent. In particular, the isolated benzene radical anion represents an unbound metastable shape resonance with a lifetime on the femtosecond time scale, which was consistently demonstrated both by *ab initio* calculations^{11,12} and by electron scattering experiments¹⁰ in the gas phase. In contrast, the feasibility of the Birch reduction and various spectroscopic experiments performed in different polar solvents,^{7–9} which measure the species over extended time scales, imply the stability of the electronic structure of the benzene radical anion as well as its thermodynamic stability in the context of a chemical equilibrium with solvated electrons.¹³ In addition to the non-trivial behavior of the electronic structure with respect to solvation, the presence of an excess electron in an initially energetically degenerate quantum state gives rise to a dynamic multimode $E \otimes e$ Jahn–Teller (JT) effect,^{14,15} which results in complex behavior of the electronic structure and the molecular geometry. In particular, the optimal

31 August 2023 13:05:29

molecular structure of the benzene radical anion is not a highly symmetric hexagonal one like that of the neutral benzene parent molecule but is rather represented by a continuum of lower-symmetry structures that form the so-called pseudorotation path.¹²

Anticipating a future XPS measurement of the benzene radical anion as a natural continuation of the metal–ammonia solution research, we have previously investigated the benzene radical anion in a liquid ammonia solution using computational methods with the aim to shed light on its structure, dynamics, and spectroscopy and to provide a theoretical basis to aid the interpretation of various experimental data. In our original work,¹⁶ we performed *ab initio* molecular dynamics (AIMD) of the explicitly solvated anion under periodic boundary conditions. These simulations were realized at the hybrid density functional theory (DFT) level of the electronic structure, which we have shown to be, despite its high computational cost, a necessary methodological component to obtain a physically meaningful description of the benzene radical anion. At this level of theory, the excess electron spontaneously localizes on the benzene ring and remains stable for the length of the simulation, indicating the presence of a bound electronic state. Based on these simulations, we then addressed the structure of the solute and tracked the systematic geometry distortions and pseudorotation due to the dynamic JT effect that persists in the thermalized bulk system. More recently, we approached the problem of the solvent-induced stability of the benzene radical anion from the point of view of molecular clusters derived from the original condensed-phase AIMD simulations.¹⁷ In that study, we calculated the excess electron vertical binding energy using explicit ionization in clusters of increasing size and found results ranging from -2.0 to -3.0 eV at the infinite cluster size limit depending on the specific methodology.

The aim of the present work is to shed light on the electronic structure of the benzene radical anion by employing advanced electronic structure calculations and analysis performed on our original AIMD thermal geometries. In the spatial domain, we describe the probability distribution of the excess electron and its correlation with the underlying JT distortions of molecular geometry using unsupervised machine learning methods.¹⁸ These methods have been used to analyze molecular dynamics trajectories and characterize representative molecular configurations of the studied systems in a bias-free way.^{19,20} Here, we employ clustering analysis not only to the distribution of nuclear configurations of the benzene radical anion, but we also use it in conjunction with dimensionality reduction to characterize the electronic structure. Then, in the energy domain, we aim to predict the binding energies of all the valence electrons in the studied system, which can be directly compared to XPS data. To avoid the unphysical orbital energies directly available from the AIMD on-the-fly Kohn–Sham (KS) DFT electronic structure, we perform computationally demanding condensed-phase G_0W_0 calculations^{21,22} on the AIMD geometries to predict the electronic densities of states (EDOS). To better understand the contributions of the individual species in the system in question, we additionally employ an approach which projects the EDOS on local atomic orbitals to resolve the calculated data by species and in space.

The rest of this paper is organized as follows. In Sec. II, we discuss the details of the performed simulations and calculations and describe the technical foundations of the employed analysis.

The main findings are then presented and discussed in Sec. III. There, we first focus on the results pertaining to the JT effect on the electronic structure and its correlation with the underlying molecular geometry. Then, we move on to the energetics of the electronic structure and the question of the stability and binding energy of the solvated benzene radical anion. These results are referenced against those for neutral benzene solvated in liquid ammonia and neat liquid ammonia itself. Finally, we summarize our results and draw conclusions in Sec. IV.

II. METHODOLOGY

A. AIMD simulations

The original AIMD simulations of the benzene radical anion and neutral benzene in liquid ammonia under periodic boundary conditions were realized using the CP2K 5.1 package^{23–25} and its Gaussian and plane wave²⁶ electronic structure module Quickstep.²⁷ Both simulated systems consisted of one solute molecule and 64 solvent molecules in a cubic box of a fixed side length of 13.745 and 13.855 Å for the benzene radical anion and neutral benzene, respectively. The nuclei were propagated with a 0.5 fs time step in the canonical ensemble at 223 K using the stochastic velocity-rescaling thermostat.²⁸ The electronic structure was calculated using the revPBE0-D3 hybrid density functional^{29–32} to limit the self-interaction error, as required for the localization of the excess electron and the stability of the benzene radical anion.¹⁶ The KS wavefunctions were expanded into the TZV2P primary basis set,³³ while the density was expanded in an auxiliary plane wave basis with a 400 Ry cutoff. Goedecker–Teter–Hutter pseudopotentials³⁴ were used to represent the core 1s electrons of the heavy atoms. Additionally, the auxiliary density matrix method³⁵ with the cpFIT3 auxiliary basis set³⁵ was used to accelerate the computationally demanding hybrid DFT electronic structure calculations. The total simulated time was 100 ps for both systems, each collected from five 20 ps trajectories initialized from decorrelated and equilibrated initial conditions.

B. G_0W_0 calculations

In this work, we use the G_0W_0 method,^{21,22} which gives access to physically meaningful one-electron energy levels of the investigated condensed-phase periodic systems. This is in contrast with orbital energies of the underlying KS DFT, which should not formally be considered as one-electron energies. For each solute, these calculations are performed on top of 205 DFT-AIMD thermal structures extracted from the AIMD trajectories with a 0.5 ps stride with the revPBE0-D3/TZV2P KS wavefunctions being used as a starting point to obtain the corrected G_0W_0 energies. These calculations are realized using the CP2K package, version 7.1. The self-energy is described analytically over the real frequency axis using the Padé approximation, and the Newton–Raphson fixed point iteration is employed for numerical solution of the corresponding algebraic equations. The influence of periodic boundary conditions on the G_0W_0 energies is minimized by employing a periodicity correction scheme.³⁶ The resulting EDOS, obtained as the distribution of the G_0W_0 energies, is described as a continuous probability density function through the kernel density estimation method

using a Gaussian kernel with a 0.02 eV bandwidth. The G_0W_0 calculations performed in periodic boundary conditions do not directly provide the absolute values of electron binding energies due to the absence of the explicit liquid-vacuum boundary. Thus, to access the absolutely positioned EDOS, the whole spectrum must be shifted on the energy axis by a suitable constant. In other works, this was achieved by auxiliary slab calculations that provide an estimate of the shift.^{37,38} In our previous work on neat liquid ammonia combining G_0W_0 calculations with liquid XPS,^{2,4} we aligned the average energy of the calculated liquid $3a_1$ peak to -9.09 eV, the average of the same peak obtained experimentally. This bypassed the need for additional *ab initio* calculations and facilitated the comparison of the whole spectrum between theory and experiment. Here, we exploit the fact that, as detailed in Sec. III, the electronic perturbation of the liquid ammonia solvent by the presence of the benzene radical anion is minor. As such, the total EDOS was shifted to match the same experimental valence liquid ammonia peak as in our previous work. The value of the shift was determined from the mean energy of the $3a_1$ ammonia peak of the total EDOS with the solute included (other options are discussed in Sec. S2 of the supplementary material).

To gain insight into the contributions of individual chemical species to the total G_0W_0 EDOS, we decompose this quantity into separate densities for each species and address the differences between the neat ammonia data and the data from systems with solutes. Specifically, we rely on the original formulation of the projected density of states (PDOS) for KS orbitals,³⁹ which projects the total EDOS on the respective part of the atomic orbital basis set of every atom in the system individually. Extending the original approach, we use these projections for the G_0W_0 -corrected binding energies since the spatial orbitals are identical between KS DFT and G_0W_0 . For each atom and each configuration, each G_0W_0 energy is assigned a weight based on the magnitude of the projection of the corresponding orbital on that atom. Naturally, these atomic contributions can be collected into molecular contributions as needed for each particular system. The total EDOS can be expressed as the following ensemble average over the contributing structures:

$$\rho(E) = \left\langle \sum_n \delta(E - E_n) \right\rangle, \quad (1)$$

where E_n are the G_0W_0 one-electron energy eigenvalues and angle brackets denote an average over the ensemble of thermal structures. To decompose it, we use a projection on an atom-centered linear combination of atomic orbitals (LCAO) basis set $\{|I\gamma\rangle\}$. This basis satisfies the completeness relation over the spanned space

$$\sum_I \sum_\gamma |I\gamma\rangle \langle I\gamma| = \hat{1}, \quad (2)$$

where the summation runs over all atoms I and all additional quantum numbers γ and $\hat{1}$ denotes the identity operator. Using the orthonormality of the original KS orbitals $|\psi_n\rangle$ that remain unchanged during the G_0W_0 calculation, we can expand the total EDOS definition as a sum over atomic projections as

$$\begin{aligned} \rho(E) &= \left\langle \sum_n \langle \psi_n | \psi_n \rangle \delta(E - E_n) \right\rangle \\ &= \left\langle \sum_n \sum_I \sum_\gamma \langle \psi_n | I\gamma \rangle \langle I\gamma | \psi_n \rangle \delta(E - E_n) \right\rangle \\ &= \sum_I \left\langle \sum_n \sum_\gamma |\langle \psi_n | I\gamma \rangle|^2 \delta(E - E_n) \right\rangle \\ &\equiv \sum_I \langle S_I(E) \rangle \equiv \sum_I \rho_I(E), \end{aligned} \quad (3)$$

where we have labeled the overlap-weighted kernel of the thermal average $S_I(E)$ and the whole thermally averaged atomic projection $\rho_I(E)$. These atomic projections can then be summed over arbitrary subsets of atoms to obtain a PDOS on any species in question. Moreover, we can further resolve the atomic contributions as a function of distance r from a chosen point of reference as the following two-dimensional distribution:

$$\rho(E, r) = \frac{1}{4\pi r^2 g(r)} \sum_I \langle S_I(E) \delta(r - r_I) \rangle, \quad (4)$$

where r_I is the distance of the I th atom from the point of reference and the normalization factor in the denominator based on the radial distribution function $g(r)$ of the chosen species around the same point of reference ensures a uniform marginal distribution in r .

C. Clustering analysis

The clustering of the relevant feature space vectors is based on the Gaussian Mixture Model (GMM) as implemented in the scikit-learn Python library.⁴⁰ For the molecular geometries, we assign features as all vibrational normal modes with JT-active symmetry to naturally describe the distortions in an 8D configuration space. For the electronic structure, we use a high-dimensional abstract feature space that relies on a Fourier decomposition of the respective spin densities. Both feature spaces are described in detail in the following paragraphs. The GMM algorithm was chosen over the commonly used k -means clustering since it allows us to reach a similar goal in a more flexible and general way and, moreover, yields a continuous parametrization of the obtained clusters in terms of high-dimensional Gaussian functions that can be used to evaluate the cluster membership probability. The full covariance in all dimensions was employed to account for possible spatial anisotropy of the clusters, and a tight convergence limit of 10^{-5} was used.

III. RESULTS AND DISCUSSION

In the following paragraphs, we focus on the spatial character of the excess electron of the benzene radical anion using the spin density, an observable quantity obtained directly from an unrestricted Kohn-Sham (KS) DFT calculation. We aim at a description of the evolution of the spin density in the context of the condensed-phase JT effect, which governs the distortions of the underlying molecular geometry of the benzene radical anion solvated in liquid ammonia.¹⁶ Specifically, we ask if the molecular distortions correlate with the immediate shape of the spin density and, therefore, if information about the JT state of the solute can be extracted directly from the electronic structure of the solvated species, similarly to how it can be extracted from its molecular geometry. Later on, we turn our attention to the energetics of the electronic structure, predict electronic

densities of states for the studied system, and then discuss in detail in the context of the question of the stability of the solvated benzene radical anion and from the perspective of interpretation of XPS data.

A. The Jahn–Teller effect on the molecular and electronic structure

The essence of the JT effect in the benzene radical anion is as follows. As the D_{6h} -symmetric benzene molecule accepts an excess electron, the formed degenerate E_{2u} electronic state of the radical anion becomes unstable since it corresponds to a conical intersection between two adiabatic potential energy hypersurfaces (APESs). This instability is resolved by a symmetry-lowering distortion along the JT-active normal modes of e_{2g} symmetry, which brings the system into a minimum on the pseudorotational path on the lower branch of the JT-split APES. At the same time, the symmetry of the initial electronic state is reduced as well, with two new possible lower symmetries, A_u and B_{1u} , corresponding to the ground state of the benzene radical anion in the two opposite distortions of the molecular geometry.

1. Clustering of molecular geometries

The natural coordinates to describe the molecular distortions are the four degenerate pairs of JT-active normal modes. These are adopted here consistently with our previous work from the vibrational normal modes of an optimized neutral benzene molecule since it shares the same molecular structure and the point group with the radical anion in its reference undistorted geometry. A physically meaningful observation of the JT pseudorotation can be made by averaging the full 8D data over all modes that do not exhibit a strong enough JT split to be observable in the thermal system. Thus, the pseudorotation can be represented as a 2D distribution in the pair of remaining e_{2g} modes at 1654 cm^{-1} , which show an appreciably strong JT effect. In this case, the free energy landscape of the pseudorotation valley is essentially flat and the path around it is described by the pseudorotation angle $\theta = \arctan2(Q_y/Q_x)$, a scalar parameter which represents the polar angle in the 2D subspace of the relevant normal mode coordinates labeled Q_x and Q_y .¹⁶

In order to analyze the full 8D distribution, we applied the GMM clustering algorithm to the normal modes dataset with the aim to find representative distortions. However, unlike in the case of the electronic structure discussed in the following paragraphs, the resulting clustering of the data is not satisfactory for several reasons. Motivated by the threefold symmetry of the reference gas-phase APES, we attempted to separate the data into both three and six clusters. In both cases, clustering of comparable quality was obtained, which implies that there is no clear number of natural clusters in the dataset. This is further supported by additional attempts to cluster the data into a number of clusters that does not respect the inherent symmetry of the problem: again, similar outputs were produced. Moreover, the clustering is generally not reproducible and inconsistent positions of clusters are obtained each time. As a measure of clustering performance, we use silhouette coefficients, which range from -1 (wrong clustering) through 0 (poor clustering) to $+1$ (excellent clustering).⁴¹ If we cluster our data into three groups, the average silhouette coefficient does not exceed the value of ~ 0.08 , which quantifies the insufficient separation of the data (a silhouette plot is presented in Sec. S2 of the [supplementary material](#)). This

demonstrated lack of clear separation in the molecular geometries suggests that the remaining modes do not bring much additional structure to the dataset in comparison to the reduced 2D distribution in Q_x and Q_y , and the essentially flat character of the probability distribution around the pseudorotation valley generalizes to the full dimensionality. Therefore, we adhere to the simpler continuous parametrization by the pseudorotation angle θ to describe molecular distortions in the following analysis.

2. Spin density dimensionality reduction

To motivate the analysis of the electronic structure of the solvated species, we consider optimized gas-phase benzene radical anion structures where the excess electron is artificially localized due to a finite orbital basis set. The spin density distributions for the two distinct JT distortions are shown in [Fig. 1](#). The spin density of the A_u state (left) is characterized by four atom-centered maxima and two less pronounced minima localized along one of the C_2 symmetry axes; the B_{1u} spin density (right) exhibits two maxima localized on distal carbon atoms along the corresponding C_2 axis and two elongated bridge-like positive deviations over a pair of carbon–carbon bonds parallel with this C_2 axis. We also have to take into account that the high symmetry of the benzene molecular geometry allows for distortion in three equivalent directions corresponding to the three horizontal apex-to-apex C_2 axes in the D_{6h} point group. These distortions are represented by two sets of three equivalent stationary points around the point of high symmetry on the pseudorotation APES, each separated by a pseudorotation angle of 60° from its opposite-kind neighbors and by 120° from its pseudorotated images. As a consequence, three equivalent A_u -type and three equivalent B_{1u} -type spin densities exist, which correspond to the six APES stationary points. The pseudorotation of the nuclear geometry between these minima is discussed in detail in [Ref. 16](#); a video file illustrating the evolution of the spin density on top of the pseudorotating geometry in the idealized gas-phase case is included in [Sec. S3 of the supplementary material](#).




FIG. 1. The spin densities of the A_u and B_{1u} electronic states of the benzene radical anion. The presented idealized geometries and spin densities were obtained from a finite basis set gas-phase calculation at the hybrid DFT level, as used for the AIMD simulations; similar spin densities are, however, observed in the condensed-phase simulations. The positive deviations of the spin density are shown in green at two contours, 0.025 a_0^{-3} (opaque) and 0.006 a_0^{-3} (transparent), while the negative deviations are shown in purple at the same isovalues with a negative sign. The molecular structure of the benzene radical anion is shown in gray as a whole.

Eventually, we want to analyze the thermal data in the condensed phase—the natural environment where the solvated benzene radical anion is electronically stable and physically relevant observation of the JT effect, and the associated spin density can be made. For this purpose, we design a two-step dimensionality reduction procedure that represents the spin densities in a feature space of reasonable dimension in such a way that the two idealized JT-distorted cases can be distinguished. The periodicity of the spin density along the aromatic ring makes it advantageous to express its spatial dependence in terms of a local spherical coordinate system r, ϑ , and φ (where φ is the polar angle ranging from 0 to 2π). These coordinates are obtained by the usual transformation from a local Cartesian system in which the x, y plane is represented by the molecular plane of the benzene radical anion and the z axis is represented by its normal with its origin at the solute center of mass (see Sec. S1 of the supplementary material for details). The spherical coordinates represent a natural description for the systems in question and allow us to reduce the dimensionality of the full spin density into a one-dimensional (1D) function by partial integration. As documented in Sec. S2 of the supplementary material, the 1D spin densities in r and ϑ show practically perfect overlap for the two spin density types and thus bring no distinction between them. The information that distinguishes the two types is contained in the remaining possible spin density in φ ,

$$\rho_s(\varphi) = \int_0^\pi \int_0^{r_{\max}} d\vartheta dr r^2 \sin \vartheta \rho_s(r, \vartheta, \varphi), \quad (5)$$

which describes the character of the spin density around the benzene ring. Its shape can be traced back to the spatial characteristics of the full spin densities through the respective sequence of the 1D maxima and minima along the aromatic ring, as shown for the idealized spin densities in Fig. 2, top panel, full lines. In terms of $\rho_s(\varphi)$, the pseudorotation of each type of the full spin density by 120° simply translates into a 120° shift on the φ axis.




FIG. 2. The relevant 1D spin densities $\rho_s(\varphi)$ for the A_u and B_{1u} states (Fig. 1). The original functions are shown as the black solid lines. The $N = 20$ Fourier-reconstructed curves are shown as the blue and orange dashed lines, respectively. The bottom panel illustrates five samples equidistant in their degree n from the employed $N = 20$ Fourier basis with the sine components shown in dark gray and cosine components in light gray.

At this point, the 3D spin density is reduced to a 1D function that is still fully capable of distinguishing between the two spin density types. An additional level of simplification that opens the door to numerical analysis is achieved by mapping the continuous 2π -periodic 1D spin densities onto discrete vectors by means of a Fourier series and noting that only the first few harmonics are necessary to achieve a highly accurate decomposition as demonstrated by the dashed curves in the top panel of Fig. 2. This set of Fourier coefficients clearly distinguishes the two idealized spin densities in relatively few dimensions. While the technical aspects of this step are discussed in detail in Sec. S1 of the supplementary material, we note here that the Fourier decomposition was performed using the first 20 harmonics, yielding an 82-dimensional Euclidean feature vector for each spin density sample [a total of $2(2N + 1)$ real coefficients are needed for a Fourier series counting N harmonic functions].

3. Clustering of the electronic structure

We are now able to represent each spin density distribution in a compact way and can move to the analysis of the electronic structure of the condensed-phase system. A visual inspection of the trajectory¹⁶ of the solvated benzene radical anion clearly reveals the presence of two limiting spin density structures similar to the optimized ones. Therefore, we aim to perform an analysis that will allow us to divide the observed ensemble of condensed-phase spin densities into six categories centered around each of the limiting spin density structures and including the surrounding thermal population. Once this is established, one can examine the correlation between the immediate electronic structure and the underlying molecular geometry of the solute.

To categorize the spin densities of the thermal solvated system, we turn again to GMM clustering to separate the data now concisely represented as feature vectors constructed out of Fourier coefficients. GMM is able not only to split the data into natural clusters but also to provide a continuous parametrization of each cluster through evaluation of posterior probabilities of cluster membership. Indeed, in this case, the dataset splits cleanly into six clusters, as shown by the cluster silhouettes presented in Fig. 3, which average to the mean silhouette coefficient of ~ 0.4 and contain no outliers for




FIG. 3. Characterization of the clustering of the electronic structure by the means of silhouette plots. The top three clusters represent the A_u clusters, and the bottom three clusters represent the B_{1u} clusters.

31 August 2023 13:05:29

the A_u state and only a small number of outliers (negative silhouette coefficients) for the B_{1u} state. Additional clustering validation is documented in Sec. S2 of the [supplementary material](#). The centers of the six clusters then correspond to the electronic structure at the six possible D_{2h} distortions, and the population of each cluster corresponds to the thermal fluctuations around these minima. This is directly shown by summing up the Fourier series defined by the coordinates of the cluster centers to obtain new 1D spin densities. These exhibit physically meaningful properties such as close-to-reference shapes (such as those shown in [Fig. 2](#)) and the expected 120° shifts within each type group (see the [supplementary material](#), Sec. S2). While these findings show that the excess electron structure is analogous to that found for the benzene radical anion in the gas phase using a comparable finite orbital basis set, it is important to keep in mind that such a system converges to an unbound state when the size of the basis set is increased. Only in the condensed phase is the species actually bound and its JT effect observable and are the electronic states physically meaningful and potentially experimentally measurable.

4. Correlation between the electronic structure and molecular geometry

To quantify the correlation between the molecular structure and the spin density, we exploit the features of the trained Gaussian mixture model to assign a posterior probability of belonging to a specific cluster to each spin density data point. Thus, a generalized single-valued parameter $p(A_u)$, which can be defined as a sum over all A_u -type cluster probabilities, gives the overall probability that a data point is of the A_u -type, including all three possible pseudorotations. Clearly, the same can be done for the B_{1u} -type clusters and the identity for complementary probabilities that $p(A_u) + p(B_{1u}) = 1$ has to hold. Now, since each spin density data point has a unique molecular geometry associated with it, the proposed probability parameters can be directly correlated with the underlying molecular distortions characterized by the pseudorotation angle θ as defined above.

We use these electronic probability parameters to weight each point contributing to the probability distribution in θ , which is almost uniform originally. This splits it into two distinct distributions, each with three well-defined peaks separated by a 120° increment. These are shown in [Fig. 4](#), exploiting a representation in polar coordinates with an offset origin. The presented complementary distributions clearly show that the individual symmetries of the molecular distortions are accompanied by spin densities of the same type, as can be deduced from the fact that the distortion at $\theta = 0^\circ$ is uniquely identified with the distortion of molecular geometry corresponding to the A_u electronic state. It thus appears that the electronic character of the JT effect of the benzene radical anion in liquid ammonia closely follows the predicted gas-phase theory, while the solvent acts as a stabilizing, but non-perturbing environment. Due to the correlation shown in [Fig. 4](#), we conclude that similar information about the JT effect can be extracted from the immediate spin density and from the immediate molecular geometry of the solute.

Even though the molecular geometries undergo almost free pseudorotation with effectively no free energy barriers and cannot therefore be clustered into distinct populations of different pseudorotations, the situation is different for the electronic state of

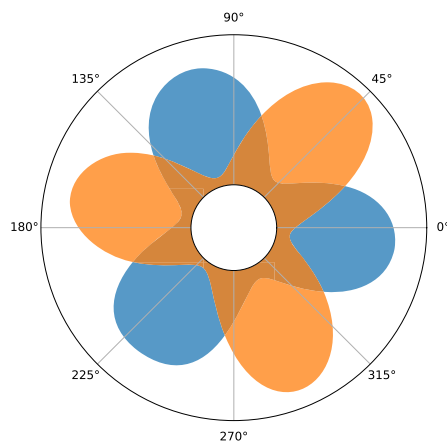


FIG. 4. Correlation of the electronic structure with the distortion of nuclear geometry. Each molecular distortion is characterized here by the value of the pseudorotation angle θ . The distributions of θ weighted by the corresponding electronic parameters $p(A_u)$ (blue) and $p(B_{1u})$ (orange) are shown in polar coordinates with an offset zero-distance.

the system. As it moves along the pseudorotation path, it transitions rather sharply between ground-state spin densities of the two possible symmetries, as revealed by our analysis.

B. Energetics of the electronic structure

At this point, we turn our attention to the energetics of the electronic structure of the whole studied system in terms of one-electron levels. The single-electron energies are calculated using the G_0W_0 method^{21,22} on an ensemble of 205 structures drawn with a 0.5 ps stride from our previously published hybrid DFT trajectories of the benzene radical anion and neutral benzene for comparison. The absolute energies of the whole spectrum were shifted, as detailed in Sec. II. The distribution of the obtained G_0W_0 quasiparticle energies, which accurately approximate electron binding energies, represents the EDOS and is shown in [Fig. 5\(a\)](#). The dominant three-peak pattern in both systems can be readily related to the neat liquid ammonia EDOS,⁴ shown here in gray shading for reference. In our systems with solutes, it is accompanied by a multitude of low-intensity features along the whole range of energies. We can now use the projection approach detailed in Sec. II to isolate these features and examine the solute and solvent spectra separately.

Focusing first on the benzene radical anion, we obtain the solute PDOS shown in [Fig. 5\(b\)](#). Clearly, this component isolates the low-intensity features that do not overlap with the neat ammonia EDOS and, moreover, uncovers additional ones that were previously contained in the high-intensity solvent peaks. Most notably, this solute PDOS suggests that the highest energy state, occupied by the excess

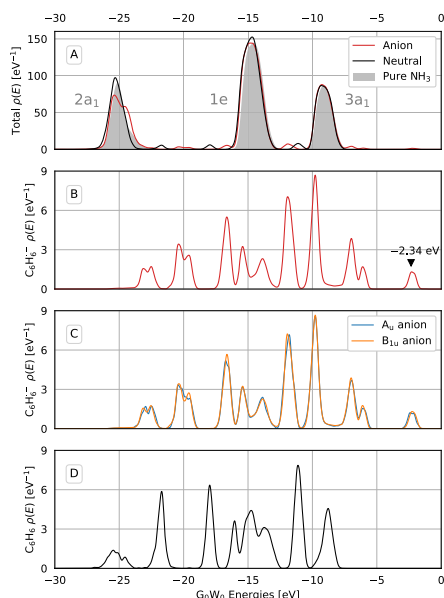


FIG. 5. The solvated benzene radical and neutral benzene total G_0W_0 EDOS and the PDOS projections on the solutes. (a) The total EDOS of the benzene radical anion (red) and neutral benzene (black) in liquid ammonia. The calculated pure liquid ammonia EDOS⁴ is shown in gray. Consistently with the published pure ammonia data, the corresponding peaks are labeled by the symmetry labels of the gas-phase ammonia molecular orbitals. (b) PDOS of the benzene radical anion. The projection shows a detailed account of the electronic structure of the anion, including the highest occupied state, marked by its binding energy and a black triangle. (c) The benzene radical anion PDOS resolved for the two types of JT-relevant electronic structure symmetries. Note that the small differences between the blue (A_u) and orange (B_{1u}) curves, caused by sampling from the corresponding smaller subsets of the calculated G_0W_0 energies, are insignificant within the available statistics. The PDOS of both JT structures is therefore identical. (d) PDOS of neutral benzene in liquid ammonia.

electron, is fully accounted for by the solute, consistent with the previously observed spatial localization of the spin density.¹⁶ Its mean binding energy of -2.34 eV and the absence of tails extending into the positive values prove that the excess electron on benzene is bound relative to the vacuum level, thus conclusively answering the question of stability of the molecular structure of the anion as long as it is solvated in liquid ammonia. This is in excellent agreement with the vertical electron binding energy of -2.30 eV obtained by explicit ionization calculations of benzene radical anion and ammonia gas-phase clusters in the infinite cluster size limit.¹⁷ Compared to neutral benzene [Fig. 5(d)], the whole G_0W_0 anion solute PDOS is systematically shifted toward weaker binding energies by several

electronvolts. Its shape is modified as well, including several peak splittings not observed in the neutral system. These are likely due to the overall lower symmetry of the anion, rather than due to the presence of two distinct JT pseudorotamers, which give rise to identical PDOS within the available statistical sampling, as shown in Fig. 5(c). Since the excess electron binding energy in the benzene radical anion is close to the binding energy of the solvated electron of -2.0 eV,⁵ an overlap might arise in an experimental photoelectron spectrum if the two species coexist in equilibrium, leading to a single broader peak or perhaps a double peak feature. This suggests that the excess electron binding energy itself might not be sufficient to prove the presence of the benzene radical anion. However, a viable workaround exists in the predicted changes of the lower electronic levels of benzene after the addition of the excess electron. These are large enough to be measured, and several bands are localized in the regions where no overlap with the solvent signal is expected, as clearly shown by the projected densities.

Next, we concentrate on the solvent subspace. In Fig. 6, the solvent PDOS shown in the left-hand side panels in gray shading features subtle differences compared to the EDOS of neat ammonia. These appear because of the changes of the electronic structure of the solvent molecules induced by the interaction with the radical anion solute. To better quantify this perturbation, we exploit the molecular resolution of the PDOS projection to resolve the solvent PDOS as a function of distance between the solute center of mass and the ammonia nitrogen atoms (Fig. 6, main panels). The uniformity of the resolved distribution along the distance axis is achieved

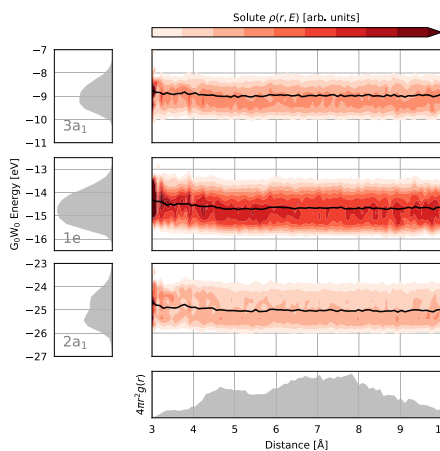


FIG. 6. Electronic density of states projected onto the solute subspace and resolved as a function of distance from the center of mass of the radical anion. The black dashed lines denote the mean of each peak again as a function of distance. The left-hand side panels show the total solute PDOS in gray.

31 August 2023 13:05:29

by factoring out the probability density in this distance. In an infinite system, this is proportional to $4\pi r^2 g(r)$, where $g(r)$ is the radial distribution function. For our finite simulation cell, this quantity is shown in the bottom panel of Fig. 6; note the decay starting after ~ 7 Å that corresponds to half the length of the simulation box. The distance resolution reveals a small systematic shift toward weaker electron binding energies in the proximity of the charged solute, up to 0.4 eV in the case of the 1e peak. The origin of this effect can be attributed to the presence of the excess electron since neutral benzene does not have a similar effect on liquid ammonia; its resolved peaks are essentially flat over the studied distance range (see Sec. S2 of the [supplementary material](#)). The small magnitude of the perturbation of solvent one-electron levels by the solute can be used to justify the alternative method of spectrum resolution by subtraction of the neat solvent that is typically used in an experimental setting where a projection is not an option.⁴² The possible causes of the observed effect include the screening of the electrostatic interaction with the excess charge by the bulk solvent and are discussed in Sec. S2 of the [supplementary material](#) in terms of molecular clusters in open boundary conditions. Additionally, we present a detailed validation of the required PDOS properties in Sec. S2 of the [supplementary material](#).

IV. CONCLUSIONS

The reported analysis of the electronic structure of the solvated benzene radical anion in liquid ammonia complements the analysis of molecular geometry from our previous work and provides results that can be directly related to future experimental measurements of the system studied here.

The JT behavior of the solvated radical anion is analogous to that predicted for the idealized gas-phase species based on fundamental theory and symmetries. The electronic state and its associated spin density strongly correlate with the dynamic distortion of the molecular geometry as it undergoes motion through the almost flat pseudorotation valley. It, thus, turns out that the presence of the solvent is key to stabilize the studied system electronically but does not perturb it substantially from the perspective of the JT effect. This sets the stage for possible experimental studies of the consequences of the JT effect on the molecular and electronic structure of the benzene radical anion, which is not an option in the gas phase where the radical anion does not exhibit long-term stability. However, such experiments would have to rely on ultrafast techniques so that the individual JT structures are observed rather than their high-symmetry average.

We quantified the solvent-induced stability of the benzene radical anion using accurate and computationally demanding condensed-phase G_0W_0 calculations performed on thermal geometries sampled from a hybrid DFT-AIMD simulation. We estimated the binding energy of the excess electron to be -2.34 eV relative to the vacuum level, clearly showing that the excess electron represents a bound quantum state in solution. Moreover, the density of states obtained from such calculations predicts the complete valence electronic structure and, thus, provides a way to interpret future photoelectron spectroscopy measurements.

The present work showcases the descriptive power of accurate molecular simulations and detailed analysis of their outputs. We captured subtle quantum effects in both the spatial and energy

domains and obtained a detailed description of the solvated benzene radical anion in liquid ammonia and a prediction of its electronic density of states that complements our previous prediction of the vibrational density of states. The immediate next step lies in exploiting the synergy between the calculations reported here and future liquid photoelectron spectroscopy measurements. Referencing the results against the baseline of the solvated neutral benzene molecule further aids the interpretation of the anticipated experimental results. This combination has the potential to experimentally corroborate the solvent-induced stability of the benzene radical anion. One remaining issue is the computational description of the thermodynamic equilibrium between the benzene radical anion and solvated electrons that will provide additional insight into the experimentally observable chemical properties of the solvated benzene radical anion as well as an entryway to the theoretical exploration of the chemistry of the Birch reduction.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional data analysis details; additional results concerning the spin density dimensionality reduction, the evaluation of the GMM clustering, and the projected densities of states; and a video file visualizing the evolution of spin density over the pseudorotating molecular structure of the benzene radical anion.

ACKNOWLEDGMENTS

K.B. acknowledges funding from the IMPRS for Many Particle Systems in Structured Environments. This work was supported by the Project SVV 260586 of Charles University. This work was partially supported by the OP RDE project (Grant No. CZ.02.2.69/0.0/0.0/18_070/0010462), International Mobility of Researchers at Charles University (MSCA-IF II). P.J. acknowledges support from the European Regional Development Fund (Project ChemBioDrug, Grant No. CZ.02.1.01/0.0/0.0/16_019/0000729) and the Humboldt Research Award. This work was supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations Project "IT4Innovations National Supercomputing Center LM2015070." The authors would like to thank Hubert Beck and Tomáš Martinek for helpful comments on the manuscript.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

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

REFERENCES

- 1 E. Zurek, P. P. Edwards, and R. Hoffmann, "A molecular perspective on lithium-ammonia solutions," *Angew. Chem., Int. Ed.* **48**, 8198–8232 (2009).
- 2 M. Faubel, B. Steiner, and J. P. Toennies, "Photoelectron spectroscopy of liquid water, some alcohols, and pure nonane in free micro jets," *J. Chem. Phys.* **106**, 9013–9031 (1997).

- ³T. Buttersack, P. E. Mason, P. Jungwirth, H. C. Schewe, B. Winter, R. Seidel, R. S. McMullen, and S. E. Bradforth, "Deeply cooled and temperature controlled microjets: Liquid ammonia solutions released into vacuum for analysis by photoelectron spectroscopy," *Rev. Sci. Instrum.* **91**, 043101 (2020).
- ⁴T. Buttersack, P. E. Mason, R. S. McMullen, T. Martinek, K. Brezina, D. Hein, H. Ali, C. Kolbeck, C. Schewe, S. Malerz, B. Winter, R. Seidel, O. Marsalek, P. Jungwirth, and S. E. Bradforth, "Valence and core-level x-ray photoelectron spectroscopy of a liquid ammonia microjet," *J. Am. Chem. Soc.* **141**, 1838–1841 (2019).
- ⁵T. Buttersack, P. E. Mason, R. S. McMullen, H. C. Schewe, T. Martinek, K. Brezina, M. Crhan, A. Gomez, D. Hein, G. Wartner, R. Seidel, H. Ali, S. Thürmer, O. Marsalek, B. Winter, S. E. Bradforth, and P. Jungwirth, "Photoelectron spectra of alkali metal-ammonia microjets: From blue electrolyte to bronze metal," *Science* **368**, 1086–1091 (2020).
- ⁶A. J. Birch, "Reduction by dissolving metals," *Nature* **158**, 585 (1946).
- ⁷T. Shida and S. Iwata, "Electronic spectra of ion radicals and their molecular orbital interpretation. III. Aromatic hydrocarbons," *J. Am. Chem. Soc.* **95**, 3473–3483 (1973).
- ⁸T. R. Tuttle and S. I. Weissman, "Electron spin resonance spectra of the anions of benzene, toluene and the xylenes," *J. Am. Chem. Soc.* **80**, 5342–5344 (1958).
- ⁹J. C. Moore, C. Thornton, W. B. Collier, and J. P. Devlin, "Vibrational spectra, Jahn-Teller distortion, and the structure of the benzene radical anion," *J. Phys. Chem.* **85**, 350–354 (1981).
- ¹⁰L. Sanche and G. J. Schulz, "Electron transmission spectroscopy: Resonances in triatomic molecules and hydrocarbons," *J. Chem. Phys.* **58**, 479–493 (1973).
- ¹¹A. L. Hinde, D. Poppinger, and L. Radom, "Ab initio study of the benzene radical anion," *J. Am. Chem. Soc.* **100**, 4681–4685 (1978).
- ¹²A. P. Bazante, E. R. Davidson, and R. J. Bartlett, "The benzene radical anion: A computationally demanding prototype for aromatic anions," *J. Chem. Phys.* **142**, 204304 (2015).
- ¹³R. A. Marasas, T. Iyoda, and J. R. Miller, "Benzene radical ion in equilibrium with solvated electrons," *J. Phys. Chem. A* **107**, 2033–2038 (2003).
- ¹⁴M. C. M. O'Brien and C. C. Chancey, "The Jahn-Teller effect: An introduction and current review," *Am. J. Phys.* **61**, 688–697 (1993).
- ¹⁵I. B. Bersuker, *The Jahn-Teller Effect* (Cambridge University Press, 2006).
- ¹⁶K. Brezina, P. Jungwirth, and O. Marsalek, "Benzene radical anion in the context of the Birch reduction: When solvation is the key," *J. Phys. Chem. Lett.* **11**, 6032–6038 (2020); [arXiv:2005.03525](https://arxiv.org/abs/2005.03525).
- ¹⁷V. Kostal, K. Brezina, O. Marsalek, and P. Jungwirth, "Benzene radical anion microsolvated in ammonia clusters: Modeling the transition from an unbound resonance to a bound species," *J. Phys. Chem. A* **125**, 5811 (2021).
- ¹⁸A. Glielmo, B. E. Husic, A. Rodriguez, C. Clementi, F. Noé, and A. Laio, "Unsupervised learning methods for molecular simulation data," *Chem. Rev.* **121**, 9722–9758 (2021).
- ¹⁹P. Gasparotto and M. Ceriotti, "Recognizing molecular patterns by machine learning: An agnostic structural definition of the hydrogen bond," *J. Chem. Phys.* **141**, 174110 (2014); [arXiv:1410.5447](https://arxiv.org/abs/1410.5447).
- ²⁰M. Ceriotti, G. A. Tribello, and M. Parrinello, "Simplifying the representation of complex free-energy landscapes using sketch-map," *Proc. Natl. Acad. Sci. U. S. A.* **108**, 13023–13028 (2011).
- ²¹F. Hüser, T. Olsen, and K. S. Thygesen, "Quasiparticle GW calculations for solids, molecules, and two-dimensional materials," *Phys. Rev. B* **87**, 235132 (2013).
- ²²J. Wilhelm, M. Del Ben, and J. Hutter, "GW in the Gaussian and plane waves scheme with application to linear acenes," *J. Chem. Theory Comput.* **12**, 3623–3635 (2016).
- ²³J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, "CP2K: Atomistic simulations of condensed matter systems," *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **4**, 15–25 (2014).
- ²⁴M. Guidon, F. Schiffmann, J. Hutter, and J. VandeVondele, "Ab initio molecular dynamics using hybrid density functionals," *J. Chem. Phys.* **128**, 214104 (2008).
- ²⁵M. Guidon, J. Hutter, and J. VandeVondele, "Robust periodic Hartree-Fock exchange for large-scale simulations using Gaussian basis sets," *J. Chem. Theory Comput.* **5**, 3010–3021 (2009).
- ²⁶G. Lippert, J. Hutter, and M. Parrinello, "A hybrid Gaussian and plane wave density functional scheme," *Mol. Phys.* **92**, 477–488 (1997).
- ²⁷J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, "Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach," *Comput. Phys. Commun.* **167**, 103–128 (2005).
- ²⁸G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *J. Chem. Phys.* **126**, 014101 (2007); [arXiv:0803.4060](https://arxiv.org/abs/0803.4060).
- ²⁹J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximation made simple," *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- ³⁰Y. Zhang and W. Yang, "Comment on 'Generalized gradient approximation made simple,'" *Phys. Rev. Lett.* **80**, 890 (1998).
- ³¹C. Adamo and V. Barone, "Toward reliable density functional methods without adjustable parameters: The PBE0 model," *J. Chem. Phys.* **110**, 6158–6170 (1999).
- ³²L. Goerigk and S. Grimme, "A thorough benchmark of density functional methods for general main group thermochemistry, kinetics, and noncovalent interactions," *Phys. Chem. Chem. Phys.* **13**, 6670–6688 (2011).
- ³³J. VandeVondele and J. Hutter, "Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases," *J. Chem. Phys.* **127**, 114105 (2007).
- ³⁴S. Goedecker and M. Teter, "Separable dual-space Gaussian pseudopotentials," *Phys. Rev. B* **54**, 1703–1710 (1996); [arXiv:9512004 \[mtrl-th\]](https://arxiv.org/abs/9512004).
- ³⁵M. Guidon, J. Hutter, and J. VandeVondele, "Auxiliary density matrix methods for Hartree-Fock exchange calculations," *J. Chem. Theory Comput.* **6**, 2348–2364 (2010).
- ³⁶J. Wilhelm and J. Hutter, "Periodic GW calculations in the Gaussian and plane-waves scheme," *Phys. Rev. B* **95**, 235123 (2017).
- ³⁷F. Ambrosio, Z. Guo, and A. Pasquarello, "Absolute energy levels of liquid water," *J. Phys. Chem. Lett.* **9**, 3212–3216 (2018).
- ³⁸A. P. Gaiduk, T. A. Pham, M. Govoni, F. Paesani, and G. Galli, "Electron affinity of liquid water," *Nat. Commun.* **9**, 247 (2018).
- ³⁹P. Hunt, M. Sprik, and R. Vuilleumier, "Thermal versus electronic broadening in the density of states of liquid water," *Chem. Phys. Lett.* **376**, 68–74 (2003).
- ⁴⁰P. Fabian, M. Vincent, G. Olivier, B. Mathieu, P. Peter, W. Ron, V. Jake, and D. Cournapeau, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
- ⁴¹P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *J. Comput. Appl. Math.* **20**, 53–65 (1987).
- ⁴²R. Seidel, S. Thürmer, J. Moens, P. Geerlings, J. Blumberger, and B. Winter, "Valence photoemission spectra of aqueous $\text{Fe}^{2+/3+}$ and $[\text{Fe}(\text{CN})_6]^{4-/3-}$ and their interpretation by DFT calculations," *J. Phys. Chem. B* **115**, 11671–11677 (2011).

C.5 | Paper V

Brezina, K., Beck, H. and Marsalek, O.: Elucidating the nature of π -hydrogen bonding in liquid water and ammonia, *arXiv:2403.12937*, 2024

Elucidating the Nature of π -hydrogen Bonding in Liquid Water and Ammonia

Krystof Brezina, Hubert Beck, and Ondrej Marsalek*
 Charles University, Faculty of Mathematics and Physics, Ke Karlovu 3, 121 16 Prague 2,
 Czech Republic

(Dated: 20 March 2024)

Aromatic compounds form an unusual kind of hydrogen bond with water and ammonia molecules, known as the π -hydrogen bond. In this work, we report *ab initio* path integral molecular dynamics simulations enhanced by machine-learning potentials to study the structural, dynamical, and spectroscopic properties of solutions of benzene in liquid water and ammonia. Specifically, we model the spatial distribution functions of the solvents around the benzene molecule, establish the π -hydrogen bonding interaction as a prominent structural motive, and set up existence criteria to distinguish the π -hydrogen bonded configurations. These serve as a structural basis to calculate binding affinities of the solvent molecules in π -hydrogen bonds, identify an anticooperativity effect across the aromatic ring in water (but not ammonia), and estimate π -hydrogen bond lifetimes in both solvents. Finally, we model hydration-shell-resolved vibrational spectra to clearly identify the vibrational signature of this structural motif in our simulations. These decomposed spectra corroborate previous experimental findings for benzene in water, offer additional insights, and further emphasize the contrast between π -hydrogen bonds in water and in ammonia. Our simulations provide a comprehensive picture of the studied phenomenon and, at the same time, serve as a meaningful *ab initio* reference for an accurate description of π -hydrogen bonding using empirical force fields in more complex situations, such as the hydration of biological interfaces.

I. INTRODUCTION

A π -hydrogen bond is a hydrogen bonding interaction formed between a partially positively charged hydrogen atom of a donor molecule and the π -electron density of an aromatic acceptor. Such interactions have been observed with a variety of donor molecules including halogen hydrides,¹ aliphatic hydrogens,² alcohols,³ liquid ammonia⁴ and water.^{5,6} The interaction of aromatics with liquid water (H_2O) bears everyday-life pertinence due to its role in the solvation of aromatic residues at biological interfaces, structural stabilization inside proteins,⁷ as well as contamination of sea waters by oil-derived aromatic substances dissolved during oil spills.⁸ For instance, the prototypical aromatic compound benzene has a surprisingly high solubility of up to 20 mM in liquid water at ambient conditions.⁹ Chemically similar to water, liquid ammonia (NH_3) is a noticeably better solvent for aromatics¹⁰ and such mixtures are significant from the synthetic and industrial viewpoint as an entryway to Birch chemistry.¹¹

In general, water and ammonia in their liquid state are considered to be similar based on their molecules both being simple first-period, isoelectronic hydrides, as well as on the formation of hydrogen-bonded structures. However, because of the less electronegative nature of the nitrogen atom in ammonia, the hydrogen bond strength is substantially lowered when compared to water, which causes numerous remarkable differences between the two liquids.^{12,13} At the macroscopic level, this can be represented, for instance, by the large difference in their boiling points of 133 K despite their nearly identical molar masses.¹⁴ At the molecular level, the different hydrogen bond strength manifests in various properties of the two liquids, including both structural and dynamic quantities such as for example, infrared (IR) vibrational spec-

tra. Here, water exhibits a pronounced red shift of the OH stretch in comparison to the gas phase from approximately 3700 to 3400 cm^{-1} owing to the softening of the effective bond potential in the presence of the hydrogen bond acceptor.¹⁵ On the other hand, ammonia remains nearly unchanged in this aspect.¹⁶

To study the nature of π -hydrogen bonds in these protic solvents, limiting oneself to benzene as the prototypical aromatic species gives a general picture of the phenomenon. In this regard, a key experimental contribution reports the measurement of the multivariate-curve resolution Raman spectrum¹⁷ of benzene in water.⁶ In this study, the authors present the so-called solute-correlated (SC) spectrum, which represents a contribution to the overall Raman intensity that captures the vibrational character of the benzene solute as well as that of its immediate surroundings. This reveals an unexpected vibrational feature at approximately 3610 cm^{-1} , which is interpreted as an imprint of the effect of the water molecules engaging in π -hydrogen bonds on the solute vibrations.⁶ The particular frequency of the observed π -hydrogen bond peak suggests that water molecules participating in a π -hydrogen bond are vibrationally blue-shifted compared to water molecules in the liquid bulk and, as such, have more of a character of an isolated molecule since they are not fully involved in the water-water hydrogen bond network. Our recent work focusing on the corresponding liquid ammonia solutions using *ab initio* molecular dynamics (AIMD) with dispersion-corrected hybrid density functionals finds the equivalent π -hydrogen bond spectral feature in the IR SC spectrum of benzene obtained by exact decomposition of the dipole-dipole autocorrelation function Fourier transforms.⁴ Additionally, the blue shift of water has been given a degree of computational support at the level of individual OH bond frequency analysis in AIMD simula-

tions using the BLYP density functional^{18,19} to describe the condensed-phase system.^{20,21}

The solvation structure of benzene at the atomistic level has been studied using MD simulations^{22,23} at the PBE level of theory²⁴ in water and, as part of our previous work,⁴ using the revPBE0-D3 density functional^{24–27} in liquid ammonia. Both works come to a qualitatively comparable conclusion concerning the shape of the spatial distribution function (SDF) which suggests a substantial degree of similarity of structural aspects in both solvents. The reported SDFs have a geoid-like shape with pronounced maxima in the pole regions above and below the planar benzene molecule, where π -hydrogen bonds were found to occur, implying a major role of π -hydrogen bonds as a binding motive in benzene solvation.

All of the above-mentioned simulations employ a classical approximation for the atomic nuclei. However, a π -hydrogen bond, just like any other hydrogen bond, is an interaction that features the very light hydrogen nucleus and, therefore, can be expected to exhibit pronounced nuclear quantum effects (NQEs) even at ambient conditions.^{28,29} The overall change in the character of a particular hydrogen bond upon inclusion of quantum nuclei into the picture is derived from a fine balance between the so-called competing NQEs, which can have a two-way effect of strengthening already strong hydrogen bonds, but also weakening initially weak ones.^{28,30} This can have profound manifestations in the structure and dynamics of a hydrogen-bonded system at the local level but can also act at longer scales by affecting, for instance, the cooperativity of hydrogen bonds. As such, the multiple research works up to this point provide a qualitative picture of various isolated aspects of π -hydrogen bonds, but a rounded computational treatment with highly accurate electronic structure and the inclusion of quantum nuclei is missing. At the same time, having access to an accurate *ab initio* reference is important as π -hydrogen bonding is prominent in biological systems, but common force fields were found to be unreliable for this purpose.²³

To provide a comprehensive picture of the phenomenon, we perform thermostatted ring polymer molecular dynamics³¹ (TRPMD) simulations of dilute solutions of benzene in liquid water (at 300 K) and ammonia (at 223 K) using revPBE0-D3 hybrid density functional theory to model the interaction potential. Such simulation exploits the classical isomorphism of the imaginary-time path-integral formalism, where quantum particles are represented as harmonically-coupled ring polymers of replicas of the classical system. This allows for the use of classical AIMD techniques at a computational cost elevated linearly with the number of path-integral replicas to address NQEs of both structural and dynamical properties of studied systems. Performing these simulations naively using the traditional on-the-fly electronic structure calculation approach to AIMD is possible, but obtaining long enough trajectories for tightly converged statistical properties is impractical due to the increased cost of the path integral simulations. Therefore, we approach

this issue mainly by combining the traditional approach with Behler–Parrinello high-dimensional neural network potentials^{32,33} (NNPs) to fit the reference *ab initio* potential energy surface and perform 2 ns long simulations at a drastically reduced computational cost while maintaining the accuracy of the simulations. The shorter 250 ps *ab initio* trajectories themselves are used for the training of the NNPs, their validation, and the calculation of molecular dipole moments, which are not directly accessible through NNPs. A comprehensive summary of the employed computational methodology is available in Section S1 of the Supporting Information. With this methodology, we address the molecular-level solvation structure, discuss the π -hydrogen bond orientation, cooperativity, lifetimes, and, finally, vibrational spectroscopy, where we provide theoretical insight with immediate relevance for the interpretation of the underlying experimental work.

II. RESULTS AND DISCUSSION

In the following paragraphs, we present and discuss the results obtained in this work. In Section II A, we analyze the molecular solvation structure with focus on the geometry of the π -hydrogen bonds. This allows us to set up a time-dependent existence criterion to assess whether a solvent molecule is or is not participating in a π -hydrogen bond at a given instant, which is then used in Section II B to address the issue of π -hydrogen bond cooperativity across the two binding sites that the benzene molecule offers. In Section II C, we exploit the existence criterion again to look into the kinetics of π -hydrogen bond formation and, finally, we look into the problem of vibrational spectroscopy in Section II D by modeling the influence of the benzene solute on the vibrational spectra from *ab initio* data, finding an effect consistent with the above experimental observations and connecting its origin to vibrational features in the solvent part of the spectrum mainly due to π -hydrogen bonded solvent molecules.

A. Solvation structure

A global picture of the solvation of benzene by water and ammonia is encoded in the SDF $G(r, z, \varphi)$ of the solvent around the solute. The SDF is a 3D histogram of the positions of a selected atomic species described conveniently in this case in standard cylindrical coordinates r , z , and φ with the origin located in the solute center of mass and the plane of the solute aligned with $z = 0$. The SDFs for both solvents were obtained as statistical averages of heavy-atom positions over the 1 ns long C-NNP trajectories and the resulting data is shown in Figure 1, in panel A for water and in panel C for ammonia. Both SDFs share similar features and thus will be discussed concurrently. Their overall shape can be com-




FIG. 1. Spatial distribution functions of water (top row) and ammonia (bottom row) around the benzene solute. Panel A shows two different contours of the 3D SDF $G(r, z, \varphi)$ for the oxygen atoms of water at 95 \AA^{-3} (transparent) and 130 \AA^{-3} (solid). Panel B shows the same SDF integrated around the polar angle φ , zoomed in on the π -hydrogen bond cap regions and resolved for the oxygen atoms (red) and the hydrogens (gray). Panels C and D show the same data for the ammonia simulations with the nitrogen-related quantities colored blue.

pared to that of a rotational ellipsoid with the shorter axis in the z direction and they comprise two distinct regions: a set of hexagonally arranged stripe-like lobes around the solute plane that feature a hydrophobic solvation regime³⁴ (this issue will not be discussed in detail in this work) and two caps at the poles surrounding $r = 0$ which represent global maxima of the SDF and which will be identified as the regions in which π -hydrogen bonds are formed. Benzene thus simultaneously offers two binding sites for solvent molecules to form π -hydrogen bonds that we will denote for later use as site A (above the ring, $z > 0$) and site B (under the ring, $z < 0$). The only noticeable difference between the two SDFs is the less structured hydrophobic part of the ammonia SDF than that of the water SDF; the π -hydrogen bond caps look comparable. This suggests that the solvation structure of the benzene solute is similar in the two solvents, and the slightly stronger polarity of water molecules does not play a structure-defining role in this instance. Note that the data shown in panel C of Figure 1 is directly comparable to our previous data shown in Reference 4. However, unlike in our previous work, where the *ab initio* data required a smoothing procedure, the C-NNP data here is presented raw and originates from TRPMD simulations.

More insight into the nature of the π -hydrogen bond caps is achieved by calculating a similar SDF for the hydrogen atoms. For graphical clarity, we integrate the 3D

SDFs around the polar angle φ and show the resulting 2D SDFs $G(r, z)$ for both the heavy atoms and the hydrogens in panels B and D of Figure 1, respectively, for the two solvents and averaged over both binding sites A and B . Note that the shown extent focuses on the π -hydrogen bond caps only. Here, each colored, heavy-atom peak has a gray, hydrogen peak approximately 0.8 \AA under it, which suggests that an oriented interaction $\text{C}_6\text{H}_6 \cdots \text{HX}$ ($X = \text{O}, \text{N}$) is formed where the XH bond faces the solute along the z axis with its hydrogen pointing towards it. This interaction, prominently featured in the simulation data in both solvents, is the π -hydrogen bond.

Next, we want to explore how are the two independent 2D SDF clouds correlated, or, in other words, what possible orientations of π -hydrogen bonds occur in our simulations. For that purpose, we define an existence criterion for π -hydrogen bonds which picks all molecules that are located inside a cylinder that isolates the π -hydrogen bond caps from the rest of the SDFs. This criterion is a time series over the simulated trajectory and is defined individually for each n -th solvent molecule and k -th imaginary time replica as

$$h_{nk}^{(A)}(t) = \Theta[r_{nk}(t) - r_0] \Theta[z_{nk}(t) - z_0] = \begin{cases} 1 & \text{in a } \pi\text{-hydrogen bond} \\ 0 & \text{out of } \pi\text{-hydrogen bond} \end{cases} \quad (1)$$

for binding site A and identically as

$$h_{nk}^{(B)}(t) = \Theta[r_{nk}(t) - r_0] \Theta[z_{nk}(t) + z_0] \quad (2)$$

for binding site B . In both definitions, Θ is the step function, and r_{nk} and z_{nk} are the cylindrical coordinates of the heavy atom X of a given molecule and replica. The parameters r_0 and z_0 are selected based on the shape and extent of the individual SDFs: for water we set $r_0 = 0.12 \text{ nm}$ and $z_0 = 0.40 \text{ nm}$ and for ammonia $r_0 = 0.14 \text{ nm}$ and $z_0 = 0.41 \text{ nm}$. Next, we define a set of angles θ_j for each solvent molecule ($j = 1, 2$ for water and $j = 1, 2, 3$ for ammonia) which denote the angles between each HX bond and the z axis where a 0° angle is defined as the HX bond being collinear with z and with the hydrogen atom pointing towards the solute. All angles are then ordered in j by the vertical distance of the hydrogen atom from the solute so that θ_1 is the angle of the π -hydrogen bond. The distributions of these angles for replicas that comply with $h_{nk}^{(S)}(t)$ are shown in Figure 2 averaged over both binding sites $S = A, B$. For water (Figure 2, top panel), we observe a V-shaped distribution with a maximum in $\theta_1 \sim 10^\circ$ and $\theta_2 \sim 110^\circ$. This corresponds to the expected, albeit slightly tilted orientation of the π -hydrogen bond with the solvent molecule pointing one of its hydrogen towards the solute: the value of θ_2 here derives from the equilibrium bond angle in water of 104.5° (Figure 2, snapshot A). The tilt originates as a result of the contribution of the orientational entropy to the shown population: since there is only a single way to realize the perfectly oriented $\theta_1 = 0^\circ$ configuration, this is

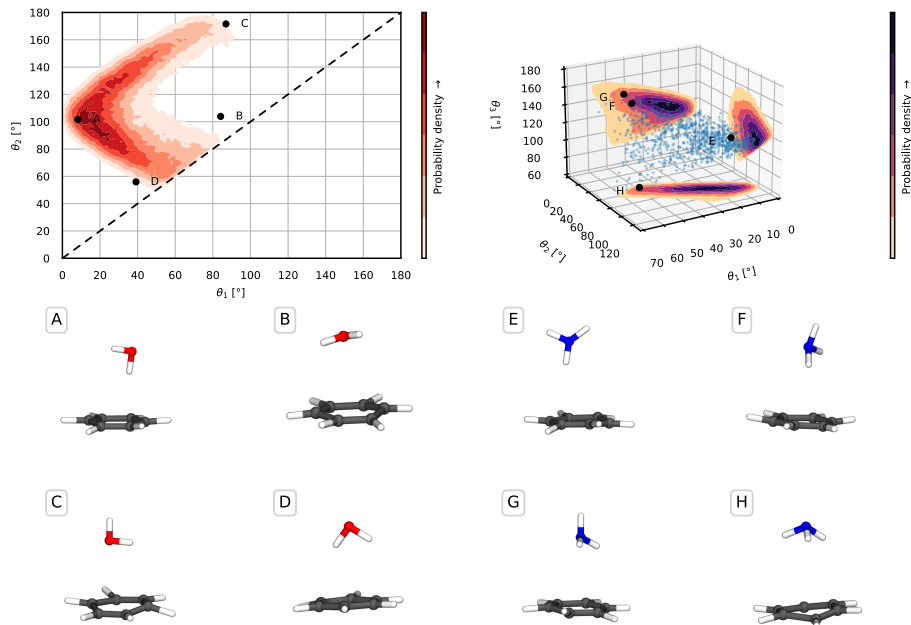


FIG. 2. Analysis of the orientation of solvent molecules around the benzene solute. Top left panel: the orientation of water molecules is described by two angles, θ_1 and θ_2 which describe the angle between the O–H bond and the solute normal with θ_1 describing the hydrogen that is closer to the solute and θ_2 the further one. Both angles reach 0° when the bond is oriented towards the solute with the hydrogen atom. Panels A, B, C, and D: snapshots of representative orientations of water molecules that occur in the cylinder defined by the criterion $h_i^{(k)}(t)$. Only panel A is clearly an oriented π -hydrogen bond. Top right panel: identical type of plot for ammonia (as for water on the left) with the addition of the angle θ_3 describing the orientation of the third hydrogen atom. The 3D distribution itself is shown as a scatter plot of a subset of points selected from the simulations. The 2D marginal distributions capturing the correlations of the orientation of two different N–H bonds are shown on the corresponding coordinate planes. Panels E, F, G, and H: representative configurations of solvent molecules in the π -hydrogen bond cylinder.

practically not seen during the simulation. The fact that it is indeed the energetically most favored configuration is shown by factoring out the entropic sine contributions as shown in Section S2 of the Supporting information. However, the extent of the distribution points to the fact that a richer pool of configurations is encountered in the simulations than just the idealistic π -hydrogen bond. In fact, the selection of molecules that comply with $h_{nk}^{(S)}(t)$ also includes molecules that are not in a π -hydrogen bond at all and just happen to be located in the region of space, but likely engaging in regular water–water hydrogen bond network (Figure 2, snapshots B, C and, as a borderline case, D). Therefore, based on this analysis, we update the selection criterion to also reflect the π -

hydrogen bond angle:

$$h_{nk}^{(S)}(t) = \Theta[\theta_{1,nk}(t) - \theta_{1,0}], \quad (3)$$

where the cutoff angle was set to $\theta_1 = 40^\circ$ for both solvents. A discussion along the same lines arises for ammonia. The 3D distribution in the corresponding angles is shown in the top right panel of Figure 2 and representative configurations in panels E, F, G, and H of the same Figure. Here, we remark that the preferred π -hydrogen bond configuration is also oriented vertically towards the solute with a slightly larger tilt of $\sim 20^\circ$.




FIG. 3. An illustration of the π -hydrogen bond existence time series. Snapshot: a single simulation frame showing the simulated periodic system consisting of water solvent and benzene solute for a single path-integral replica. Water molecules that fulfill the existence criterion are highlighted in the van der Waals representation and in green. Bottom panel: A graphical representation of the time dependence of $h_{n,k=0}^A(t)$ (Equation 1) over a 100 ps long segment of one of the simulated C-NNP path-integral trajectories and for a single replica. The molecular index n is represented by the different colors of the curves, showcasing the sampling of the exchange of molecules in the binding site A .

B. π -hydrogen bond cooperativity

As established above, the benzene molecule can engage in two π -hydrogen bonds at a time by binding a solvent molecule in sites A and B . Here, we ask what the populations of the zero-, one- and two-bonded configurations are in the simulated trajectories and to what extent the two coexisting hydrogen bonds cooperate or whether they can be considered independent.

To begin such analysis, we calculated the π -hydrogen bond existence criterion for the solute in our *ab initio* simulations (not C-NNP, see below) by summing over the time series of individual solvent molecules and sites

as

$$h_k(t) = \sum_{n,S} h_{nk}^{(S)}(t) = \begin{cases} 0 & \pi\text{-hydrogen bonds} \\ 1 & \pi\text{-hydrogen bond} \\ 2 & \pi\text{-hydrogen bonds.} \end{cases} \quad (4)$$

The probabilities for the individual species ($p(0)$ for 0 π -hydrogen bonds, $p(1)$ for a single π -hydrogen bond, and $p(2)$ for both binding sites occupied) were calculated as a simple average over the time variable and the P path integral replicas

$$p(N) = \frac{1}{PT} \int_0^T dt \sum_{k=1}^P \delta[h_k(t) - N] \quad (5)$$

for each $N = 0, 1, 2$. The obtained numbers $p(N)$ are shown in Figure 4 using the darker bars. In both solvents, the most likely configuration is that with only a single π -hydrogen bond populated at 51.9 % in water and 48.6 % in ammonia. This is followed by the configuration with no π -hydrogen bonds populated at 37.0 % in water and 33.3 % in ammonia, and the least probable configuration is the one with the maximal occupancy of 2 π -hydrogen bonds with 11.1 % in water and 18.0 % in ammonia. The probabilities $p(N)$ are related to the probabilities of occupying the specific binding sites A and B as follows (for this purpose, we borrow the notation A and B for referring to the sets of microstates corresponding to a solvent molecule being bound at the given binding site). By the symmetry of the problem, we clearly have $p(A) = p(B)$. Then, $p(2) = p(A \cap B) = p(A)p(A|B)$, $p(1) = p(A \cup B) - p(A \cap B) = 2p(A) - p(2)$, and, by completeness, $p(0) = 1 - p(1) - p(2)$.

The possible cooperativity was examined by comparing the observed populations to a model that assumes the independence of both binding sites. This model was constructed based on the premise that independent occurrences of binding to sites A and B (the corresponding probabilities are denoted as p_{ind}) comply with $p_{\text{ind}}(A|B) = p_{\text{ind}}(A)$, which implies that $p_{\text{ind}}(2) = p_{\text{ind}}^2(A)$. This leads to a single equation for two unknown independent probabilities, $p_{\text{ind}}(0)$ and $p_{\text{ind}}(1)$, giving us the freedom to fix one of these quantities. In this work, we choose to preserve the free-energy difference between states 0 and 1, which requires setting $p_{\text{ind}}(1)/p_{\text{ind}}(0) = p(1)/p(0)$. This allows us to estimate all the independent probabilities $p_{\text{ind}}(N)$. These results are shown in Figure 4 using the lighter shading next to the bars for the true observed probabilities $p(N)$. For benzene in water, comparing the real and independent probabilities is curious because it points to an anticooperativity effect between the two binding sites: the observed 11.1 % population of the double-bonded state is significantly lower compared to the value of 17.0 % predicted by the independent model. We explain this observation by considering the stronger impact a water molecule has on the solute, drawing the π -electron density of the aromatic ring to its side upon the formation of a π -hydrogen




FIG. 4. Probabilities of finding the zero-, one- and two-bonded configurations within our simulated trajectories. The darker bars show real, observed populations. The lighter bars then show probabilities that are calculated based on the premise of binding site independence. The error bars have been estimated using the block averaging method individually for each $\delta[h_k(t) - N]$ time series.

bond. This way, it weakens the other binding site and makes it less prone to forming a second π -hydrogen bond at the same time. In ammonia, this effect does not arise, and the independent model represents an excellent approximation of the real system. An illustration of this effect on electron density in gas-phase clusters consisting of a single benzene molecule and two π -hydrogen bonded solvent molecules is presented in Section S2 of the Supporting information.

Finally, we comment on the performance of C-NNP models in the simulation of π -hydrogen bond cooperativity. Our initial attempts to investigate the π -hydrogen bond cooperativity in C-NNP trajectories revealed that the anti-cooperativity effect for aqueous systems is reproduced poorly by the committee of Behler–Parrinello (BP) NNPs — it is considerably lower than in the original *ab initio* trajectories. π -hydrogen bond cooperativity is a very subtle effect and, hence, more difficult to accurately reproduce than other properties such as spatial distribution functions or vibrational spectra. A quantitative investigation into the accuracy of the employed models can be found in Section S2 the Supporting information and shows that the force errors are below the threshold typically needed to reproduce most standard properties. One challenge for the C-NNP description of π -hydrogen bonds is that the typical distance between the center of the benzene ring and the oxygen atom of a water molecule is between 3.0 and 3.5 Å (as shown in Figure 1). This suggests that in configurations with two simultaneous π -hydrogen bonds, the water molecule at the opposite site is at the very edge of the neighborhood cutoff radius of the BP NNP, which we set, as per usual, to $12 a_0$ (≈ 6.35 Å).




FIG. 5. Time ACFs of the π -hydrogen bond existence criterion. The top panel shows the ACF for both water and ammonia on the interval of lag time from 0 to 20 ps on a logarithmic scale. The bottom panel shows the integrated ACFs to estimate the correlation times as the long time limit of these ACFs — this is shown as the dashed lines.

Since the resolution of the atom-centered symmetry functions towards the edge of the cutoff radius is poor, an accurate prediction of this long-range effect is unlikely by these local MLPs. The more advanced equivariant message-passing graph NNPs such as NequIP can mitigate this downside, as the message-passing architecture extends the effective field of view beyond the set cutoff radius.³⁵ For this reason, we trained a NequIP model³⁶ for the classical benzene–water system and ran the same MD simulations with it as well. However, we observed no improvement in these trajectories, as the anti-cooperativity effect remains severely underestimated. Further research will be necessary to disentangle the reasons why both BP and equivariant NNPs fail to reproduce this effect quantitatively. One plausible explanation lies in the training data set, which has been selected using query by committee (QbC) with BP NNPs and was used to train both model architectures. It is possible that due to the short-range nature of BP NNPs, structures required for an accurate description of the anti-cooperativity effect have been left out of the training data set and that only a NequIP-specific QbC will lead to an NNP that reproduces all aspects of this system correctly.

C. π -hydrogen bond lifetimes

The time dependence of the solvent existence criteria $h_{nk}^{(S)}(t)$ (Equation 3) can be used to estimate the characteristic lifetime of the π -hydrogen bond. To calculate the time correlation function, we first evaluate the RPMD time-dependent observable³⁷ as

$$h_n^{(S)}(t) = \frac{1}{P} \sum_{k=1}^P h_{nk}^{(S)}(t) \quad (6)$$

by averaging over all the path-integral replicas for each molecule and each site separately. This quantity no longer has discrete values between 0 and 1 but can acquire a continuum of values between these two bounds: this can be interpreted as the quantum-delocalized solvent molecule only partially fulfilling the π -hydrogen bond existence criterion due to its replicas entering and leaving the relevant region. The autocorrelation function is then

$$c(\tau) = \frac{1}{2N_{\text{sol}}} \sum_{n,S} \frac{\langle h_n^{(S)}(t_0) h_n^{(S)}(t_0 + \tau) \rangle}{\langle [h_n^{(S)}(t_0)]^2 \rangle} \quad (7)$$

and, as such, gives the probability that a π -hydrogen bond still exists after a delay τ if it initially existed at time t_0 . Such autocorrelation functions for both solvents are shown in Figure 5 over the range of 20 ps and feature a typical shape with a sharper decrease at very short delays followed by a less abrupt, long, and monotonously decaying exponential regime. Integration of these functions over the time variable reveals the time scale at which these correlations decay, which can be interpreted as the characteristic lifetime of π -hydrogen bonds. Interestingly, these lifetimes are comparable in both solvents: for water, we obtained an autocorrelation time of approximately 1.8 ps, while for ammonia, a just slightly shorter value of 1.7 ps. We propose that this agreement is coincidental as the solutions are at different temperatures and have different viscosities. To set the ground for comparison, we additionally calculated the hydrogen bond lifetimes for the neat solvents as well using preexisting trajectories: the lifetime of water-water hydrogen bonds is approximately 3.9 ps and ammonia-ammonia hydrogen bonds approximately 1.4 ps. This indirectly suggests that solvent-solvent hydrogen bonds are stronger than π -hydrogen bonds in water but comparable or even ever so slightly weaker in ammonia.

D. Vibrational spectroscopy

Finally, we turn our attention to the simulation of vibrational spectra of both solutions in order to relate to the previous experimental findings⁶ and to gain insight into the influence of the aromatic solute on the solvent. In this regard, we present a qualitative perspective based on vibrational densities of states (VDOS) and infrared (IR) intensities. VDOS is defined through the Fourier transform of the autocorrelation function of atomic velocities \mathbf{v}_i ,

$$C_{vv,i}(t) = \langle \mathbf{v}_i(t_0) \cdot \mathbf{v}_i(t_0 + t) \rangle_{t_0}, \quad (8)$$

$$I_{\text{VDOS}}(\omega) \propto \sum_i \int dt e^{-i\omega t} C_{vv,i}(t),$$

where we write the total VDOS as a sum of atomic contributions. While it has no experimentally measurable counterpart, it is the most direct and local probe into

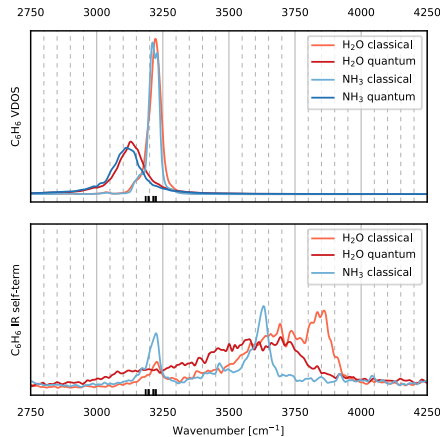


FIG. 6. Vibrational characterization of the benzene solute in liquid water (light red for classical and dark red for TRPMD trajectories) and ammonia (light blue for classical and dark blue for TRPMD trajectories). Top panel: Vibrational density of states of the solute carbon and hydrogen atoms combined. Bottom panel: IR absorption self-term due to the benzene molecular dipole moment. Harmonic frequencies of gas-phase benzene calculated at the revPBE0-D3 hybrid DFT level are shown as black ticks at the bottom of each panel.

the mechanical vibrations of molecules. Similarly, IR intensities, which correspond to experimentally measurable IR signals, can be related to the Fourier transform of the autocorrelation function of the total dipole moment \mathbf{M} of the studied system,

$$C_{MM}(t) = \langle \mathbf{M}(t_0) \cdot \mathbf{M}(t_0 + t) \rangle_{t_0}, \quad (9)$$

$$I_{\text{IR}}(\omega) \propto \omega^2 \int dt e^{-i\omega t} C_{MM}(t).$$

Given a partitioning method to decompose the total dipole to molecular contributions as $\mathbf{M} = \sum_n \boldsymbol{\mu}_n$, the total IR intensity can also be decomposed to contributions of pairs of molecules as

$$I_{\text{IR}}(\omega) = \omega^2 \sum_{n,m} \int dt e^{-i\omega t} \langle \boldsymbol{\mu}_n(t_0) \cdot \boldsymbol{\mu}_m(t_0 + t) \rangle_{t_0}, \quad (10)$$

where not only the molecular self-terms ($n = m$), but also cross-correlation terms ($n \neq m$), contribute to the total intensity. In the following analysis, the quantum vibrational spectra are considered to be the primary outcome, but we will frequently relate to the corresponding classical simulations, owing to the insight they provide into the interpretation of the presented spectral curves.

To motivate the discussion, we first inspect the vibrational character of the solute in both solvents. The solute component of the VDOS in the region of benzene CH stretch modes between 2750 and 4250 cm^{-1} is shown in the top panel of Figure 6 for water as well as ammonia. Classically, the CH vibrations present in the VDOS as a relatively narrow peak located just above 3200 cm^{-1} with an almost identical shape in both systems. Its position closely corresponds to the position of gas-phase harmonic frequencies of the isolated benzene molecule (Figure 6), suggesting that the vibrational motion of the solute is not strongly affected by its solvation. The TRPMD VDOS exhibits a typical vibrational manifestation of NQEs, where the quantum spectrum is, compared to the classical one, broader and redshifted — in this case by approximately 100 cm^{-1} . There is no substantial difference between this vibrational motion of the benzene molecule in liquid water or ammonia and no higher-frequency vibrations of the solute than the CH stretches are present. The bottom panel of Figure 6 shows the component of the IR absorption spectrum due to the solute-solute dipole moment correlations, with the aim of isolating and inspecting the solute-only contribution to the total IR spectrum. The classical solute-solute self-terms (light blue and light red) contain the expected CH vibration peak at approximately 3200 cm^{-1} that is consistent with the VDOS peak as well as the harmonic frequencies. However, the spectra now contain new additional features at higher frequencies, too. In water, these form a broad region of vibrational intensity between 3250 – 4000 cm^{-1} in the classical spectrum that can be interpreted as a double-peak structure with maxima located at 3700 cm^{-1} and 3850 cm^{-1} . With a quantum description, we can once again see NQEs in the form of a broadening and redshift of the spectral features to the point where the CH peak nearly fuses with the new features, almost leading to a single broad structure, where the original peaks are basically unrecognizable. In contrast, in ammonia, the additional feature presents as a single sharp peak at 3625 cm^{-1} under the classical approximation; the TRPMD IR spectrum of the ammonia solution was not calculated for reasons explained below. Since these new features do not correspond to any peak in the solute VDOS, they cannot be ascribed to the vibrational motion of the atoms of the solute. Rather, they must arise as an effect due to the solvent, which, through its own vibration, affects the dipole moment of the solute, leading to this feature in the solute-solute self-term spectrum. This is consistent with the interpretation of the Raman experiment, which reports an equivalent feature in the SC spectrum and assigns it to the water molecule that forms a π -hydrogen bond.

However, the results presented until this point do not clarify whether the solvent-induced features in the solute-only IR spectrum are caused directly by solvent molecules that are in a π -hydrogen bond with the solute, as implied by the interpretation of the experiment. To gain insight into the origin of the effect, we time-resolve the correla-

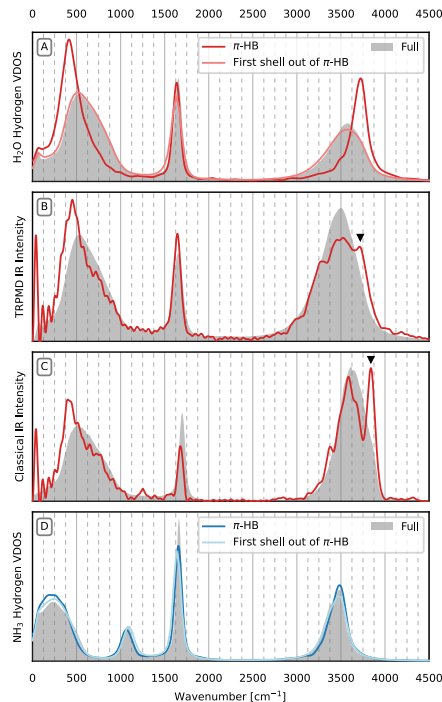


FIG. 7. Vibrational characteristics of solvent molecules engaging in a π -hydrogen bond. Panel A: TRPMD VDOS of water hydrogen atoms for the full system (gray shading) and for selected hydrogens that engage in a π -hydrogen bond. In addition, the same quantity is shown for hydrogen atoms in the first solvent shell that do not engage in a π -hydrogen bond. Panel B: TRPMD IR water self-term, again shown for the full system (gray shading) and the corresponding π -hydrogen bond subselection (red). The black triangle marks the relevant blueshifted feature. Note that the oscillation at the low-frequency end of the spectrum is a numerical artifact related to the practical execution of the Fourier transforms for the filtered spectra (see Section S1 of the Supporting information for details). Panel C: same data as in panel B, but obtained using classical mechanics. Note the pronounced split-peak feature in the OH stretch peak for molecules in a π -hydrogen bond. Panel D: same data as in Panel A, but for the liquid ammonia simulations.

tion functions that contribute to the *solvent* VDOS and the whole IR spectrum and correlate them with the aforementioned π -hydrogen bond existence criteria. This allows us to isolate the vibrational signature of molecules that are in a π -hydrogen bond at a given moment: this is

possible given the lifetime of a π -hydrogen bond is much larger than the vibrational period. The key results of this analysis for the simulations with quantum nuclei are summarized in Figure 7.

Let us first turn our attention to the aqueous solution. The overall (unresolved) solvent hydrogen TRPMD VDOS is shown in panel A of Figure 7 as the gray-shaded reference; it has the typical three-peak structure routinely observed for liquid water.³⁸ Restricting the presented VDOS to those solvent hydrogen atoms actively participating in a π -hydrogen bond shows a contribution that differs strongly from the total spectrum (Figure 7, panel A, dark red). Notably, the maximum of the stretch band is blueshifted from 3575 cm^{-1} to 3750 cm^{-1} , suggesting that the vibration in a π -hydrogen bond is stiffer than those in the bulk water-water hydrogen bond network. In turn, this implies that the π -hydrogen bond vibration is more like a dangling or isolated OH bond, the vibration of which is also not softened by the presence of a relatively strong hydrogen bonding partner. In contrast, the terahertz peak is shifted to lower frequencies, which is again consistent with the π -hydrogen bond being weaker than a water-water hydrogen bond and thus hindering rotation of the water molecule less. To verify that the described VDOS shifts are specific for π -hydrogen bonds rather than a general effect of molecules in proximity to the solute, we also inspect the VDOS of solvent molecules that belong to the first solvent shell in terms of distance but are excluded by the π -hydrogen bond existence criteria. We find that such molecules possess a vibrational signature essentially identical to the overall VDOS with no observable shifts (Figure 7, panel A, light red).

To explain these solvent-induced features in the solute-only IR component, we now turn our attention to the solvent-solute self-term IR spectra. The overall (unresolved) TRPMD water-water self-term IR spectrum is shown in panel B of Figure 7 as a starting point, again in gray. In this case, the π -hydrogen bond resolved counterpart brings around a more complicated structure of the OH stretch band (Figure 7, panel B, dark red) because, unlike VDOS, IR spectra cannot be resolved all the way to individual atoms and the effect of both OH-bond stretches is imprinted in the observed spectral feature. As such, the faster vibration of the OH bond that participates in a π -hydrogen bond manifests as a high-frequency shoulder just below 3750 cm^{-1} on the OH stretch peak (marked by a black arrow), while the rest of the dipole moment, including the second OH bond that points away from the solute and participates in the regular bulk hydrogen bond network, leads to a maximum at 3500 cm^{-1} . The same split of the stretch peak due to different chemical environments of the two OH bonds is better visible in the classical IR spectrum (Figure 7, panel C, dark red), where a pronounced double-peak feature arises with maxima located at 3600 cm^{-1} and 3850 cm^{-1} . The frequencies of the observed double-peaks in both the classical and TRPMD case of the π -hydrogen bond resolved solvent-only IR spectrum correspond to those observed in the π -

hydrogen bond resolved solute-only IR spectrum shown in Figure 6. This suggests that the effect observed in the solute-only IR spectrum is caused by the vibrating dipole of the solvent molecule inducing a dipole on the benzene solute at the frequency of this vibration. Thus, the observed phenomenon in the solute-only spectrum is not specifically due to the π -hydrogen bond, but rather represents a general effect due to solute-solvent correlations. However, since the interactions responsible for these correlations decay with distance and the π -hydrogen bonded molecule is the closest a solvent molecule can get to the solute (as demonstrated in Figure 1), it is expected that the π -hydrogen bonded molecules have the strongest contribution to the effect; this interpretation is further supported by the fact that we found the imprint of the vibration of the distal OH bond, which does not directly participate in the formation of the π -hydrogen bond, in the solute-only IR spectrum as well. The extent of participation of π -hydrogen bonded molecules in forming the solute-only spectral features could be directly tested, in principle, by calculating the solute-only spectrum in the absence of π -hydrogen bonds. However, this is not technically possible with the present direct MD simulations: given the vibrational period, the characteristic length of the window where no π -hydrogen bonds exist is too short to be able to get a resolved spectrum corresponding to this state. Simulations with the addition of a suitable biasing potential would be needed to obtain continuous trajectories free of π -hydrogen bonds.

The ammonia TRPMD VDOS (Figure 7, panel D) has one extra peak compared to the water spectra at around 1100 cm^{-1} due to the umbrella vibration. Interestingly, no frequency shifts are observed in the ammonia solutions after resolving the total solvent spectrum into the π -hydrogen bond and non- π -hydrogen bond first-shell contributions: all spectra are essentially identical. This result has a simple physical explanation: unlike in water, where the bulk hydrogen bond structure is strong, in ammonia, it is nearly non-existent, and, as such, the molecule in a π -hydrogen bond does not differ substantially from the other solvent molecules in the bulk. For this reason, the ammonia counterpart of the solvent-induced peak in the solute-only spectrum (Figure 6, bottom panel in blue) is represented by a single narrow peak and not a complicated double-peak structure as it is in water.

In summary, the vibrational character of molecules engaging in π -hydrogen bonds is different in both solvents, but this difference is due to the difference in the hydrogen bond strength within the respective solvent bulks, not a difference in the nature of the π -hydrogen bonds themselves.

III. CONCLUSIONS

The reported TRPMD simulations of solutions of benzene in liquid water and ammonia provide a detailed com-

putational perspective on the phenomenon of π -hydrogen bonding. To this end, we relied on the combination of accurate *ab initio* electronic structure with high computational efficiency in the form of C-NNPs to converge the statistical properties of interest. In this work, we found clear evidence of π -hydrogen bonding in the solvation structure of benzene in both solvents and characterized it in several complementary ways. We discussed the orientational flexibility of π -hydrogen bonded solvent molecules, described the strength of the interaction by its equilibrium population, showed evidence of its anticooperativity in water, and determined its lifetimes in both solvents. Furthermore, we modeled the VDOS and IR spectra of these systems and, employing a spatial and temporal decomposition that is only available in a computational approach, found and explained imprints of the π -hydrogen bonding interaction that are consistent with previous experimental findings in water. Our simulations point to the fact that π -hydrogen bonding is a prominent interaction present in the studied systems most of the time. This suggests relevance for situations where aromatic species come into contact with water, such as in the solvation of biological residues; some authors even suggest a possible role of π -hydrogen bonding in biological scenarios such as signaling.⁶ An accurate description by empirical force fields would thus be instrumental, yet it is currently not available.²³ Our work provides insight as well as an extensive *ab initio* reference for the future development of adequate force fields.

The question of why π -hydrogen bond anticooperativity presents a challenge for machine learning potentials, as identified by our *ab initio* simulations, remains open. Rectifying this issue will provide insight into the functionality and shortcomings of these potentials and will enable accessing other hard-to-converge aspects of π -hydrogen bonding, such as the further details of the thermodynamics of the formation of this bond. An additional level of improvement will be achieved by employing machine-learned models for predicting molecular dipoles and polarizabilities,³⁹ which will give access to better converged IR spectra and open the doors to computationally accessible Raman spectroscopy as well.

ACKNOWLEDGMENTS

O.M. thanks Dor Ben-Amotz for stimulating discussions of vibrational spectroscopy of π -hydrogen bonds in solution. The authors acknowledge support from the Czech Science Foundation, project No. 21-27987S. This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90254).

¹A. R. Nekoei and M. Vatanparast, *Physical Chemistry Chemical Physics* **21**, 623 (2019).

²M. Nishio, *Physical Chemistry Chemical Physics* **13**, 13873 (2011).

³M. Mons, E. G. Robertson, and J. P. Simons, *Journal of Physical Chemistry A* **104**, 1430 (2000).

⁴K. Brezina, P. Jungwirth, and O. Marsalek, *Journal of Physical Chemistry Letters* **11**, 6032 (2020).

⁵S. Suzuki, P. G. Green, R. E. Bumgarner, S. Dasgupta, W. A. Goddard, and G. A. Blake, *Science* **257**, 942 (1992).

⁶K. P. Gierszal, J. G. Davis, M. D. Hands, D. S. Wilcox, L. V. Slipchenko, and D. Ben-Amotz, *Journal of Physical Chemistry Letters* **2**, 2930 (2011).

⁷T. Steiner and G. Koellner, *Journal of molecular biology* **305**, 535 (2001).

⁸J. Gros, D. Nabi, B. Würz, L. Y. Wick, C. P. Brussaard, J. Huisman, J. R. V. D. Meer, C. M. Reddy, and J. S. Arey, *Environmental Science and Technology* **48**, 9400 (2014).

⁹D. S. Arnold, C. A. Plank, E. E. Erickson, and F. P. Pike, *Journal of Chemical and Engineering Data* **3**, 253 (1958).

¹⁰H. C. Schewe, K. Brezina, V. Kostal, P. E. Mason, T. Buttersack, D. M. Stemer, R. Seidel, W. Quevedo, F. Trinter, B. Winter, and P. Jungwirth, *Journal of Physical Chemistry B* **126**, 229 (2022).

¹¹A. J. Birch, *Pure and Applied Chemistry* **68**, 553 (1996).

¹²A. D. Boese, A. Chandra, J. M. Martin, and D. Marx, *The Journal of Chemical Physics* **119**, 5965 (2003).

¹³A. Krishnamoorthy, K. I. Nomura, N. Baradwaj, K. Shimamura, R. Ma, S. Fukushima, F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashishta, *Journal of Physical Chemistry Letters* **13**, 7051 (2022).

¹⁴S. Glazier, N. Marano, and L. Eisen, *Journal of Chemical Education* **87**, 1336 (2010).

¹⁵E. M. Kosower and G. Borz, *RSC Advances* **1**, 1506 (2011).

¹⁶M. N. Slipchenko, B. G. Sartakov, and A. F. Vilesov, *Journal of Chemical Physics* **128**, 134509 (2008).

¹⁷P. Perera, M. Wyche, Y. Loethen, and D. Ben-Amotz, *Journal of the American Chemical Society* **130**, 4576 (2008).

¹⁸A. D. Becke, *Physical Review A* **38**, 3008 (1988).

¹⁹C. Lee, W. Yang, and R. G. Parr, *Physical Review B* **37**, 785 (1988).

²⁰B. S. Mallik, A. Semparathi, and A. Chandra, *Journal of Physical Chemistry A* **112**, 5104 (2008).

²¹A. Choudhary and A. Chandra, *Journal of Physical Chemistry B* **119**, 8600 (2015).

²²M. Allesch, E. Schwegler, and G. Galli, *Journal of Physical Chemistry B* **111**, 1081 (2007).

²³M. Allesch, F. C. Lightstone, E. Schwegler, and G. Galli, *Journal of Chemical Physics* **128**, 14501 (2008).

²⁴J. P. Perdew, K. Burke, and M. Ernzerhof, *Physical Review Letters* **77**, 3865 (1996).

²⁵Y. Zhang and W. Yang, *Physical Review Letters* **80**, 890 (1998).

²⁶C. Adamo and V. Barone, *Journal of Chemical Physics* **110**, 6158 (1999).

²⁷L. Goerigk and S. Grimme, *Physical Chemistry Chemical Physics* **13**, 6670 (2011).

²⁸X. Z. Li, B. Walker, and A. Michaelides, *Proceedings of the National Academy of Sciences of the United States of America* **108**, 6369 (2011).

²⁹T. E. Markland and M. Ceriotti, *Nature Reviews Chemistry* **2018**, 2:3 2, 1 (2018).

³⁰S. Habershon, T. E. Markland, and D. E. Manolopoulos, *Journal of Chemical Physics* **131** (2009), 10.1063/1.3167790/938289.

³¹M. Rossi, M. Ceriotti, and D. E. Manolopoulos, *Journal of Chemical Physics* **140** (2014), 10.1063/1.4883861/73314.

³²J. Behler and M. Parrinello, *Physical Review Letters* **98**, 146401 (2007).

³³C. Schran, K. Brezina, and O. Marsalek, *The Journal of Chemical Physics* **153**, 104105 (2020).

³⁴D. Ben-Amotz, *Annual Review of Physical Chemistry* **67**, 617 (2016).

³⁵K. T. Schütt, O. T. Unke, and M. Gastegger, "Equivariant message passing for the prediction of tensorial properties and molecular spectra," (2021).

³⁶S. Batzner, A. Musaelian, L. Sun, M. Geiger, J. P. Mailoa, M. Ko-

- rnbluth, N. Molinari, T. E. Smidt, and B. Kozinsky, *Nature Communications* **13** (2022), [10.1038/s41467-022-29939-5](https://doi.org/10.1038/s41467-022-29939-5).
- ³⁷I. R. Craig and D. E. Manolopoulos, *The Journal of Chemical Physics* **121**, 3368 (2004).
- ³⁸T. Morawietz, O. Marsalek, S. R. Pattenau, L. M. Strecker, D. Ben-Amotz, and T. E. Markland, *Journal of Physical Chemistry Letters* **9**, 851 (2018).
- ³⁹A. M. Lewis, P. Lazzaroni, and M. Rossi, *Journal of Chemical Physics* **159**, 14103 (2023).